Essays in Information Economics

DISSERTATION

zur Erlangung des akademischen Grades
doctor rerum politicarum
(Doktor der Wirtschaftswissenschaft)

eingereicht an der
Wirtschaftswissenschaftlichen Fakultät
der Humboldt-Universität zu Berlin

von
Dipl.-Math. Jonas von Wangenheim.

Präsidentin der Humboldt-Universität zu Berlin:
Prof. Dr.-Ing. Dr. Sabine Kunst

Dekan der Wirtschaftswissenschaftlichen Fakultät:
Prof. Dr. Christian D. Schade

Gutachter:
1. Prof. Dr. Roland Strausz
2. Prof. Dr. Paul Heidhues

Eingereicht am: 22.03.2018
Tag des Kolloquiums: 18.07.2018

Acknowledgements

First and foremost I would like to express my gratitude to my supervisors Paul Heidhues and Roland Strausz for their invaluable guidance throughout the last years. Their advice and support only made this work possible.

Moreover, I am very grateful to Daniel Krähmer, Francoise Forges, Helmut Bester, Radosveta Ivanova-Stenzel, and Ran Spiegler for their important suggestions.

I would like to thank many fellow PhD students and postdocs in Berlin for their valuable feedback, in particular Andreas Asseyer, Johannes Johnen, Lucien Frys, Philipp Dimakopoulos, Matthias Lang, Thomas Mettral, Sebastian Schweighofer-Kodritsch, Vincent Meisner, and Yves Breitmoser. I thank Nathan Bechhofer for suggestions on the Introductions.

For their administrative support, I thank Regine Hallmann, Viviana Lalli, Myrna Selling and Sandra Uzman.

I am deeply grateful to my wife Jana for always supporting me and patiently listening to my newest research ideas.

For supporting my time at the Toulouse School of Economics, I thank to Bruno Jullien and Thomas Mariotti.

Financial support from the German Research Foundation is gratefully acknowledged.

Inhaltsverzeichnis

Introduction **2**

1 Consumer-Optimal Information Design **6**

1.1 Introduction . . . 6
1.2 Related Literature . . . 8
1.3 The model . . . 10
1.4 The Uniform Case . . . 11
1.5 The General Case . . . 16
1.6 The Limits of Surplus Distribution . . . 18
1.7 Conclusion . . . 21
1.8 Appendix to Chapter 1 . . . 22

2 English vs. Vickrey Auctions with Loss Averse Bidders **34**

2.1 Introduction . . . 34
2.2 Related Literature . . . 37
2.3 The Vickrey Auction . . . 39
2.3.1 Auction Rules . . . 39
2.3.2 Preferences . . . 39
2.3.3 Equilibrium Concept . . . 40
2.3.4 The Equilibrium . . . 42
2.4 The English Auction . . . 43
2.4.1 The Model . . . 43
2.4.2 Analysis . . . 46
2.5 Revenue Comparison . . . 54
2.6 Extensions and Robustness . . . 56
2.6.1 Loss Aversion in the Money Dimensions . . . 56
2.6.2 False Beliefs or Heterogeneous Preferences . . . 59
2.6.3 Reserve Price . . . 59
2.6.4 Generalization to n bidders . . . 60

2.7 Conclusion 63
2.8 Appendix to Chapter 2 64

3 Persuasion against Self-Control Problems **86**
3.1 Introduction 86
3.2 Related Literature 88
3.3 The Model 90
3.4 Analysis 93
3.5 Conclusion 101
3.6 Appendix 102

Bibliography **112**

Abbildungsverzeichnis

1.1 Partition of the optimal signal structure (stylized) 17
1.2 All potential pairs of surplus division 18
1.3 The signal to induce $(u_B, u_S) = (BS, SS)$ 20
1.4 Signal $\tau(\theta)$ for the distribution in Figure 1.3 (stylized) 21

2.1 Updating in the English Auction 48
2.2 Total Expected Loss for $F \sim U[0,1]$ 51
2.3 Ex-ante strategy and time-cons. constraint, $G(\theta) \sim U[0,1], \eta = 0.3, \lambda = 4$ 52
2.4 Equilibria in the English auction, $G(\theta) \sim U[0,1]$, $\eta = 0.3$, $\lambda = 4$ . . . 54
2.5 Equilibrium bidding function, $G(\theta) \sim U[0,1]$, $\eta = 0.3$, $\lambda = 4$ 55
2.6 Equilibrium bidding functions, $G(\theta) \sim U[0,1]$, $\eta = 0.4$, $\lambda = 3$ 58

3.1 Welfare loss under complete information 94
3.2 Gain and Loss of a Deviation from Self-t's Perspective 100

Introduction

Abstract (Deutsch)

Diese Dissertation besteht aus drei unabhängigen Artikeln aus dem Forschungsbereich der Informationsökonomik. Das verbindene Motiv aller Artikel ist die zweischneidige Rolle von Information. In Kontrast zur klassischen Entscheidungstheorie, in der mehr Information Individuen niemals schlechter stellen kann, analysiere ich drei Marktsituationen, in denen mehr Information schädlich für Konsumenten sein kann.

Abstract

This dissertation comprises three independent chapters in the field of information economics. The recurrent theme of all three chapters is the ambiguous role of information: While in standard decision theory additional information enables individuals to weakly increase utility through making better choices, I analyze three different environments in which more information to consumers may actually be detrimental to consumer utility.

Introduction

Chapter 1 is concerned with a monopolistic market, where consumers only learn their valuation for the product over time. As an example one may consider a consumer buying on an online platform. While the consumer may have some ex-ante product information prior to trade, many consumer-specific details of the product remain vague until the consumer receives the good and can properly inspect it.

I analyze the role of the ex-ante information—the information obtained prior to the purchasing decision—for consumer surplus in such environments. The effect of this ex-ante information is ambiguous as it impacts not only the decision of consumers, but also the price set by the profit maximizing seller. I identify the optimal amount of ex-ante information from the consumers' perspective, and show that it keeps consumers to some extent uninformed about their actual value for the product. It induces efficient trade and leaves the seller with the same profit as she would receive under full information.

Further, I emphasize the role of ex-ante information by characterizing, for any given distribution of buyer valuations, all possible divisions of buyer surplus and seller surplus arising for different ex-ante information structures. I show that any surplus pair which satisfies some basic feasibility constraints may arise as a market outcome.

In Chapter 2, I study the impact of dynamically revealing information in auctions with bidders who are expectation-based loss averse. As experimental evidence suggests, many individuals evaluate losses with respect to their expectations stronger than gains. I analyze the effect of this expectation-based loss aversion on bidding behavior in the second-price (Vickrey) auction and in the ascending-clock (English) auction.

While in both auction formats the highest bidder ultimately wins the auction and pays the price of the second highest bid, the two auctions differ in the way information is revealed during the auction. In the Vickrey auction bidders submit one sealed bid, and only learn more after the auction is resolved, whereas in the English auction prices are incrementally increased and all bidders continuously observe which of the opponents are still willing to buy at the current price. Under standard preferences, this information about remaining bidders is immaterial whenever bidders' private valuations of the auctioned object do not depend on other bidders' valuations: In both formats it is optimal for bidders to respectively submit or respectively bid up to their private valuation, and thus the seller receives the same revenue. This result is broadly known as revenue equivalence.

In my work, I show that the revenue equivalence breaks down if bidders are expectation-based loss averse. Intuitively, while information revelation during the auction does not affect the intrinsic value of the good, it changes the expectations of winning. The stronger loss-averse bidders believe they are winning the auction, the more they are threatened by a potential loss, and the more they are willing pay in order to reduce that risk. Consequently, the optimal bidding strategy does not only depend on the value of the good, but also on the evolution of expectation during the auction process, which differs between the two auction formats.

I show that in equilibrium there is a clear revenue ranking between the two auction formats: The Vickrey auction always yields strictly higher revenue than the English auction, a result in line with most of the experimental literature.

For the dynamic model of reference-dependent preferences it is assumed that any change in expectations instantaneously gives rise to gains and losses in utility which can be interpreted as elation and disappointment. As losses weigh stronger than gains any fluctuation in beliefs will create negative utility in expectation. For the English auction this implies that the permanent arrival of new information is in expectation detrimental for bidders' utility. However, the source of disutility from detailed information is quite different from that in Chapter 1. In Chapter 1 the ambiguous role of information arises from the strategic interaction with other market players, while for dynamically loss-averse consumers it is inherently given by their individual preferences.

In Chapter 3, I analyze the optimal information policy aimed at helping consumers who have self-control problems. Such consumers strongly value instantaneous gratification compared to their long-term utility. These preferences lead to time-inconsistent consumption decisions: Consumers may smoke, drink, or procrastinate due to their present bias, even though they would have liked to commit to abstention or hard work beforehand.

I contribute to the literature, which analyzes strategic ignorance in such an environment, i.e. how staying uninformed about the consumption risk can lead to consumption decisions that are more aligned with a consumer's long-term preferences. When forming his consumption decision, a consumer faces a trade-off between the instantaneous consumption utility and the expected future consumption risk. Hence, the consumer's assessment about his personal risk determines his consumption decision.

I identify a region of intermediate values of the risk parameter, for which the consumer would ex-ante like to commit to abstention, but where he consumes under complete information due to the self-control problem. I then derive the information signal about the risk parameter, which induces the maximal abstention for this region. The signal simply provides information on whether the risk parameter is above or below some threshold. Such a signal may be implement by simple regulatory measures as threshold-based consumption recommendations, certified labels, or medical definitions of risk groups.

Furthermore, I show that such a signal can be the informational outcome in an equilibrium of a dynamic game where the consumer can acquire costless additional information on the risk parameter before each consumption decision. Intuitively, while additional information may help the consumer evaluate the instantaneous gains from consumption, any information is shared with his future selfs, and may therefore induce a suboptimal high level of future consumption due to the present bias. As a result, in equilibrium the consumer abstains from any further information acquisition.

In this chapter the benefits of strategic ignorance are of a slightly different nature compared to the other chapters. As consumers cannot commit to future actions, and consumption decisions are time inconsistent, different incarnations of the consumer play a dynamic non-cooperative game against each other. Strategic ignorance provides an advantage in this intra-personal game.

Chapter 1

Consumer-Optimal Information Design

This Chapter is based on von Wangenheim (2017a).

1.1 Introduction

Over the recent decade, trade has increasingly shifted towards online markets. Commonly in these markets, consumers do not observe all product characteristics at the time of purchase. For instance, if a consumer buys clothes he may only have a vague idea of the cut and the color. If he books a hotel online, he may learn some coarse information from the hotel's number of stars or the reviews of other customers, but many consumer-specific details will remain unclear until his actual arrival. Consequently, despite having some ex-ante information, the buyer will learn whether the good matches his private taste sufficiently only ex post, ie., after he contracts with the seller and gets access to the good.

A monopolist may exploit this partial uninformedness of the consumer, and offer contracts which leave only small information rents to consumers. This chapter addresses the question to what extent regulation that obliges the monopolist to provide consumer-information, can protect consumers against an exploitation of the monopolist's market power. In particular, I derive the buyer-optimal information design in such markets.

On the simple intuition that more information cannot hurt the buyer, one might expect that consumer surplus can only be increasing in the amount of buyers' ex-ante private information. This is, however, not the case. Since sellers respond in their contract offers to the structure of buyers' private information, the choice of information exhibits a strategic effect on the subsequent contracting game. I show that the buyer-optimal ex-ante information keeps the buyer to some extent uninformed about his valuation. The buyer-optimal information signal induces efficient trade and distributes all rents in excess of the classical static monopoly profit with fully informed buyers, to the buyer. Moreover, a seller-optimal contract for the buyer-optimal information always consists of a simple buy-now offer without refund. In a second step, I characterize all divisions of buyer surplus and seller surplus that can arise under different information signals for a given prior.

While a buyer may have incomplete information about his value for a product at the time of contracting, I assume that he learns his exact valuation ex post by inspecting the product after delivery.[1] Due to the sequential information structure, the seller faces a sequential screening problem. She optimally screens buyers by offering several contracts, which differ in price and refund conditions, as studied in Courty and Li (2000).[2] I extend their sequential screening framework, by allowing the buyer to decide how much he wants to learn about his valuation for the good before contracting. More specifically, the buyer, or the regulator on behalf of the buyer, first chooses a signal about the valuation for the good. The seller observes the signal distribution, but not its realization. This assumption expresses that the seller can observe what the buyer learns, but not how it translates into the buyer's valuation.[3] Then, the seller offers a contract, before the buyer learns his true valuation.

The seller screens the buyer with respect to his signal realization. She optimally offers a menu of option contracts, each specifying a price, and the refund conditions. Intuitively, buyers with higher valuation uncertainty are more attracted by contracts with high refund flexibility.[4]

To better understand why a partial information revelation can be beneficial to the buyer, assume that trade is efficient for all buyer types, and consider the full information benchmark:

If the buyer learns his exact valuation by choosing a fully informative signal, there is no further learning. The seller charges the static monopoly price, leaving an information rent to the buyer. The rent, however, may come at the cost of trade

[1] Alternatively, assume that inspection only reveals some additional information, and interpret buyer's valuation as his updated value estimate. As buyers are risk neutral, this leaves all insights of the chapter unchanged.

[2] The optimality of sequential screening also features, among others, in Baron and Besanko (1984), Battaglini (2005), Eső and Szentes (2007), Hoffmann and Inderst (2011), Krähmer and Strausz (2011), Nocke et al. (2011), and Pavan et al. (2014).

[3] I am only interested in product information that relates to subjective taste. By the unraveling argument of Viscusi (1978), every seller will disclose any information on quality, if he can credibly and costlessly do so.

[4] Due to the buyer's freedom to design information signals, the regularity conditions, imposed in Courty and Li (2000), may be violated. Thus, we cannot rely on their analysis to find the optimal contract.

inefficiencies, since the monopoly price in general induces only high types to buy.

This benchmark naturally lead to the question, whether there is an only partially informative signal that induces efficient trade, and distributes the additional rents of this more efficient allocation to the buyer. Indeed, I show that for any prior distribution there is a suitable signal structure, such that the seller chooses a contract for which

1. trade is efficient, and

2. the seller only receives the static monopoly profit of fully informed buyers.

Note that this static monopoly profit is always a lower bound on the seller's profit, since she can always charge the static monopoly price and allow full refund, after the buyer learns his type. Hence, such a signal is buyer optimal in the sense that it maximizes the buyer surplus.

The optimal signal keeps low types partly uninformed, while high types have full information. Indeed, if different low types obtain the same signal, the seller can sell to these types by providing less information rent, since they have to break even only on average, rather than individually. Consequently, the seller has an incentive to lower the price below the static monopoly price, in order to increase participation. Lower prices increase efficiency as well as rents for high types.

Moreover, I fully characterize the possible combinations of buyer surplus and seller surplus that can arise in the sequential screening model for different signal distributions. Similar to Bergemann et al. (2015), I show that the only limits are imposed by the natural constraints that

1. buyer utility is nonnegative,

2. the seller receives at least the static monopoly profit, and

3. aggregate surplus does not exceed the first-best gains from trade.

The remainder of the chapter is structured as follows: After discussing the relevant literature in Section 1.2, I introduce the model in Section 1.3. Section 1.4 covers the case of a uniform distribution and provides an illustrative example. In Section 1.5, I construct the buyer-optimal signal. Section 1.6 characterizes all possible surplus division, before Section 1.7 concludes the chapter. All proofs are relegated to the appendix.

1.2 Related Literature

This chapter contributes to the growing literature on dynamic mechanism design, in which private information is learned over time. Baron and Besanko (1984) were the first to study dynamic price discrimination in a two-period procurement model with auditing. My model builds on the framework of Courty and Li (2000) who analyze optimal price discrimination for monopolistic markets in a two-period model. Battaglini (2005) and Pavan et al. (2014) provide general models on optimal dynamic mechanism design for longer time horizons.

A recent branch of the literature, building on the pioneer work of Lewis and Sappington (1994), studies sellers' strategic information revelation. Bergemann and Pesendorfer (2007) analyze auctions, where the seller can choose the accuracy by which the buyers learn their private valuations. They identify a trade-off between allocation efficiency and information rents. Eső and Szentes (2007) show that the trade-off disappears when the information provision is part of the contractual relationship, and argue that the seller should always disclose all relevant information. Li and Shi (2017) show that this no longer holds when the seller can use discriminatory information disclosure. If buyers have different ex ante types, and the provided information can depend to the reported types then partial information disclosure may be optimal. Hoffmann and Inderst (2011) characterize optimal contracts for the case where the buyer's and the seller's information are stochastically independent.

This chapter conversely analyzes buyers' optimal information acquisition. The agent may acquire private information costlessly and observably *before* the contractual relationship, to obtain a strategic advantage in the contracting game. This timing is in contrast to the classical literature on buyer's information acquisition in principal-agent relationships, where the principle aims to contractually provide incentives for costly learning.[5]

The model is probably closest related to Roesler and Szentes (2017), who characterize the buyer-optimal signal in a classical static one-unit trade environment. In contrast to their setup, I assume that after delivery the buyer receives additional information that affects his valuation. As a result, the seller may combine the contract with refund options, which—different to Roesler and Szentes—induces a lower bound on seller profit, and results in efficient trade for the buyer optimal signal. I consider my framework to be more appropriate in the context such as internet markets, where consumer typically receive additional information upon the good's delivery (see Krähmer and Strausz (2011)), while the context of Roesler and Szentes seems more appropriate in markets in which this learning effect is negligable.

Kessler (1998) analyzes the value of ignorance in a classical adverse selection

[5]E.g., Lewis and Sappington (1997), Crémer et al. (1998), Szalay (2009), Krähmer and Strausz (2011).

model with two types. She finds that, even if the agent can learn his type costlessly, he will choose a signal that is uninformative with some positive probability, in order to receive a more favorable contract.

Bergemann et al. (2015) analyze trading contracts, where the seller has information beyond the prior distribution. In particular, they characterize the buyer-optimal seller information structure. In contrast to my model the *seller* receives a signal, while the buyer is fully informed. In my model, the seller has to elicit information on the signal via an incentive compatible mechanism.

The idea that one party can choose arbitrary information signals to influence another party's decision has lately drawn a lot of attention, and produced a vast literature on Bayesian persuasion, based on the work of Kamenica and Gentzkow (2011). My setup is different from persuasion, since the buyer himself is uninformed. We can, however, make use of the tools from the framework of arbitrary signal choices. (One interpretation of my model is that the buyer tries to "persuade" himself, in the sense that he wants to manipulate his beliefs to obtain a strategic advantage towards the seller.)

1.3 The model

A seller can produce one unit of a good at zero cost. The valuation of the good for a buyer is drawn from a commonly known prior distribution $F(\theta)$ on some positive support $[\underline{\theta}, \overline{\theta}]$ with positive, continuous density $f(\theta)$.[6] Before contracting and learning the valuation, the buyer (or the regulator in the buyer's interest) can choose a signal structure to gain some information on the valuation. The signal distribution is commonly observed, the realization is private information to the buyer. I allow for any general signal structure in form of a Borel-measurable signal space $T \subseteq \mathbb{R}$, together with a probability measure μ on the Borel σ-algebra of $[\underline{\theta}, \overline{\theta}] \times T$. The buyer observes a signal $\tau \in T$, which is distributed according to the signal distribution

$$G(\tau) = \int_{t \leq \tau} \int_{\theta \in [\underline{\theta}, \overline{\theta}]} \mathbb{1}(t, \theta) d\mu.$$

The only restriction on the signal is the "consistency" with the prior F in the sense that

$$\int_{T \times [\underline{\theta}, \theta]} \mathbb{1} d\mu = F(\theta)$$

[6]The restriction to a positive support is only to keep the exposition transparent and tractable. It does not change the results. Indeed, if trade is inefficient for some buyer types, one can interpret an optimal learning process as a two step procedure. First the buyer learns whether $\theta > 0$, and then applies the optimal learning process described in this chapter to the conditional distribution on θ being larger than 0.

for all $\theta \in [\underline{\theta}, \overline{\theta}]$.[7]

The setup includes the common examples of a finite signal space $T = \{\tau_1, ..., \tau_n\}$ with $p_i = Prob(\tau_i)$, and the restriction that

$$\sum_{i=1}^{n} F(\theta|\tau_i)p_i = F(\theta),$$

as well as a continuous signal space $T = [\underline{\tau}, \overline{\tau}]$ with some distribution $G(\tau)$, and the restriction that

$$\int_{[\underline{\tau}, \overline{\tau}]} F(\theta|\tau)dG(\tau) = F(\theta).$$

The timing of the game is as follows:

1. the buyer publicly chooses a signal structure
2. the signal realization is privately observed by the buyer
3. the seller offers a contract, the buyer accepts / rejects
4. the buyer observes his type
5. transfers are made according to the rules of the contract

For any signal structure that reveals at least some information to the buyer, the seller in Stage 3 faces a classical sequential screening problem, as described in Courty and Li (2000). They show that any optimal deterministic contract can be implemented as a menu of option contracts from which the buyer can choose at the contracting stage. An option contract specifies an upfront payment a to the seller, and an option price p, for which the buyer can decide to buy, after he learns his true valuation.[8]

In the following section, I derive the buyer-optimal signal, which achieves the upper bound of buyer utility, for a uniform prior. In Section 3.1, I show how the construction generalizes to arbitrary prior distributions if we restrict the seller to the use of option contracts.

1.4 The Uniform Case

It is instructive to analyze first the case of a uniform prior, as it catches the main intuitions.

[7] We explicitly do not make common restrictions on the signal distribution, such as non-shifting support or an order by first-order stochastic dominance.

[8] Equivalently, one can interpret such a contract as a buy price of $a + p$, together with the option to return the good for a refund of p.

Let the prior $F(\theta)$ be the uniform distribution on $[0,1]$. Consider, as a benchmark, that the buyer fully learns his type θ under signal τ. The seller will then charge the monopoly price of

$$p^M = \arg\max_p p(1 - F(p)) = 1/2.$$

She will therefore sell to the buyer if and only if the buyer's valuation exceeds $1/2$. The seller's profit is $\pi^M = 1/4$, while the buyer's expected surplus is $1/8$.

Note that the seller can always ignore the possibility to exploit the signal for ex-ante screening, and just charge the monopoly price after the buyer learns the true valuation, i.e. $(a,p) = (0, p^m)$. Hence, the static monopoly profit of $\pi^M = 1/4$ defines a lower bound for the seller's utility.

Since trade is always efficient, the upper bound for buyer surplus is achieved, if trade always occurs, and the seller is left with her monopoly profit π^M. The main result of this section is that such a contract can be induced by the following signal.

$$\tau(\theta) = \begin{cases} 0, & \theta \leq \frac{1}{2}, \\ \theta, & \theta > \frac{1}{2}. \end{cases} \tag{1.1}$$

The buyer only learns his valuation if it is above $1/2$. Buyers with valuation below $1/2$ are pooled in one signal of $\tau = 0$, which induces an expected valuation of $\mathbb{E}[\theta|\tau = 0] = 1/4$.

Suppose the seller offers a single contract $(a,p) = (1/4, 0)$, which means she offers the good at a price of $1/4$ before the buyer learns θ with certainty. Since $\mathbb{E}[\theta|\tau] \geq 1/4$ for all τ, this offer will attract all buyers. Using the tools of mechanism design, I show in the appendix that, given this signal structure, there is no contract that generates a higher seller utility.

Proposition 1.1. *Given signal τ, there is no mechanism which generates a seller utility above $\frac{1}{4}$. In particular, the contract $(\frac{1}{4}, 0)$, which sells to all buyers ex ante at a price of $\frac{1}{4}$, is a seller-optimal trading mechanism.*

Since the seller is left with her lower bound utility of $1/4$, and social surplus is maximized, the signal τ implements the upper bound of buyer utility. It is therefore a buyer-optimal signal.

Even though the above construction of the optimal signal is specific to the uniform distribution, the main intuitions from this example carry over to the general case. It is suboptimal for the buyer to be fully informed about his valuation. If buyers with relatively low valuations remain partly uninformed, then the seller has to provide less information rent to sell to these types. To include lower types in trade, the seller must set low prices for *all* buyers. While low types' individual rationality constraints bind, and they make zero profits on average, high types benefit from lower prices and buyer surplus increases. Since more types trade, efficiency increases as well.

Applications and Discussion

There are numerous ways in which a regulator or intermediary can control the amount of product information exposed to consumers prior to trade.

One natural application for the use of information design are internet platforms. Especially in the hospitality and travelling industry it is common to offer car rentals, holiday packages, hotel stays, or airline tickets on internet platforms, such as online travel agencies.[9] By collecting personalized data, platforms can gather a profound understanding of consumers preferences. Further, they are able to discriminate product information with respect to individual consumers.

If the platform has to grant standard monopoly profits to hotels, but aims to maximize consumer surplus due to platform competition, it faces exactly the information design problem described in the model. As seen in Proposition 1.1 and generalized in Theorem 1.1, the platform optimally does not provide all product details to consumers, but leaves details of the deal somewhat opaque, and sells at low prices. Indeed, such "opaque deals", are common practice in online travel agencies such as priceline and hotwire: The platforms offer discounted deals, which guarantee specific features such as the number of hotel stars or location at the city center, but reveal the identity and other details of the hotel only after payment.[10]

Note that in contrast to Shapiro and Shi (2008) and Balestrieri et al. (2015), where opaque selling is a result of firm's profit maximizing behavior, in my model it appears as the natural tool to maximize consumer surplus, which provides a novel perspective for the use of opaque goods.[11]

Alternatively, a regulator may control the amount of product information by specific labelling requirements, certification standards, or—as Hoffmann et al. (2017) argue— by regulating the length of trial periods.

For instance, both the USA and the EU require sellers to label food ingredients on the package in descending order of predominance by weight, yet not by the exact

[9] According to Green and Lomanno (2012), in 2010 about 11 percent of all revenues in the US hotel industry were generated by online travel agencies like Expedia, Priceline, and Orbitz.

[10] Green and Lomanno (2012) find that about one quarter of all hotel bookings in online travel agencies involve opaque goods.

[11] In Shapiro and Shi (2008), opaque selling arises as an equilibrium under competition with differentiated consumers, whereas Balestrieri et al. (2015) show that a monopolist's optimal selling strategy for substitutes may feature opaque options.

amount.[12] The same EU regulation requires firms to label the nutrition value with a *Guideline Daily Amount* (GDA), however Grunert et al. (2010) find that only about 70 percent of customers have a conceptional understanding of its meaning. They find that the understanding is positively correlated with the interest in healthy eating, which suggests that the information design is particularly informative to consumers who have a high value for healthy food.

A different information design approach is taken by the Food Standard Agency (FSA) in the UK. In 2006 they introduced the traffic light rating system, under which nutrition values—such as sugar or saturated fat—are highlighted in red (high), amber (medium) or green (low). In their literature review on food labelling Hawley et al. (2013) conclude that traffic light ratings have “most consistently helped consumers to identify healthier products”.

The following example depicts how certification regulation may provide individual consumers with the information signal that generates the consumer-optimal information structure.

Example: Hotel Certifications

This stylized example aims to illustrate that, in the terms of this chapter, the current european system of hotel certification can be understood as providing consumers with only partial information. The example, however, also serves to show that, based on this chapter’s arguments, the information provision underlying the certification—designed by the hotel associations—is not consumer optimal. The consumer-optimal information regulation is derived.

In 2009 hotel associations from seven European countries founded the “Hotelstars Union” to harmonize the national standards of hotel certifications. By 2017 the system was adapted by 17 countries within the European Union, with only very slight differences between participating countries.

The grading system mainly consists of five different quality levels, represented by one to five stars[13]. Participating hotels gather points by providing features from a list of over 200 possible criteria, divided into categories as *reception, services, gastronomy,* and *leisure*. Hotels who want to certify a certain number of stars must achieve a respective number of points. Besides some minimum requirements for each star, hotels are entirely free in how to achieve the number of points.

While the number of stars may be a good measure for the overall hotel quality, it is quite uninformative about the match value with respect to private taste. While

[12] USA: 21CFR §101.4, EU: Regulation 1169/2011

[13] Sometimes there are intermediate grades denoted with the label “superior”, in addition to the number of stars, for details on the national certification gradings see https://www.hotelstars.eu/

business travellers may be exceptionally concerned about reception opening hours and good Wi-Fi, leisure travellers may have a higher valuation for available sports equipment and wellness services.

For the following example, consider a hotel that has certified a certain number of stars. Guests who consider to book online can infer from the number of stars alone only how many total points the hotel achieved in the grading system. The provided features remain unknown to the guests until they arrive at the hotel and can inspect it.

Potential guests g are—with equal probability—either of type *business* ($g = B$), which only care about features in the categories *reception and services*, or of type *leisure* ($g = L$), which only care about features in the categories *gastronomy and leisure*. Moreover, any type $g \in \{B, L\}$ has private preferences over the specific features within her relevant categories, which can be represented by a location x_g on the circumference of a circle with perimeter one (Salop's circle). Assume for each type $g \in \{B, L\}$ that the location is uniformly distributed on the circumference. Guests are risk neutral and the realization of (g, x_g) is their private information.

The hotel h provides—with equal probability—either mainly features for the business type (focus $h = B$), or mainly for the leisure type (focus $h = L$). The exact features that the hotel offers for each, business and leisure type, may be represented by two locations (y_B, y_L) on the circumferences of two distinct circles, each with perimeter one. Assume, again, that y_B and y_L are each ex ante uniformly distributed on their circumference. Marginal costs for additional guests are zero. The values of (h, y_B, y_L) are private information of the hotel, but become observable after the guest arrives at the hotel.

Let the utility of a guest (g, x_g) who pays price p to stay in a hotel (h, y_B, y_L) be given by

$$u((g, x_g), (h, y_B, y_L), p) = 0.5 \cdot \mathbb{1}_{g=h} + (0.5 - d(x_g, y_g)) - p,$$

where $d(x_g, y_g) \in [0, 0.5]$ describes the distance of x_g and y_g on the circumference of the circle. In other words, the guest receives a utility of 0.5 if the hotel has a focus suitable for his type, and an additional utility up to 0.5 if the hotel offers preferred features *within* the relevant categories. As neither the guest nor the hotel have ex ante any information on the match value, the value is ex ante uniformly distributed on $[0, 1]$, with values in $[0, 0.5]$ for types $g \neq h$ and values in $[0.5, 1]$ for types $g = h$.

Since guests' valuation is ex ante uniform on $[0, 1]$ they have an expected value of 0.5. The hotel can set $p = 0.5$ without refund option, and all guests will accept the offer, leading to an efficient outcome where all rents are realized by the hotel.

Consider now the role of a regulator who is solely interested in consumer surplus and has the power to precisely regulate the information the hotel has to provide online. If the regulator forces the hotel to disclose its focus $h \in \{B, L\}$, as well as the available features only in its focal categories—thus the exact position of y_h—the guest receives a signal exactly as defined in (1.1): Guest types $g = h$ with a valuation above 0.5 learn their exact valuation as they learn the hotel's position y_g, while guest types $g \neq h$ don't learn the relevant position y_g and remain pooled with valuations in $[0, 0.5]$. As we have seen in Proposition 1.1, such a signal is guest-optimal, and implements the upper bound of the guests' utility.

The intuition of the example is simple: The regulator should exclusively allow information that is relevant to high types. High types receive high information rents, while the low types, who are relatively uninformed, induce the seller to set low prices.

1.5 The General Case

The main result of the chapter is that the buyer can achieve his first-best for any arbitrary prior distribution:

Theorem 1.1. *Let π^M be the standard static monopoly profit the seller can achieve if the buyer privately learns his valuation before contracting. Then there exists an information signal such that for the seller's optimal menu of option contracts*

- *trade is efficient, and*
- *the seller receives π^M.*

Such a signal is buyer-optimal, since it maximizes aggregate surplus and leaves the seller with her lower bound of utility π^M.

The result follows immediately from the more general Theorem 1.2. I will, however, provide the intuition how to construct such a signal structure.

If a menu of contracts induces trade for all types θ, then each type will trade at the lowest available total payment $a + p$. This means that any signal which induces full trade and a seller profit of π^M must necessarily sell to all buyers at a uniform price of $a + p = \pi^M$.

We start by defining y as the type such that

$$\mathbb{E}[\theta | \theta \leq y] = \pi^M.$$

Consider first a signal structure similar as that defined in Equation (1.1) of Section 1.4, where types above y learn their valuation, while types $\theta \leq y$ are pooled in one signal realization. By construction, the pooled buyers are ex ante willing to pay at most π^M for the good. Therefore, a uniform ex ante price of π^M would attract all buyers, and the seller would make monopoly profit π^M.

However, for many prior distributions F such a uniform price with full participation is suboptimal for the seller, given the signal. The seller may make higher profit, if she chooses to exclude types below some threshold $\hat{\theta} < y$ from trade. She could achieve this for example with a menu offer of $\mathcal{M} = \{(a, \hat{\theta})\}$, where a is chosen such that types $\theta \leq y$ in the pooling area receive a utility of zero in expectation. Such a contract comes at the cost of losing the participation of types $\theta < \hat{\theta}$, but at the gain of higher prices and thus more revenue from all types $\theta > y$.

One can modify the signal structure such that such contracts are never optimal. The idea is to give buyers more information by maintaining the property that $\mathbb{E}[\theta|\tau] = \pi^M$ for all $\theta \leq y$. Figure 1.1 illustrates how to achieve this goal.

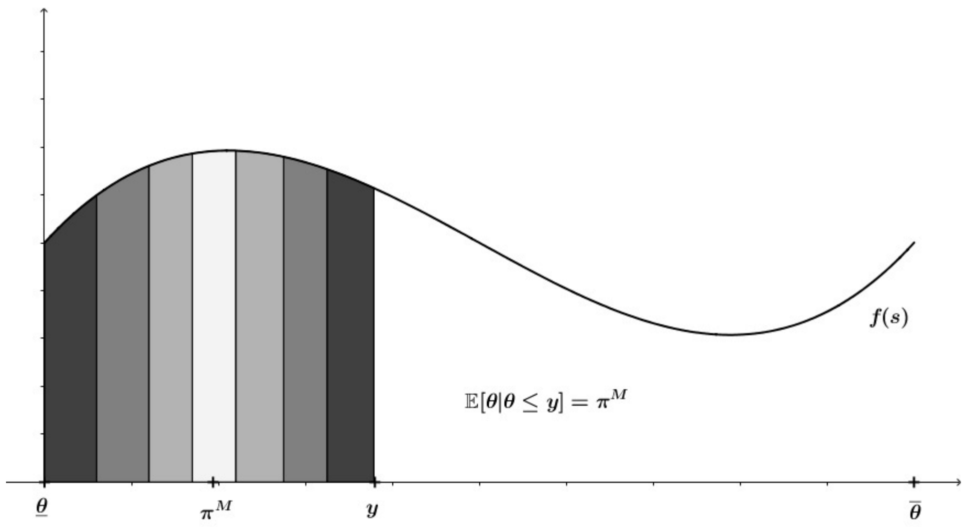

Figure 1.1: Partition of the optimal signal structure (stylized)

Types above y fully learn their valuation, while types below y learn that their type is in a certain pooling region, represented by the shade of gray, assigned to their type in Figure 1.1. The shaded areas are constructed in such a way that for any shade signal τ

$$\mathbb{E}[\theta|\tau] = \pi^M.$$

Further, if τ_1 is darker than τ_0, then $F(\cdot|\tau_1)$ is a mean preserving spread of $F(\cdot|\tau_0)$. Now, we let the area of each shade signal shrink, while we let the number of different shades goes to infinity. In the limit, we obtain a continuum of shades, where each signal τ only pools two types $\{\theta^L_\tau, \theta^H_\tau\}$ with $\theta^L_\tau < \pi^M < \theta^H_\tau$, and $\mathbb{E}[\theta|\tau] = \pi^M$.

If the seller now aims to sell to any type $\hat{\theta} < \pi^M$, then—by ex ante individual rationality—she has to offer a contract that charges at most $a + p = \mathbb{E}[\theta|\tau(\theta) = \tau(\tilde{\theta})] = \pi^M$. Such a contract would attract all types, and generate a profit of π^M to the seller. It turns out, this is the best the seller can do: suppose that the seller aims for higher prices at the cost of participation. If it were optimal for the seller to use a menu for which the lowest type that buys satisfies $\hat{\theta} > \pi^M$, then one can show that the best the seller could do is to offer a contract $(0, \hat{\theta})$. Such a contract is equivalent to an ex post take-it-or-leave-it offer with price $\hat{\theta}$, which certainly cannot generate more profit to the seller than the optimal static monopoly price with fully informed buyers.

1.6 The Limits of Surplus Distribution

In the previous section, I have analyzed a signal which maximizes buyer surplus. In this section, I characterize which combinations of buyer surplus and seller surplus are feasible in a sequential screening framework. Similarly to Bergemann et al. (2015), Figure 1.2 characterizes the natural constraints to this problem graphically.

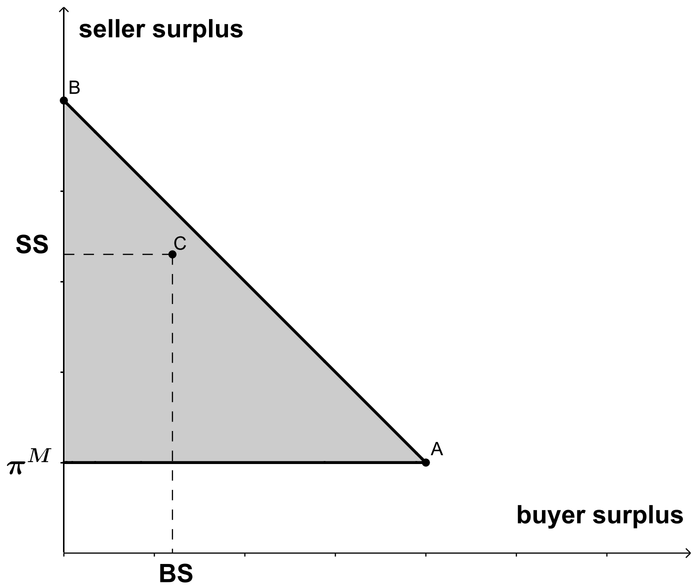

Figure 1.2: All potential pairs of surplus division

First of all, by buyer's individual rationality, expected buyer surplus will never be negative. Second, as argued in the previous section, seller surplus can never fall below the static monopoly profit, since the seller can always use a static mechanism after the buyer learned his true type. Finally, aggregate surplus cannot exceed first-best welfare, which is sketched as the diagonal pareto frontier. Consequently, any surplus pair must lie in the gray shaded triangle. Point A corresponds to the buyer-optimal signal, as constructed in Section 3.1. Point B corresponds to the case, where the seller has no ex ante information upon the prior distribution. In this case, the seller can extract the entire surplus by selling ex ante at a price of $\mathbb{E}[\theta]$.

We will see that *any* arbitrary point C in the triangle can be implemented as the solution to the seller's problem for an appropriate signal distribution.

Theorem 1.2. *There exists a signal and an optimal sequential selling mechanism with seller surplus u_S and buyer surplus u_B if and only if*

- $u_B \geq 0$,
- $u_S \geq \pi^M$, *and*
- $u_S + u_B \leq \mathbb{E}[\theta]$,

where π^M is the standard static monopoly profit the seller can achieve, if the buyer has full information.

A full proof can be found in the appendix. I will sketch the main steps here. Take an arbitrary surplus pair (BS, SS) which satisfies the above constraints. We will construct a corresponding signal, such that, indeed, buyer surplus and seller surplus are given by $(u_B, u_S) = (BS, SS)$.

Call $AS = BS + SS$ the aggregate surplus we want to construct. Define the threshold x by

$$AS = \int_x^{\overline{\theta}} f(\theta)\theta d\theta = (1 - F(x))\mathbb{E}[\theta|\theta \in [x, \overline{\theta}]].$$

Note that the gains from trade $u_S + u_B$ are indeed given by AS, if we can construct a signal for which exactly all types above x trade.
Next, we define the threshold $y \leq \overline{\theta}$ by

$$SS = (1 - F(x))\mathbb{E}[\theta|\theta \in [x, y]].$$

Further, define $\overline{a} \in [x, y]$ by

$$\overline{a} := \mathbb{E}[\theta|\theta \in [x, y]].$$

We will use a similar construction as in Section 3.1 to build a signal for which the seller chooses to sell ex ante to all types $\theta \geq x$ at a uniform price of $\overline{a}$. In this case, seller surplus u_S is indeed given by SS, and buyer surplus is $u_B = AS - SS = BS$.

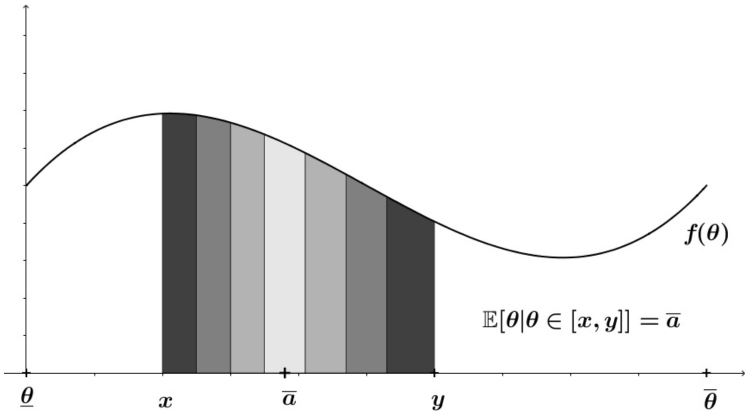

Figure 1.3: The signal to induce $(u_B, u_S) = (BS, SS)$

Figure 1.3 depicts the construction of the signal structure. Types below x and above y fully learn their valuation. Types in the interval $[x, y]$ again learn their pooling region that is assigned to their type, illustrated by the corresponding shade of grey. The shaded areas are constructed in such a way, that for any shade signal τ we have

$$\mathbb{E}[\theta|\tau] = \overline{a}.$$

If we let the number of different shades go to infinity, we obtain a continuum of shades. In the limit, the signal can be represented by

$$\tau(\theta) = \begin{cases} \theta - \overline{\theta}, & \theta < x, \\ \int_{\theta}^{\overline{a}} f(s)(\overline{a} - s)ds, & \theta \in [x, y], \\ \theta, & \theta > y. \end{cases}$$

While types $\theta \notin [x, y]$ learn their valuation, the signal $\tau(\theta)$ for each $\theta \in [x, y]$ corresponds to exactly two types $\{\theta^L_\tau, \theta^H_\tau\}$, which satisfy $\theta^L_\tau \leq \overline{a} \leq \theta^H_\tau$, and $\mathbb{E}[\theta|\tau] = \overline{a}$. Figure 1.4 depicts signal τ as a function of θ.

Now, consider an optimal menu of contracts, the seller will offer to the buyer. Let $\hat{\theta}$ be the lowest type which buys given this menu. If $\hat{\theta} \notin [x, y]$, then the buyer learns $\hat{\theta}$ with certainty under τ, and the seller can charge at most $a+p = \hat{\theta}$ from this type. Since this contract is available to all types, expected profits cannot exceed

$$(1 - F(\hat{\theta}))\hat{\theta} \leq \max_p\{(1 - F(p))p\} = \pi^M \leq SS.$$

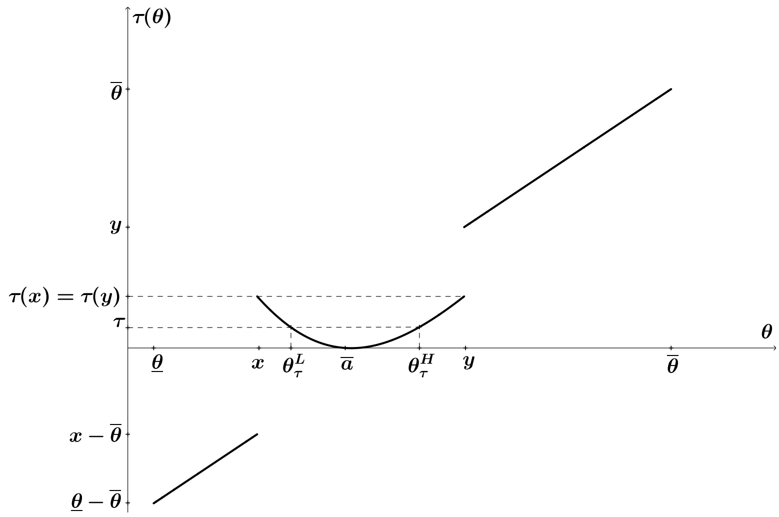

Figure 1.4: Signal $\tau(\theta)$ for the distribution in Figure 1.3 (stylized)

The seller can do (weakly) better if she decides to sell ex ante to all types in $[x, y]$. Since for all these types we have $\mathbb{E}[\theta|\tau] = \overline{a}$, a contract $(a, p) = (\overline{a}, 0)$ attracts exactly all types $\theta \geq x$, and seller's profit is $(1 - F(x))\overline{a} = SS$. One can show that this contract is indeed optimal from the seller's perspective.

The boundaries x and y define three partitions of types. In the optimal contract, types $\theta < x$ don't trade and therefore induce an efficiency loss. Types $\theta \in [x, y]$ trade, but make no surplus in expectation. Buyers $\theta > y$ receive an information rent and extract the entire buyer surplus. Intuitively, by moving the boundaries x and y one can realize any buyer and seller surplus that satisfies the natural constraints in Theorem 1.2.

1.7 Conclusion

This chapter emphasizes the role of private information in sequential screening. It shows that there are almost no restrictions to the division of buyer and seller surplus, that can arise in sequential screening for different ex-ante information.

The buyer-optimal signal keeps the buyer to some extent uninformed about his valuation at the time of contracting. Thus, the ability of a monopolist to screen buyers sequentially may not necessarily harm consumers, but lead to lower prices, and increase efficiency.

It is worth noting that the European Union has taken a clear stand on the issue of consumer rights in online markets. By Directive 2011/83/EU, any consumer is granted the right to withdraw from online contracts within 14 days after delivery. As Krähmer and Strausz (2015) point out, this policy effectively destroys the ability of a monopolist to screen ex ante, granting the consumers the same information rent, as under full information.[14,15] This chapter shows, that regulators, who care about consumer utility, can improve on this regulation, if they have sufficent control over information provided to individual buyers. The regulation of information may be more flexible and powerful than the regulation of contracts, and therefore deserves further study.

1.8 Appendix to Chapter 1

Proof of Proposition 1.1. By the revelation principle for dynamic games (e.g., Myerson (1986)), we can restrict attention to direct, incentive compatible mechanisms:

The buyer reports his private information sequentially. After learning τ, he reports its realization to the seller. If $\tau \in (0.5, 1]$ then $\tau = \theta$, thus a truthful report of τ reveals θ already. If the buyer reports $\tau = 0$ then the seller asks for a report of $\theta \in [0, 0.5]$ after the buyer observes its realization.[16]

A direct mechanism specifies the trading rules as a function of the buyer's reports. Formally, the allocation rule

$$q : (\{0\} \times [0, 0.5]) \cup (0.5, 1] \to [0, 1],$$

assigns to each complete report a probability of receiving the good. The transfer rule

$$t : (\{0\} \times [0, 0.5]) \cup (0.5, 1] \to \mathbb{R},$$

assigns to each complete report a monetary transfer from the agent to the principal.

Note that since τ defines a partition on θ, each feasible report corresponds exactly to one claim of being some type $\theta \in [0, 1]$. Identifying the report space

[14] Krähmer and Strausz (2015) already find that the welfare effect of such a policy is ambiguous, and depends on consumers' ex ante private information.

[15] The buyer can be forced to bear the shipping cost of returning the good, so there may be scope for ex-ante screening to some very limited extent.

[16] In Myerson (1986), the agent's report space is the entire support of his private information at each stage. That is, if the agent lies about the value of τ, he may still report θ truthfully and inconsistent with the report of τ. Since in our case τ defines a partition on all types θ, however, the seller can immediately detect and punish any untruthful report (τ, θ) with $\tau(\theta) \neq \theta$, such that the buyer would never choose such a report. It is therefore without loss of generality to consider only direct mechanisms, which restrict the reports of θ to values that are admissible for the reported τ.

with the type space, let

$$u(\hat{\theta}|\theta) = \theta q(\hat{\theta}) - t(\hat{\theta})$$

be the utility of buyer of type θ reporting as if being type $\hat{\theta}$. Let us further simplify notation by $u(\theta) := u(\theta|\theta)$. The incentive constraints, which guarantee truthful reporting, read

$$\forall\theta \in (0.5, 1], \hat{\theta} \in [0, 1] \qquad u(\theta) \geq u(\hat{\theta}|\theta) \qquad \text{(IC } \tau \neq 0\text{)}$$
$$\forall\hat{\theta} \in [0.5, 1] \qquad \mathbb{E}[u(\theta)|\tau = 0] \geq \mathbb{E}[u(\hat{\theta}|\theta)|\tau = 0] \qquad \text{(IC } \tau = 0\text{)}$$
$$\forall\theta \in [0, 0.5], \hat{\theta} \in [0, 0.5] \qquad u(\theta) \geq u(\hat{\theta}|\theta). \qquad \text{(IC } \theta\text{)}$$

First period individual rationality reads

$$\forall\theta \in [0.5, 1] \qquad u(\theta) \geq 0 \qquad \text{(IR } \tau \neq 0\text{)}$$
$$\mathbb{E}[u(\theta)|\tau = 0] \geq 0 \qquad \text{(IR } \tau = 0\text{)}$$

Since the seller's utility equals social surplus minus buyer's utility, her program is

$$\mathcal{P}: \qquad \max_{(q,t)} \int_0^1 (\theta q(\theta) - u(\theta))d\theta$$
$$s.t. \qquad \text{(IC } \tau \neq 0\text{),(IC } \tau = 0\text{),(IC } \theta\text{),(IR } \tau \neq 0\text{),(IR } \tau = 0\text{).}$$

We will derive the optimum for a so called "relaxed" problem $\mathcal{P}'$ with less constraints, and verify ex post that the remaining constraints are satisfied for the derived solution. The constraint (IC $\tau \neq 0$) directly implies the weaker condition

$$\forall\theta \in (0.5, 1], \hat{\theta} \in [0.5, 1] \qquad u(\theta) \geq u(\hat{\theta}|\theta). \qquad \text{(IC' } \tau \neq 0\text{)}$$

We now define

$$\mathcal{P}': \qquad \max_{(q,t)} \int_0^1 (\theta q(\theta) - u(\theta))d\theta$$
$$s.t. \qquad \text{(IC' } \tau \neq 0\text{),(IC } \theta\text{),(IR } \tau = 0\text{).}$$

The solution to program $\mathcal{P}'$ must implement weakly higher seller surplus than program $\mathcal{P}$, as it faces less constraints.

By Revenue Equivalence (e.g., Myerson (1981)), (IC θ) is equivalent to

1. $q(\theta)$ is increasing on $[0, 0.5]$, and
2. $u(\theta) = u(0) + \int_0^\theta q(s)ds$ for all $\theta \in [0, 0.5]$.

From 2. it follows that

$$u(\theta) = u(0.5) - \int_\theta^{0.5} q(s)ds.$$

Further, in any optimal solution, (IR $\tau = 0$) must bind, because otherwise the seller could uniformly raise the transfer for all types. Using integration by parts we obtain

$$\begin{aligned} 0 &= \int_0^{0.5} u(\theta)d\theta \\ &= \int_0^{0.5} \left(u(0.5) - \int_\theta^{0.5} q(s)ds \right) d\theta \\ &= 0.5u(0.5) - \left[\theta \int_\theta^{0.5} q(s)ds \right]_0^{0.5} + \int_0^{0.5} -\theta q(\theta)d\theta \\ &= 0.5u(0.5) - \int_0^{0.5} \theta q(\theta)d\theta, \end{aligned}$$

or equivalently

$$0.5u(0.5) = \int_0^{0.5} \theta q(\theta)d\theta. \tag{1.2}$$

Again by Revenue Equivalence, (IC' $\tau \neq 0$) implies that on any closed interval $[\tilde{\theta}, 1] \subset (0.5, 1]$, the allocation $q(\theta)$ is weakly increasing, and further for any $\theta \in [\tilde{\theta}, 1]$

$$u(\theta) = u(\tilde{\theta}) + \int_{\tilde{\theta}}^{\theta} q(s)ds.$$

Therefore, $u(\theta)$ is continuous on $(0.5, 1]$, and, because $q(\theta)$ is weakly positive and bounded, $\lim_{\theta \searrow 0.5} u(\theta)$ exists. Since by (IC' $\tau \neq 0$) for all $\theta \in (0.5, 1]$

$$u(\theta) \geq u(0.5, \theta) = (\theta - 0.5)q(0.5) + u(0.5) \geq u(0.5),$$

we have necessarily

$$\lim_{\theta \searrow 0.5} u(\theta) \geq u(0.5).$$

Moreover, if we had $\lim_{\theta \searrow 0.5} u(\theta) = u(0.5) + \varepsilon$, for some $\varepsilon > 0$, the seller could increase all transfers of types $\theta \in (0.5, 1]$ uniformly by ε and still satisfy all constraints of $\mathcal{P}'$. A mechanism with $\lim_{\theta \searrow 0.5} u(\theta) > u(0.5)$ can therefore not be optimal. It follows that any solution to $\mathcal{P}'$ must satisfy

$$u(\theta) = u(0.5) + \int_{0.5}^{\theta} q(s)ds$$

for all $\theta \in [0.5, 1]$.

For the seller's objective function in $\mathcal{P}'$ we obtain

$$\begin{aligned}
\int_0^1 (\theta q(\theta) - u(\theta))d\theta &= \int_0^1 \theta q(\theta)d\theta - \underbrace{\int_0^{0.5} u(\theta)d\theta}_{=0} - \int_{0.5}^1 u(\theta)d\theta \\
&= \int_0^1 \theta q(\theta)d\theta - \int_{0.5}^1 \left(u(0.5) + \int_{0.5}^{\theta} q(s)ds \right) d\theta \\
&= \int_0^1 \theta q(\theta)d\theta - 0.5u(0.5) - \int_{0.5}^1 \int_{0.5}^{\theta} q(s)dsd\theta \\
&\overset{(1.2)}{=} \int_0^1 \theta q(\theta)d\theta - \int_0^{0.5} \theta q(\theta)d\theta - \int_{0.5}^1 \int_{0.5}^{\theta} q(s)dsd\theta \\
&= \int_{0.5}^1 \theta q(\theta)d\theta - \left[\theta \int_{0.5}^{\theta} q(s)ds \right]_{0.5}^1 + \int_{0.5}^1 \theta q(\theta)d\theta \\
&= \int_{0.5}^1 \theta q(\theta)d\theta - \int_{0.5}^1 q(s)ds + \int_{0.5}^1 \theta q(\theta)d\theta \\
&= \int_{0.5}^1 (2\theta - 1)q(\theta)d\theta.
\end{aligned}$$

Note that the seller's utility is independent of the allocation for types $\theta \leq 0.5$. Indeed, any attempt to increase surplus from these types equally increases the information rent the seller has to provide to types $\theta \in [0.5, 1]$.

Since $(2\theta - 1) > 0$ for $\theta > 0.5$, the seller maximizes her utility by setting $q(\theta) = 1$ for any $\theta > 0.5$. The seller's maximal utility under $\mathcal{P}'$ therefore is

$$\int_{0.5}^1 (2\theta - 1)d\theta = [\theta^2 - \theta]_{0.5}^1 = 0.25.$$

If the seller chooses to set $q(\theta) = 1$ for all $\theta \leq 0.5$ then the direct mechanism takes the form

$$q(\theta) \equiv 1,$$

and

$$t(\theta) \equiv 0.25,$$

which corresponds exactly to the offer to sell the product ex ante at a uniform price of 0.25. Since all buyer types obtain the same offer in this contract, it satisfies all incentive constraints of $\mathcal{P}$. Moreover, the contract yields positive profit for all $\theta \geq 0$, therefore it satisfies the constraint (IR $\tau \neq 0$) of program $\mathcal{P}$ as well. $\square$

Proof of Theorem 1.2. Take some arbitrary $SS \geq \pi^M$ and $BS \geq 0$, with $BS + SS \leq \mathbb{E}[\theta]$. We need to construct a signal such that the seller's optimal mechanism induces seller utility $u_S = SS$ and buyer utility $u_B = BS$.

Constructing the signal

Define x implicitly by

$$BS + SS = \int_x^{\overline{\theta}} \theta dF(\theta) = (1 - F(x))\mathbb{E}[\theta | \theta \in [x, \overline{\theta}]]. \tag{1.3}$$

Since f has full support, the right hand side is strictly decreasing in x for $x \in [\underline{\theta}, \overline{\theta}]$, with $\int_{\underline{\theta}}^{\overline{\theta}} \theta dF(\theta) = \mathbb{E}[\theta]$, and $\int_{\overline{\theta}}^{\overline{\theta}} \theta dF(\theta) = 0$. Since

$$0 < BS + SS \leq \mathbb{E}[\theta],$$

there is exactly one $x \in [\underline{\theta}, \overline{\theta}]$, for which (1.3) is satisfied.[17]

Define now y implicitly by

$$SS = (1 - F(x))\mathbb{E}[\theta | \theta \in [x, y]]. \tag{1.4}$$

The right hand side is strictly increasing in y and since

$$(1-F(x))\mathbb{E}[\theta|\theta \in [x,x]] = (1-F(x))x \leq \pi^M \leq SS \leq BS+SS = (1-F(x))\mathbb{E}[\theta|\theta \in [x,\overline{\theta}]],$$

there is exactly one $y \in [x, \overline{\theta}]$, which satisfies (1.4). Further, we call

$$\overline{a} := \mathbb{E}[\theta | \theta \in [x, y]].$$

Finally, we define the following signal structure:

$$\tau(\theta) = \begin{cases} \theta - \overline{\theta}, & \theta < x, \\ \int_{\theta}^{\overline{a}} f(s)(\overline{a} - s)ds, & \theta \in [x, y], \\ \theta, & \theta > y. \end{cases}$$

The signal prescribes full learning for $\theta < x$ and $\theta > y$. For $\theta \in [x, y]$ the function $\tau(\theta)$ is continuous and strictly decreasing on $[x, \overline{a}]$, and strictly increasing on $[\overline{a}, y]$,

[17] The assumption that F is continuous and increasing is innocuous and only for mathematical convenience. If F has atoms, then $\tau(\theta)$ is not deterministic. If F is not increasing, we loose uniqueness of x and y. None of the results or intuitions hinge on these assumptions.

with

$$
\begin{aligned}
\tau(x) &= \int_x^{\overline{a}} f(s)(\overline{a}-s)ds \\
&= \int_x^{y} f(s)(\overline{a}-s)ds + \int_y^{\overline{a}} f(s)(\overline{a}-s)ds \\
&= (F(y)-F(x)) \underbrace{\left(\overline{a} - \frac{\int_x^y f(s)sds}{F(y)-F(x)}\right)}_{=0} + \int_y^{\overline{a}} f(s)(\overline{a}-s)ds \\
&= \tau(y).
\end{aligned}
$$

Thus, for any τ with $0 < \tau \leq \tau(x)$ there are exactly two types $\theta_\tau^L, \theta_\tau^H$ with $\tau = \tau(\theta_\tau^L) = \tau(\theta_\tau^H)$, where without loss of generality $\theta_\tau^L < \overline{a} < \theta_\tau^H$. Let us call $\theta^L(\tau)$ the inverse function of $\tau(\theta)$ on $[x, \overline{a}]$, and $\theta^H(\tau)$ the inverse function of $\tau(\theta)$ on $[\overline{a}, y]$. This means that the distribution of τ is given by

$$G(\tau) = F(\theta^H(\tau)) - F(\theta^L(\tau)).$$

It follows[18] that for any $\tau \in (0, \tau(x)]$

[18] We denote by $\mathbb{P}(A|\tau)$ the regular conditional probability for A given τ. This notion extends the concept of conditional probabilities to the case where one conditions on events of probability zero. The regular conditional probability is defined by the condition that for any measurable sets A, B the equality $\mathbb{P}(\theta \in A, \tau \in B) = \int_B \mathbb{P}(A|\tau)dG(\tau)$ holds. It is unique almost surely. Since we are interested in expectations only, this restriction is innocuous. For formal details see for example §7 on regular conditional distributions in Shiryaev (1996).

$$\begin{aligned}
\mathbb{P}(\theta^H_\tau|\tau) &= \mathbb{P}(\theta > \overline{a}|\tau) \\
&= \lim_{\varepsilon\to 0} \mathbb{P}(\theta > \overline{a}|\tau(\theta) \in [\tau, \tau+\varepsilon]) \\
&= \lim_{\varepsilon\to 0} \frac{F(\theta^H(\tau+\varepsilon)) - F(\theta^H(\tau))}{F(\theta^H(\tau+\varepsilon)) - F(\theta^H(\tau)) + F(\theta^L(\tau)) - F(\theta^L(\tau+\varepsilon))} \\
&= \frac{f(\theta^H_\tau)\theta^{H\prime}(\tau)}{f(\theta^H_\tau)\theta^{H\prime}(\tau) - f(\theta^L_\tau)\theta^{L\prime}(\tau)} \\
&= \frac{f(\theta^H_\tau)/\tau'(\theta^H_\tau)}{f(\theta^H_\tau)/\tau'(\theta^H_\tau) - f(\theta^L_\tau)/\tau'(\theta^L_\tau)} \\
&= \frac{1/(\theta^H_\tau - \overline{a})}{1/(\theta^H_\tau - \overline{a}) + 1/(\overline{a} - \theta^L_\tau)} \\
&= \frac{\overline{a} - \theta^L_\tau}{\theta^H_\tau - \theta^L_\tau}.
\end{aligned}$$

Similarly, we have

$$\mathbb{P}(\theta^L_\tau|\tau) = \frac{\theta^H_\tau - \overline{a}}{\theta^H_\tau - \theta^L_\tau}.$$

It follows that

$$\mathbb{E}[\theta|\tau] = \frac{\theta^H_\tau - \overline{a}}{\theta^H_\tau - \theta^L_\tau}\theta^L_\tau + \frac{\overline{a} - \theta^L_\tau}{\theta^H_\tau - \theta^L_\tau}\theta^H_\tau = \overline{a}. \tag{1.5}$$

This means that for any $\tau_1, \tau_2 \in [0, \tau(x)]$ with $\tau_1 < \tau_2$, the distribution $F(\cdot|\tau_2)$ is a mean-preserving spread of $F(\cdot|\tau_1)$.[19]

The menu

We turn to the seller's decision problem to choose an optimal menu of option contracts, given τ. Consider the menu $\mathcal{M} = \{(\overline{a}, 0)\}$. All buyers with $\theta < x$ receive a fully informative signal $\tau < 0$, and know with certainty that their valuation satisfies $\theta < \overline{a}$, so they would reject the contract. Types $0 \leq \tau \leq \tau(x)$ satisfy $\mathbb{E}[\theta|\tau] = \overline{a}$, and types $\tau > \tau(x)$ satisfy $\mathbb{E}[\theta|\tau] = \tau > \overline{a}$, so they would both accept the contract $(\overline{a}, 0)$, which sells ex ante at a uniform price of $\overline{a}$. This means that under contract $\mathcal{M}$ we have

$$u_S = \overline{a}(1 - F(x)) = SS,$$

[19] Note however, that the common assumption in Courty and Li (2000) of „non-shifting support"is violated. Thus, we cannot use their standard procedure to solve the seller's maximization problem.

and

$$u_B = \int_x^{\overline{\theta}} \theta dF(\theta) - u_S = (BS + SS) - SS = BS.$$

This shows that the menu $\mathcal{M}$ indeed implements the buyer and seller utility we want to construct. It remains to show, that $\mathcal{M}$ is an optimal menu for the seller for the given signal τ.

The optimality of the menu

Let $\tilde{\mathcal{M}} = \{(a_i, p_i)\}_{i \in I}$ be an arbitrary menu of option contracts. We need to show that it does not generate higher seller utility than SS.

Let $\hat{\theta}$ be the lowest type who purchases the good under $\tilde{\mathcal{M}}$, in the sense that he chooses some $(a, p) \in \tilde{\mathcal{M}}$ to pay the upfront fee a, and decides to buy the good at the price p, after he learns his type.

Case 1: $\hat{\theta} < x$ or $\hat{\theta} > y$

In this case $\hat{\theta}$ learns his type with certainty under τ. Since, by assumption, he accepts the contract (a, p), we can conclude that

$$a + p \leq \hat{\theta}.$$

Further, any buyer's signal $\tau(\theta)$ reveals to the buyer with certainty whether his type satisfies $\theta > \hat{\theta}$. This means, that any buyer with $\theta > \hat{\theta}$ learns from his signal realization that he receives positive utility from contract (a, p). Consequently no type $\theta > \hat{\theta}$ will accept a contract at higher total cost than $a + p$. Since $\hat{\theta}$ is by assumption the lowest type that buys, we can conclude that

$$u_S \leq (a + p)(1 - F(\hat{\theta})) \leq \hat{\theta}(1 - F(\hat{\theta}) \leq \max_p \{(1 - F(p))p\} = \pi^M \leq SS.$$

Case 2: $\hat{\theta} \in [x, \overline{a}]$

Then $\hat{\theta}$ is the low type for the respective signal realization, ie. $\hat{\theta} = \theta^L_{\tau(\hat{\theta})} < \theta^H_{\tau(\hat{\theta})}$. Thus, since type $\theta^L_{\tau(\hat{\theta})}$ purchases the good under (a, p), so will type $\theta^H_{\tau(\hat{\theta})}$. By buyer's ex ante individual rationality we have

$$a + p \leq \mathbb{E}[\theta | \tau(\hat{\theta})] = \overline{a}.$$

The contract (a, p) is therefore in particular also profitable to all types $\theta > y$, who learn their valuation ex ante with certainty. Hence, any of these types will as well pay at most $a + p \leq \overline{a}$. Thus, even if the seller extracts all surplus from types

$\theta \in [\hat{\theta}, y]$, her surplus is bounded by

$$\begin{aligned} u_S &\leq \int_{\hat{\theta}}^{y} \theta dF(\theta) + (1 - F(y))\overline{a} \\ &\leq \int_{x}^{y} \theta dF(\theta) + (1 - F(y))\overline{a} \\ &= (F(y) - F(x))\overline{a} + (1 - F(y))\overline{a} \\ &= (1 - F(x))\overline{a} \\ &= SS \end{aligned}$$

Case 3: $\hat{\theta} \in [\overline{a}, y]$
Then $\hat{\theta}$ is the high type for the respective signal realization, ie. $\hat{\theta} = \theta^H_{\tau(\hat{\theta})}$. Moreover, we have $\theta^H_{\tau(\hat{\theta})} \geq p > \theta^L_{\tau(\hat{\theta})}$, because otherwise $\theta^L_{\tau(\hat{\theta})}$ would purchase the good for p whenever $\theta^H_{\tau(\hat{\theta})}$ does, violating that $\theta^H_{\tau(\hat{\theta})}$ is the lowest type who purchases the good. Lemma 1.1 shows that since the ex-ante participation constraint is satisfied for $\tau(\hat{\theta})$, it can't bind for any higher $\tau \in [\tau(\hat{\theta}), \tau(y)]$.

Lemma 1.1. *If for signal types* $0 \leq \tau_1 < \tau_2 \leq \tau(y)$ *and some contract* (a, p) *with* $p > \theta^L_{\tau_1}$ *we have*

$$-a + \mathbb{P}(\theta^H_{\tau_1} | \tau_1)(\theta^H_{\tau_1} - p) \geq 0, \qquad \text{(IR } \tau_1\text{)}$$

then we have

$$-a + \mathbb{P}(\theta^H_{\tau_2} | \tau_2)(\theta^H_{\tau_2} - p) > 0. \qquad \text{(IR } \tau_2\text{)}$$

proof of Lemma 1.1. Call $\alpha_1 := \mathbb{P}(\theta^H_{\tau_1} | \tau_1)$ and $\alpha_2 := \mathbb{P}(\theta^H_{\tau_2} | \tau_2)$.
We thus need to show that

$$\alpha_1(\theta^H_{\tau_1} - p) < \alpha_2(\theta^H_{\tau_2} - p)$$

If $\alpha_2 > \alpha_1$ this is immediate, since $\theta^H_{\tau_2} > \theta^H_{\tau_1}$. Assume therefore in the following that $\alpha_2 \leq \alpha_1$.
Equation (1.5) can be rewritten as

$$(1 - \alpha_1)\theta^L_{\tau_1} + \alpha_1 \theta^H_{\tau_1} = \overline{a},$$

or respectively

$$(1 - \alpha_2)\theta^L_{\tau_2} + \alpha_2 \theta^H_{\tau_2} = \overline{a}.$$

It follows that

$$\alpha_1(\theta^H_{\tau_1} - \theta^L_{\tau_1}) = \overline{a} - \theta^L_{\tau_1} = (\overline{a} - \theta^L_{\tau_2}) + (\theta^L_{\tau_2} - \theta^L_{\tau_1}) = \alpha_2(\theta^H_{\tau_2} - \theta^L_{\tau_2}) + (\theta^L_{\tau_2} - \theta^L_{\tau_1}).$$

Now, since $\theta^L_{\tau_2} < \theta^L_{\tau_1} < p$ and $\alpha_2 < 1$, we have

$$\begin{aligned}\alpha_1(\theta^H_{\tau_1} - p) &= \alpha_1(\theta^H_{\tau_1} - \theta^L_{\tau_1}) + \alpha_1(\theta^L_{\tau_1} - p)\\ &= \alpha_2(\theta^H_{\tau_2} - \theta^L_{\tau_2}) + (\theta^L_{\tau_2} - \theta^L_{\tau_1}) + \alpha_1(\theta^L_{\tau_1} - p)\\ &< \alpha_2(\theta^H_{\tau_2} - \theta^L_{\tau_2}) + \alpha_2(\theta^L_{\tau_2} - \theta^L_{\tau_1}) + \alpha_2(\theta^L_{\tau_1} - p)\\ &= \alpha_2(\theta^H_{\tau_2} - p).\end{aligned}$$

□

Further, any type $\theta > y$, who learns his type with certainty under τ, obtains a utility of

$$u_B = -a + (\theta - p) > -a + (\hat\theta - p) > -a + \mathbb{P}(\hat\theta|\tau(\hat\theta))(\hat\theta - p) \geq 0$$

from contract (a, p). The contract thus generates positive expected utility to all $\tau > \tau(\hat\theta)$, and positive utility to all types $\theta > \hat\theta$. This means that the contract (a, p) alone induces all types $\theta \geq \hat\theta$ to purchase the good. Since, by assumption, $\hat\theta$ is the lowest type who purchases the good for menu $\tilde{\mathcal{M}}$, any additional contract in the menu does not increase trade efficiency. It could therefore only decrease seller utility, since a buyer would only take it if it yielded higher rents to him than the contract (a, p), and thus lower rents to the seller. Therefore, if $\tilde{\mathcal{M}}$ is an optimal menu, we can assume $\tilde{\mathcal{M}} = \{(a, p)\}$, and seller utility is given by

$$u_S = (1 - G(\tau(\hat\theta)))a + (1 - F(\hat\theta))p \quad = (1 - F(\hat\theta) + F(\theta^L_{\tau(\hat\theta)}) - F(x))a + (1 - F(\hat\theta))p.$$

Since by ex ante IR we have $a \leq \mathbb{P}(\hat\theta|\tau(\hat\theta))(\hat\theta - p)$, it follows that

$$u_S \leq \big(1 - F(\hat\theta) + F(\theta^L_{\tau(\hat\theta)}) - F(x)\big)\mathbb{P}(\hat\theta|\tau(\hat\theta))(\hat\theta - p) + (1 - F(\hat\theta))p.$$

Recall that $0 \leq \theta^L_{\tau(\hat\theta)} < p \leq \hat\theta$, since $\hat\theta$ is the lowest type who buys. If

$$\big(1 - F(\hat\theta) + F(\theta^L_{\tau(\hat\theta)}) - F(x)\big)\mathbb{P}(\hat\theta|\tau(\hat\theta)) > 1 - F(\hat\theta),$$

then

$$\begin{aligned} u_S &\leq \big(1 - F(\hat{\theta}) + F(\theta^L_{\tau(\hat{\theta})}) - F(x)\big)\mathbb{P}(\hat{\theta}|\tau(\hat{\theta}))\hat{\theta} \\ &\leq (1 - F(x))\mathbb{P}(\hat{\theta}|\tau(\hat{\theta}))\hat{\theta} \\ &\leq (1 - F(x))\big(\mathbb{P}(\theta^H_{\tau(\hat{\theta})}|\tau(\hat{\theta}))\hat{\theta} + \mathbb{P}(\theta^L_{\tau(\hat{\theta})}|\tau(\hat{\theta}))\theta^L_{\tau(\hat{\theta})}\big) \\ &= (1 - F(x))\overline{a} \\ &= SS. \end{aligned}$$

Alternatively, if

$$\big(1 - F(\hat{\theta}) + F(\theta^L_{\tau(\hat{\theta})}) - F(x)\big)\mathbb{P}(\hat{\theta}|\tau(\hat{\theta})) \leq 1 - F(\hat{\theta}),$$

then

$$\begin{aligned} u_S &\leq \big(1 - F(\hat{\theta}) + F(\theta^L_{\tau(\hat{\theta})}) - F(x)\big)\mathbb{P}(\hat{\theta}|\tau(\hat{\theta}))(\hat{\theta} - p) + (1 - F(\hat{\theta}))p \\ &= (1 - F(\hat{\theta}))(\hat{\theta} - p) + (1 - F(\hat{\theta}))p \\ &= (1 - F(\hat{\theta}))\hat{\theta} \\ &\leq \max_p (1 - F(p))p \\ &= \pi^M \\ &\leq SS \end{aligned}$$

This concludes the proof that there is no menu $\tilde{\mathcal{M}}$ which yields the seller a surplus above SS. Consequently $\mathcal{M}$ is a seller-optimal contract. □

Chapter 2

English versus Vickrey Auctions with Loss Averse Bidders

This Chapter is based on von Wangenheim (2017b).

2.1 Introduction

Auctions are a universal tool to organize sales in markets. At the core of auction theory stand the famous revenue equivalence results. In particular, Vickrey (1961) notes the strategic equivalence between the dynamic English and the static Vickrey auction: if values are independent and private, there is no effect of sequential information and it is a weakly dominant strategy to bid (up to) one's private valuation in both formats.[1] These powerful theoretical predictions, however, stand in contrast to the experimental literature, which mostly finds lower revenues for the English auction.[2] I identify endogenous preferences in form of expectation-based loss aversion as a possible explanation for this phenomenon.

In my model, bidders evaluate any auction outcome relative to their reference point, formed by rational expectations. Consequently, neither in the second-price (Vickrey), nor in the ascending-clock (English) auction it is optimal any more to bid (up to) the own intrinsic valuation. In particular, loss aversion leads to strong overbidding for high types in the Vickrey auction. Moreover, if agents update their reference point with respect to new information, opponents' behavior influences bidders' reference point, and thus their endogenous preferences. Hence,

[1] Myerson (1981) extends the results to show that all main auction formats give rise to the same expected revenue.

[2] For a summary of the experimental literature, see Kagel (1995).

even if valuations for the object are entirely private, sequential information affects the bidding behavior. Consequently, the English and the Vickrey auction are no longer strategically equivalent. I demonstrate that—consistent with most of the experimental evidence—the English auction yields lower revenue. I establish that this effect is driven by a time-inconsistency problem, which dynamic expectation-based loss averse bidders face when forming their bidding strategy.

Following the concept of loss aversion by Kőszegi and Rabin (2006), I assume bidders experience—in addition to classical utility —gain-loss utility from comparing the outcome to their expectations. Further, I assume that bidders assign gains and losses separately to money and good (narrow bracketers). For the ease of exposition, I consider mostly bidders who are only loss averse with respect to the object.[3] If they win the auction, they experience a feeling of elation, increasingly in the extent to which winning was unexpected. Similarly, they perceive a feeling of loss if they lose, increasingly in their expectations to win. Taking that into account, bidders will overbid their intrinsic valuation. Since losses with respect to expectations weigh stronger than gains, high types—who expect to win—overbid more aggressively than low types in the Vickrey auction.

To model the impact of dynamic information on the reference point in the dynamic English auction, I take the continuous-time limit of Kőszegi and Rabin (2009): every clock increment bidders observe whether opponents drop out from the auction. This information permanently updates expectations about winning the auction and about how much to pay. If the changes in beliefs immediately update the bidders' reference points, they instantaneously perceive gain-loss utility, which means that they assign gains and losses to changes in the reference distribution.

I consider the two extreme cases as benchmarks: if the reference-point updating is sufficiently lagged with respect to changes in beliefs, there is no updating during the auction process and therefore no impact of sequential information. The English auction remains equivalent to the Vickrey auction in that case.

If the new information immediately updates the reference point, however, bidders' utility depends on the observed signals about opponents' bidding strategies during the auction process, even though values are private.

Kőszegi and Rabin interpret the reference point as lagged beliefs. Recent experimental findings, however, suggest that the reference point adjusts quickly to new information. Whether instantaneous reference-point updating is a realistic approximation may depend on the exact auction environment, e.g. the speed at which the price augments, which can differ immensely across different English auc-

[3] I show in section 2.6.1 that the main insights generalize to the case where bidders assign gains and losses separately to the money and good dimension.

tions. Altogether, instantaneous updating constitutes a natural and important benchmark.

Since losses weigh stronger than gains, expected reference dependent utility is always negative. In particular, bidders dislike fluctuation in beliefs. As bidders are forward looking, they will account for these costs when they form their bidding strategy. In principle, an aggressive bid would to some extent insure against belief fluctuations during the auction process. However, as the auction prevails, bidders' beliefs to win the auction eventually decline. They become less attached to the auctioned object, and at the point they would have to bid aggressively, it is time inconsistent to do so. They eventually perceive themselves as a low type with respect to the active bidders in the remaining auction. This leads to only moderate overbidding - similarly as for low types in the Vickrey auction. Therefore, bidding is less aggressive in the English auction with updated reference points.

Since bidders dislike belief fluctuations, they would prefer to refrain from observing the auction process and rather use proxies to bid on their behalf. The logic is related to Benartzi et al. (1995) and Pagel (2016), who explain the equity premium puzzle by loss aversion: since stock prices fluctuate, an investor who regularly checks her portfolio will experience negative reference-dependent utility in expectation. This disutility makes stocks relatively less attractive to bonds.

Lange and Ratan (2010) highlight that in the presence of loss aversion in hedonic dimensions most laboratory results may not be transferable to the field: the effects of loss aversion are mainly driven by the assumption that bidders account losses and gains separately in the money and the good dimension (narrow bracketing). In order to control for private values, most auction experiments, however, use auction tokens, which can be interchanged for money at the end of the experiment. In context of these induced value experiments, bidders might not evaluate gains and losses to tokens and money separately, as they are in fact both money.[4] Since I assume narrow bracketing throughout this chapter, my results are more likely to apply to real commodity auctions, rather than to experiments on induced value auctions. It can therefore explain the revenue gap between the Vickrey auction and the English auction in the induced-value experimental literature, only if we assume that bidders don't perceive the tokens as money.

There is surprisingly little experimental literature that compares revenues of the English auction and the Vickrey auction for real commodities.[5] The only

[4]Indeed, Shogren et al. (1994) run Vickrey auctions to sell different goods and show that an endowment effect is strongest for non-market goods with imperfect substitutes.

[5]The only field experiment I am aware of is conducted by Lucking-Reiley (1999), who trades magic cards on an internet auction platform. He finds no significant difference in revenues, though he admits himself that he cannot entirely control for a potential selection bias and endogenous entry.

laboratory controlled experiment that I am aware of, is conducted by Schindler (2003). She reports 14 percent lower revenues in the English auctions, therefore confirming the findings of the induced-value literature, as well as my theoretical predictions.

The remainder of the Chapter is structured as follows: Section 2.2 discusses the related literature, Section 2.3 analyzes the Vickrey auction with loss averse bidders, while Section 2.4 analyzes the English auction with loss averse bidders. In Section 2.5, I discuss the revenue comparison of both auction formats. Section 2.6 discusses several extensions, while Section 2.7 concludes this chapter. All proofs are relegated to the appendix.

2.2 Related Literature

Kahneman et al. (1990) establish the *endowment effect* that agents' valuation for goods increase with ownership. It has since been experimentally replicated under many different circumstances, for summaries see Camerer (1995) and Horowitz and McConnell (2002). Tversky and Kahneman (1991) propose loss aversion with respect to the status quo to explain the endowment effect.

Kőszegi and Rabin (2006) suggest recent rational expectations as reference point. The hypothesis that expectations play a role in individual's preferences have been supported in recent experiments (Ericson and Fuster (2011) and Abeler et al. (2011)), as well as challenged (Heffetz and List (2014)).[6]

The idea that the reference point is determined by recent beliefs leads to the natural question of the speed of reference-point adjustment. Strahilevitz and Loewenstein (1998) provide early evidence that the time span for which individuals hold beliefs has an impact on the reference point. Gill and Prowse (2012) use a real effort task to measure loss aversion and find that in their framework "the adjustment process is essentially instantaneous". Smith (2012) induces different probabilities of winning an item across groups of individuals. After the uncertainty resolves, he measures the willingness to pay for the item among bidders who have not won. In contrast to Ericson and Fuster (2011), who elicit valuations *before* the uncertainty resolves, Smith finds no significant difference between different groups, which suggests that the reference point is not so much determined by lagged beliefs, but rather adjusts quickly to the new information.[7]

For static environments Kőszegi and Rabin (2006) has arguably become the standard model of reference-dependent preferences, and been successfully applied to

[6] For a literature revue on related evidence, see Ericson and Fuster (2014).

[7] Smith's confidence intervals are, however, rather wide.

various fields, like mechanism design (Eisenhuth (2012)), contract theory (Herweg et al. (2010)), industrial organization (Heidhues and Kőszegi (2008)), and labor markets (Eliaz and Spiegler (2014)). Heidhues and Kőszegi (2014) show that buyers in monopolistic markets may face a similar form of time inconsistency as I establish for bidders in the English auction: ex ante they would like to commit not to buy. If the seller induces low prices with some probability, this plan, however, is time inconsistent. As a result, the consumer ends up buying for a high prices as well. Dato et al. (2017) extend the equilibrium concepts of Kőszegi and Rabin (2006) to strategic interaction in static games.

In the context of auctions with reference-dependent preferences, Lange and Ratan (2010) point out that loss-averse bidders may behave differently in laboratory experiments than in the field; bidders may not bracket narrowly in induced-value experiments. Further, they calculate the equilibrium bidding function of loss averse bidders in the first-price auction and Vickrey auction for a different equilibrium concept than I use in this chapter. (For a more detailed discussion of the equilibrium concepts see section 3.)

Ehrhart and Ott (2014) introduce a model of the Dutch and English auction, where sequential information updates the reference point, but—in contrast to Kőszegi and Rabin (2009)—does not induce gain-loss utility. As a result, in equilibrium there is never any feeling of loss in the English auction, since by the time a bidder drops out she expects to lose. Eisenhuth and Ewers (2010) show that the all-pay auction yields higher payoffs than the first-price auction for narrow-bracketing bidders, since loss-averse bidders dislike payment uncertainty.

For dynamic environments Kőszegi and Rabin (2009) propose a model of dynamic loss aversion, where updates of expectations carry reference-dependent utility. This model has so far only been applied sparsely. First applications nevertheless seem promising. Rosato (2014) uses a two-period dynamic model to show that revenues are decreasing in sequential auctions with loss-averse bidders, due to a discouragement effect. To my best knowledge, Pagel is the first to rigorously apply Kőszegi and Rabin (2009) to dynamic problems with a long time horizon. Pagel (2016) shows that dynamic reference-dependent preferences can explain the historical levels of equity premiums and premium volatility in asset prices. Related to the logic in the English auction, loss-averse agents dislike price fluctuations, which makes assets relatively unattractive. Pagel (2017) shows that dynamic reference-dependent preferences can explain empirical observations about saving schemes for life-cycle consumption.

To my best knowledge, my model is the first to analyze strategic interaction of loss-averse players in a dynamic game with more than two periods.

2.3 The Vickrey Auction

2.3.1 Auction Rules

The second-price auction or Vickrey auction is a static, sealed-bid auction format. We assume that there are N loss averse bidders participating in the auction for a non-divisible good. Bidder i's valuation θ_i is privately observed and independently drawn from a common distribution

$$\theta_i \sim G,$$

where G has a differentiable density g which is strictly positive on its support $[\theta^{\min}, \theta^{\max}]$, with $0 \leq \theta^{\min} < \theta^{\max}$. After learning their private valuation, every participant submits a sealed bid. The bidder with the highest bid is assigned the object and has to pay the amount of the second highest bid.

2.3.2 Preferences

I assume that bidders are loss averse in the sense of Kőszegi and Rabin (2006). In addition to classical utility from an endowment $x \in \mathbb{R}$, bidders perceive a feeling of gain or loss, depending on whether the endowment is higher or lower than their reference point $r \in \mathbb{R}$. If we assume that classical utility is linear in x, this means:

$$u(x|r) = x + \mu(x - r),$$

where μ characterizes the gain-loss utility. In the Vickrey auction there are two commodity dimensions—money and good. We assume that bidders are narrow bracketers: utility is additively separable and gains and losses are evaluated separately across the two different dimensions: for any endowment level $x = (x^m, x^g)$ and any reference level $r = (r^m, r^g)$, agents utility is given by

$$u(x|r) = \sum_{k \in \{m,g\}} x^k + \mu_k(x^k - r^k),$$

where we allow for different loss specifications across dimensions.

If the bidder is loss averse she will weigh losses with respect to her reference point stronger than gains. Following Section IV in Kőszegi and Rabin (2006) and most of the literature, I assume μ_k to be a piecewise linear function with a kink at zero:

$$\mu_k(y) = \begin{cases} \eta_k y, & y \geq 0, \\ \lambda_k \eta_k y, & y < 0, \end{cases}$$

where $\eta_k > 0$, $\lambda_k > 1$, and $\Lambda_k := \lambda_k\eta_k - \eta_k < 1$ for $k \in \{m, g\}$.[8] Because it suffices for demonstrating the novel economic effect and allows for a significantly simpler exposition, I first focus on the case in which bidders are loss averse in the good dimension only, i.e. $\eta_m = 0$. In the extensions, I show that my results generalize to the case where we allow for loss aversion in the money dimension as well.[9]

The key feature in Kőszegi and Rabin (2006) is that the reference point is stochastic and endogenously determined by rational beliefs over future endowment levels. Consider an agent, who faces an uncertain payoff of x in some commodity dimension, which is distributed according to some distribution F. Let the reference point be determined by the agent's beliefs F' about the outcome. A realization $\overline{x}$ of x then yields an ex post utility in this commodity dimension of

$$u(\overline{x}|F') = \overline{x} + \int \mu(\overline{x} - r)dF'(r).$$

Then the ex-ante expected utility of the endowment x is given by

$$U(F|F') := \mathbb{E}u(x, |F') = \int \left(x + \int \mu(x - r)dF'(r)\right) dF(x).$$

If the agent has rational expectations, we have $F = F'$, and the expected utility of the lottery is

$$U(F|F) := \mathbb{E}u(x, |F) = \int \left(x + \int \mu(x - r)dF(r)\right) dF(x).$$

2.3.3 Equilibrium Concept

I adapt Kőszegi and Rabin's equilibrium concept under uncertainty to allow for interactive decision problems. I take an interim approach in the sense that each bidder i forms her strategy *after* she learns her private valuation θ_i. Fixing all opponents' behavior, we summarize their strategy in the distribution H of the maximal opponent bid. Given θ_i and H, any bid b induces some distribution of auction outcomes and therefore payoff distribution $F^k = F^k(b, \theta_i, H)$ in the respective commodity dimensions $k \in \{m, g\}$.

Definition 2.1. A bid $b^* \in \mathbb{R}_+$ constitutes an *unacclimated personal equilibrium*

[8] The condition $\Lambda < 1$ is referred to as "no dominance of gain-loss utility" by Herweg et al. (2010) It ensures that the dislike for uncertainty isn't too strong. If $\Lambda > 1$ a bidder could potentially prefer a strictly dominated safe outcome to a lottery.

[9] Horowitz and McConnell (2002) conclude in their summary that the endowment effect is "highest for non-market goods, next highest for ordinary private goods, and lowest for experiments involving forms of money." In this sense it may be plausible that loss aversion mainly applies to the good dimension.

(UPE) in the Vickrey auction for bidder i if for all $b \in \mathbb{R}_+$,

$$\sum_{k\in\{m,g\}} U(F^k(b^*,\theta_i,H)|F^k(b^*,\theta_i,H)) \geq \sum_{k\in\{m,g\}} U(F^k(b,\theta_i,H)|F^k(b^*,\theta_i,H)).$$

In other words b^* is a UPE if, given the reference point generated by the action b^*, there is no profitable deviation b. It contrasts the definition of a *choice-acclimating personal equilibrium (CPE)*, where we require

$$\sum_{k\in\{m,g\}} U(F^k(b^*,\theta_i,H)|F^k(b^*,\theta_i,H)) \geq \sum_{k\in\{m,g\}} U(F^k(b,\theta_i,H)|F^k(b,\theta_i,H))$$

for all $b \in \mathbb{R}_+$. Thus, in contrast to the UPE-bidder, a CPE-bidder—which is analyzed in Lange and Ratan (2010)—already internalizes the effects of her deviation on the reference point. I believe the UPE is the appropriate equilibrium concept in the Vickrey auction, mainly for two reasons.

Firstly, I apply the model as proposed by Kőszegi and Rabin who suggest that the UPE is more appropriate if the bidder "anticipates the decision she faces but cannot commit to a choice until shortly before the outcome" (Kőszegi and Rabin (2007)). In auction settings bidders may know her valuation and form expectations long before the auction starts. Bids are, however, typically placed only shortly before the auction uncertainty resolves, and may depend on characteristics specific to the environment, such as the number of bidders actually participating in the auction.

Secondly, the UPE is the static analog of the concept of a personal equilibrium, which will be introduced in Section 4 to analyze the dynamic English auction. In this context one can gather another (dynamic) interpretation for the UPE: the decision maker ex ante forms a plan before the auction actually starts. This plan will determine her reference-point. The plan is a UPE if it is time-consistent in the sense that the decision maker is willing to carry it through at the time of action.

In a joint equilibrium, the first order statistic of the $n-1$ opponent bids H is endogenously determined by the equilibrium bidding strategy. Thus, if $b(\theta)$ constitutes a symmetric increasing equilibrium bidding function, we necessarily have

$$H(b(\theta)) = G^{n-1}(\theta).$$

Definition 2.2. In the Vickrey auction with n loss averse bidders, an increasing function $b(\theta)$ constitutes a symmetric UPE if for all θ and all b'

$$\sum_{k\in\{m,g\}} U(F^k(b(\theta),\theta,G^{n-1}(b^{-1}(\cdot)))|F^k(b(\theta),\theta,G^{n-1}(b^{-1}(\cdot))))$$
$$\geq \sum_{k\in\{m,g\}} U(F^k(b',\theta,G^{n-1}(b^{-1}(\cdot)))|F^k(b(\theta),\theta,G^{n-1}(b^{-1}(\cdot)))).$$

2.3.4 The Equilibrium

In this section we restrict attention to agents who are loss averse only in the good dimension. A more elaborate analysis of the general case, which allows for loss aversion in the money dimension is relegated to the extensions. Consider a bidder of type θ who plans to submit a bid of b^*. Given the distribution H of the highest opponent bid, the plan induces the reference distribution to win a utility of θ with probability of $H(b^*)$. Suppressing some notation, the utility of bidding b if planning to bid b^* is given by

$$\begin{aligned} u(b,\theta|b^*) &:= \sum_{k\in\{m,g\}} U(F^k(b,\theta,H)|F^k(b^*,\theta,H)) \\ &= \underbrace{H(b)}_{\text{Prob to win}} \big(\theta + \underbrace{(1-H(b^*))\mu(\theta)}_{\text{feeling of gain}}\big) + \underbrace{(1-H(b))}_{\text{Prob to lose}}\underbrace{H(b^*)\mu(-\theta)}_{\text{feeling of loss}} + \underbrace{\int_0^b -sH(s)}_{\text{money dimension}} \\ &= \underbrace{\int_0^b (\theta-s)dH(s)}_{\text{classical utility}} + \underbrace{H(b)(1-H(b^*))\mu(\theta) + (1-H(b))H(b^*)\mu(-\theta)}_{\text{total reference-dependent utility}}. \end{aligned}$$

In any symmetric equilibrium, $b = b^*$ must be the utility maximizing bid, where H is given by opponents' symmetric bidding behavior.

Theorem 2.1. *The unique symmetric increasing continuously differentiable UPE in the Vickrey auction with n bidders who are loss averse with respect to the good is given by*

$$b(\theta) = \big(1 + \eta(1 - G^{n-1}(\theta)) + \lambda\eta G^{n-1}(\theta)\big)\theta.$$

Note that all types overbid with respect to their intrinsic valuation θ. This should not be too surprising since we have assigned loss aversion only to the good dimension, and therefore made the good relatively more important, compared to money. More interestingly, the degree of overbidding is increasing in the type. The lowest type moderately overbids by

$$b(\theta^{\min}) = (1 + \eta)\theta^{\min},$$

while the highest type aggressively overbids by

$$b(\theta^{\max}) = (1 + \lambda\eta)\theta^{\max}.$$

The reason is the so called attachment effect: high types believe to win. Not winning would create a feeling of loss, which they try to prevent by placing an aggressive bid. As we will see section 2.6.1, this intuition remains intact, if we allow for loss aversion in money as well.

2.4 The English Auction

2.4.1 The Model

The Auction Format

The English auction format I am considering is sometimes referred to as the "Ascending Clock Auction" or the "Japanese Auction". In contrast to the "Open Outcry Auction" bidding starts at a fixed price and is raised incrementally by the auctioneer each time period. Each bidder signals—for example by raising or dropping her hand—when she wishes to drop out of the auction. Once a bidder dropped out she cannot bid again. The auction ends if there is only one active bidder left. This bidder has to pay the price, at which the last of her opponents dropped out.

For simplicity, we assume that there is no reservation price—the clock starts with a price of zero. The effect of a reserve price is analyzed in extension 2.6.3.

Preferences

We assume that bidders' intrinsic valuations θ_i for the object are privately observed and independently drawn

$$\theta_i \sim G$$

from a distribution G that has a differentiable and strictly positive density g on a positive support $[\theta^{\min}, \theta^{\max}]$. The distribution G is common knowledge. Bidders are assumed to be loss averse.

I follow Kőszegi and Rabin (2009) in how to model loss aversion in a dynamic discrete-time environment: agents hold rational beliefs about winning the auction and the respective transfers made after the auction is over. Every period, the agent observes, whether any opponents drop out at the current price and thus receives an information signal about the outcome. We denote by F_t^k, the beliefs over final transfers in $k \in \{money, good\}$, as anticipated at time t. As the signal at any time t changes beliefs over the auction outcome, this instantaneously gives rise to psychological gain-loss utility, denoted by $N(F_t^k|F_{t-1}^k)$, separately to changes in money and good.

For the evaluation of gain-loss utility, agents are assumed to assign gains and losses to changes in the respective quantiles of the distribution function. The intuition is that the agents rank possible outcomes from worst to best and then evaluate changes to the worst, the second worst ,..., until the best outcome. Let us denote with $c_{F_t^k}$ the quantile function of F_t^k, which is mathematically just the inverse of F_t^k. Then

$$N(F_t^k|F_{t-1}^k) = \int_0^1 \mu_k(c_{F_t^k}(p) - c_{F_{t-1}^k}(p))dp,$$

where again

$$\mu(x) = \begin{cases} \eta_k \cdot x, & x \geq 0, \\ \eta_k \cdot \lambda_k \cdot x, & x < 0, \end{cases}$$

$\eta_k > 0$, $\lambda_k > 1$, and $\Lambda_k := \lambda_k \eta_k - \eta_k < 1$.
In other words, during the auction process bidders accumulate information about the auction outcome. They absorb this information in their reference-point, which instantaneously exposes them to (possibly mixed) feelings of gains and losses. The total utility perceived in the auction process is given by the accumulated gain-loss utility and the classical utility from trade if the auction is won. In the following analysis, it is convenient to index the distributions with the current price rather than with the time period. After learning her type θ_i, bidder i forms a bidding strategy, which induces beliefs F_0^m and F_0^g about the auction outcome. If the auction runs for at most T increments of ε, we can write the total utility of the auction as

$$u_i = \sum_{t=1}^{T} \left(N(F_{t\varepsilon}^m|F_{(t-1)\varepsilon}^m) + N(F_{t\varepsilon}^g|F_{(t-1)\varepsilon}^g) \right) + (\theta_i - x)$$

if bidder i wins the auction at a price of x, and as

$$u_i = \sum_{t=1}^{T} \left(N(F_{t\varepsilon}^m|F_{(t-1)\varepsilon}^m) + N(F_{t\varepsilon}^g|F_{(t-1)\varepsilon}^g) \right)$$

if bidder i loses the auction. Note that the upper bound of T in the sum is without loss of generality; if the auction terminates early, all subsequent periods can be regarded as uninformative, and carry no further reference-dependent utility.

Equilibrium Concepts

I concisely sketch the equilibrium concept of Kőszegi and Rabin (2009). For full details and a psychological justification of the specific dynamic modeling choices, I refer to their paper.

Definition 2.3. An **action plan** specifies an action for every realization of information at every point in time. An action plan constitutes a **personal equilibrium** (PE) if, given the reference point resulting from the plan, it maximizes expected utility at any point in time among all plans that the agent is willing to carry through.

This means in particular:

- The bidder can only make credible plans in the sense that she cannot commit to plans that her future self does not want to carry through at the time of actions. Committing to unfavorable actions could be profitable, because it would manipulate beliefs, and therefore the own reference point.
- In suppressed notation, an action plan that induces a distribution F is an equilibrium if and only if at any point in time $u(F_t|F_t) \geq u(F_t'|F_t)$ for any distribution F_t' that would result from another credible plan.
- Given the opponents' behavior, an agent determines her set of personal equilibria by backward reasoning: she evaluates any action in $T-1$ with respect to her optimal actions in period T, and proceeds backwards.

The only constraint on initial beliefs is that they are rational, given the action plan. In general, there may be multiple personal equilibria.

Definition 2.4. A personal equilibrium is a **preferred personal equilibrium** if it is the utility maximizing PE at time zero.

The set of personal equilibria depends on the belief about other players' actions. To analyze the interaction between multiple bidders, we focus on symmetric personal equilibria.

Definition 2.5. A **strategy** $b(\theta)$ assigns to each possible type θ an action plan. A strategy constitutes a **(preferred) symmetric equilibrium** in the English auction if for each type θ and the belief that all opponents bid according to strategy b the action plan $b(\theta)$ constitutes a (preferred) personal equilibrium.

Timing

First bidders privately learn their valuation θ_i for the object. Then each bidder forms an action plan, which prescribes for any time (clock price) and any opponent drop-out history, the decision whether to drop out or to remain. Rational beliefs induced by this action plan form the bidder's reference point. Finally, the auction takes place. Any period during the auction process is characterized by the following timing:

- The price on the clock ascends and bidders simultaneously signal whether they stay in or drop out. If a bidder deviates from her action plan, she updates her reference point according to new rational beliefs. The update instantaneously induces reference-dependent utility

- Bidders observe, whether opponents drop out and update their reference point about payoffs. The update instantaneously gives rise to gain-loss utility.

- If there is at most one bidder remaining active, the auction is terminated. The remaining bidder is assigned the object and pays the current clock price.[10]

2.4.2 Analysis

Illustrative Example of Updating

This example aims to provide an illustration how gain-loss utility is formed during the auction process, and to show why bidders would always prefer a proxy to bid on their behalf in the English auction—taken behavior of opponents as given.

Consider an English auction with two bidders. Let bidder 1—in the following referred to as the bidder—have a valuation of θ for the object. Assume that the bidder plans to drop out at a price of 8 and knows that the drop-out price of bidder 2 — in the following called opponent—is ex ante uniformly distributed on $[0, 10]$ (we do not consider here, under which circumstances this behavior would be optimal). Ex ante, the bidder has a probability of 0.8 to win the auction and to have a payoff of θ in the good dimension. In the money dimension she faces a probability of 0.2 to pay nothing. Prices between 0 and 8 are uniformly distributed and have a mass of 0.8 all together (if we assume arbitrary small increments on the clock for mathematical convenience). Thus, the ex ante quantile functions are given by

[10]For mathematical convenience, I abstract from tie breaking rules and assume that the good is not sold, if the remaining bidders drop out simultaneously. With our assumption of continuous density of types, as we let the increment size go to zero, this becomes equivalent to a tie breaking rule by coin-flip.

$$c_{F_0^g}(p) = \begin{cases} 0, & p \leq 0.2, \\ \theta, & p > 0.2 \end{cases}$$

in the good dimension, and

$$c_{F_0^m}(p) = \begin{cases} -8 + 10p, & p \leq 0.8, \\ 0, & p > 0.8, \end{cases}$$

in the money dimension.

Assume the opponent drops out at a price of 6. While the clock price ascends, the bidder permanently updates her beliefs. Let us look at the good dimension: for any increment below the price of 6, the bidder realizes that the opponent didn't drop out at that price, which reduces her beliefs to win the auction by some small amount. This means that during the auction process she accumulates perceived losses in the good dimension. Figure 2.1 shows the quantile functions at different clock prices.

At a clock price of 0—that is before the auction starts—the bidder holds her prior belief to win the auction with a probability of 0.8. The respective quantile function is a step function which is zero with probability 0.2, and θ with probability 0.8 (dotted line). At a price of 4 the bidder knows that the opponent hasn't dropped out between 0 and 4. Therefore bidder's updated belief to win is given by the probability that the opponent will drop out at price between 4 and 8, conditional on the fact that he will drop out between 4 and 10. It has thus decreased to two third which is indicated by the dashed quantile function. The medium grey shaded area is proportional to the loss the bidder has accumulated up to the price of 4 as the difference of the initial and current quantile function. Just before the opponent drops out at 6, bidder's belief has further decreased to almost one half—she wins if opponent drops out between 6 and 8, but loses if the opponent drops out between 8 and 10 (solid quantile function with jump at 0.5). The light shaded area shows the additional loss just before a price of 6 is announced. The losses have to be weighted with a factor of $\lambda\eta$. The moment the price increases to 6, the opponent drops out and the bidder wins with certainty. The quantile function jumps to the constant function $c_{F_6^g} = \theta$, inducing a feeling of gain of η times the three combined shaded areas.

Thus, the net gain-loss utility in the good dimension is $(0.2\eta + 0.3(\eta - \lambda\eta))\theta = (0.2\eta - 0.3\Lambda)\theta$. Since losses weight stronger than gains, the relief of winning the light gray and medium gray area after all can only partly make up for the disappointment felt during the auction process. If the bidder could use a bidding proxy that enabled

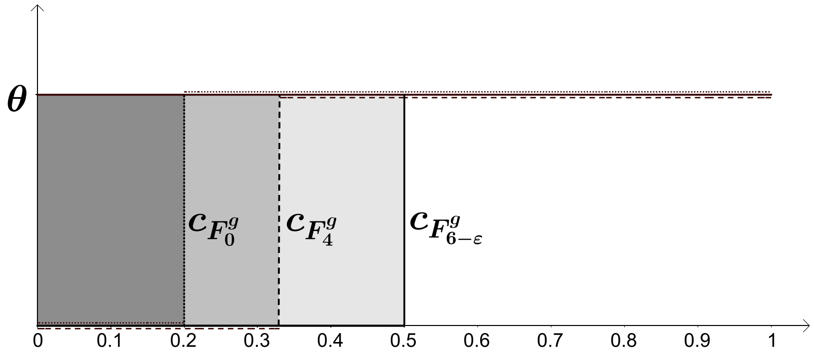

Figure 2.1: Updating in the English Auction

her to ignore new information until the auction was over, she would forgo the unpleasant variation in beliefs, which causes disutility of $-0.3\Lambda\theta$. This logic is due to Kőszegi and Rabin (2009), who find that, ceteris paribus, any collapse of information signals weakly increases agents' utility. Note that the use of a proxy in our two-bidder example is equivalent to submitting a sealed maximum bid. The example thus illustrates that——fixing her strategy and other bidders behavior—a loss averse bidder obtains weakly higher utility in the Vickrey auction than in the English auction.

The updating with respect to money is a bit more complex than the updating in the good dimension: if an opponent does not drop out at some price, the probability of losing and paying nothing increases as well as the probability of paying a high price. Nevertheless the same intuition applies: fluctuations in beliefs are costly, and loss-averse bidder would prefer to get all information at once. To summarize:

Corollary 2.1. *Loss-averse agents would prefer the use of proxies to bid on their behalf in the English auction. Thus, for a given set of bidders' maximal bids, any loss-averse bidder receives weakly higher utility in a Vickrey auction than in an English auction.*

Equilibrium Behavior for 2 Bidders

In the following, I analyze the set of equilibria in the English auction with two bidders, who are loss averse in the good dimension, as the increment size goes to zero. In section 2.6.4, I show that the main insights generalize to the n bidder auction. While the history-dependend strategy space in an n-bidder English auction is huge, it is fairly simple in a two-bidder game. Given type θ, an action plan prescribes the price at which the bidder plans to drop out, provided that the opponent is still active.

Each period the bidder observes whether her opponent remains in the auction. This information permanently updates her reference point, which induces gain-loss utility in each increment. An optimal bidding strategy will take the expected gain-loss utility from news into account.

For calculating the ex-ante expected gain-loss utility, it is more convenient to work with distribution functions rather than with quantile functions. This is possible, since they are inverse functions of each other, and the integral between functions equals the integral between their inverses up to the sign:

Lemma 2.1. *Let F_1 and F_2 be continuous distributions on an interval $[a, b]$ and let c_{F_1}, c_{F_2} be the respective quantile functions. Then*

$$\int_a^b (F_1(x) - F_2(x))dx = \int_0^1 (c_{F_2}(p) - c_{F_1}(p))dp.$$

With this result, one can look at the expected disutility from news.

Proposition 2.1. *Assume that a loss-averse bidder's payoff is distributed according to some distribution F_1 with a probability of Δ, and according to distribution F_2 with a probability of $1 - \Delta$. Let $[a, b]$ be the common support of F_1 and F_2. We denote with $F = \Delta F_1 + (1 - \Delta)F_2$ the ex ante distribution of the payoff. Then the ex ante expected reference-dependent utility from learning, whether the true distribution is F_1 or F_2, is given by*

$$\mathbb{E}(N(F_i|F) = -\Delta\Lambda \int_0^1 |c_{F_1}(p) - c_F(p)|dp,$$

or equivalently by

$$\mathbb{E}(N(F_i|F) = -\Delta\Lambda \int_a^b |F(x) - F_1(x)|dx.$$

The intuition for the result is as follows: on average, there is "as much good news as bad news". If gains and losses weighted equally, one would have zero gain-loss utility in average. Since losses loom larger than gains, variation will give us negative utility in expectation where the amount of negative utility is proportional to the expected variation and the loss dominance parameter Λ.

With this result we can calculate the accumulated expected loss due to gain-loss utility, as the increment size goes to zero. Let us denote with F the distribution of the opponent's drop-out price, in the sense that an opponent with drop-out price y remains in the auction at any clock price $t < y$, and drops out at prices $t \geq y$.

Proposition 2.2. *Consider a loss-averse bidder of type θ in the English auction with increments of ε and one opponent. Let the opponent's drop-out price be distributed according to distribution F with density f. Assume the bidder plans to drop out at x, and the opponent hasn't dropped out until time $t < x$. Then, for ε going to zero, in the limit the ex ante expected marginal gain-loss utility at time t is given by*

$$\ell_t(x, \theta, F) = \frac{-f(t)}{(1 - F(t))^2}(1 - F(x))\Lambda\theta.$$

Expected gain-loss utility for the remaining auction at time t is in the limit given by

$$L_t(x, \theta, F) = \ln\left(\frac{1 - F(x)}{1 - F(t)}\right)\frac{1 - F(x)}{1 - F(t)}\Lambda\theta.$$

Since losses weight stronger than gains, expected gain-loss utility is always negative. Note that the amount of marginal disutility is decreasing in x: an aggressive strategy induces less belief fluctuation at each information update, and thus partly insures against high gain-loss disutility in each increment. There is, however, a countervailing effect on total gain-loss disutility: the higher bidder's drop-out price, the longer she may stay in the auction and be exposed to gain-loss disutility. Figure 2.2 shows total expected gain-loss disutility at the beginning of the auction for $F \sim U[0, 1]$. We see that losses are the strongest for intermediate bids who face the highest uncertainty. Bidding 0 or 1 induces no uncertainty, and therefore no gain-loss utility.

In the following, we refer to the limit result as we let the increment size go to zero as the *continuous English auction.*[11]

With $L_t(x, \theta, F)$ we have established a function for the expected gain-loss utility *on the equilibrium path* for the strategy x. We now calculate the instantaneous gain-loss utility that the bidder perceives, if she decides to *deviate* from strategy x to strategy y at some point in time:

Lemma 2.2. *Consider a loss averse bidder in an English auction with one opponent. Let the opponent's drop-out price be distributed according to F. If at time t the bidder changes her strategy from dropping out at $x \geq t$ to dropping out at $y \geq t$, this deviation induces an instantaneous gain-loss utility of*

$$N(F_t^y | F_t^x) = \frac{\mu(F(y) - F(x))}{1 - F(t)}\theta.$$

Let us denote with $u_t(y, \theta, F|x)$ for $t \leq x, y$ the remaining expected utility of

[11] This notion does not intend to refer to the concept of *continuous games* by Simon and Stinchcombe (1989). One should still regard the game as one with discrete increments on the clock which are, however, arbitrarily small.

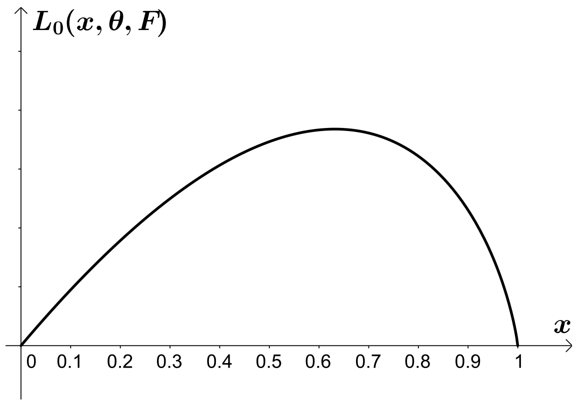

Figure 2.2: Total Expected Loss for $F \sim U[0,1]$

the agent at time t in the continuous English auction if she deviates at time t from strategy x to strategy y. Then, summarizing Proposition 2.2 and Lemma 2.2 we obtain

$$u_t(y,\theta,F|x) = \underbrace{\frac{\int_t^y(\theta-s)dF(s)}{1-F(t)}}_{\text{classical utility}} + \underbrace{\frac{\mu(F(y)-F(x))}{1-F(t)}\theta}_{\text{gain/loss from one-time update}} + \underbrace{L_t(y,\theta,F)}_{\substack{\text{expected gain-loss utility}\\ \text{of remaining auction}}}.$$

All three terms change if a bidder deviates to another strategy. Note that the deviation utility is non-differentiable at $y = x$, since μ has a kink at zero.

With this notation and the above results, we can restate the condition for a strategy to be an equilibrium as we let the increment size go to zero.

Corollary 2.2. *In the continuous English auction a bidding strategy x is a personal equilibrium if and only if*

$$u_t(y,\theta,F|x) \leq u_t(x,\theta,F|x)$$

for all $0 \leq t \leq x, y$ and all strategies y that are credible at all times $s > t$.

Since the equilibrium concept restricts to strategies x that the agent wants to carry through at any time, it is in particular necessary that the agent does not want to drop out just before x is reached. This leads to the following constraint on time consistent plans.

Lemma 2.3. *Consider a loss-averse bidder of type θ in the continuous English auction with one opponent. Let the opponent's drop-out price be distributed according to distribution F with nonzero density f on some positive support $[a, b]$. Then, any time consistent bidding strategy $x \in (a, b)$ satisfies*

$$x \leq (1 + \eta)\theta.$$

To understand the significance of this result, it is insightful to look at plans the bidder would choose if she could commit to a bidding strategy before the auction starts. She would not like to deviate from a strategy ex ante if and only if

$$u_0(y, \theta, F|x) \leq u_0(x, \theta, F|x)$$

for all y.

Proposition 2.3. *If two loss averse bidders could commit ex ante to a bidding strategy in the continuous English auction, the lowest symmetric increasing differentiable equilibrium would satisfy*

$$b(\theta) = \big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta)))\big)\theta.$$

Figure 2.3 shows the ex ante optimal strategy (solid function) and the boundary of time-consistent strategies (dashed line) for two loss averse bidders.

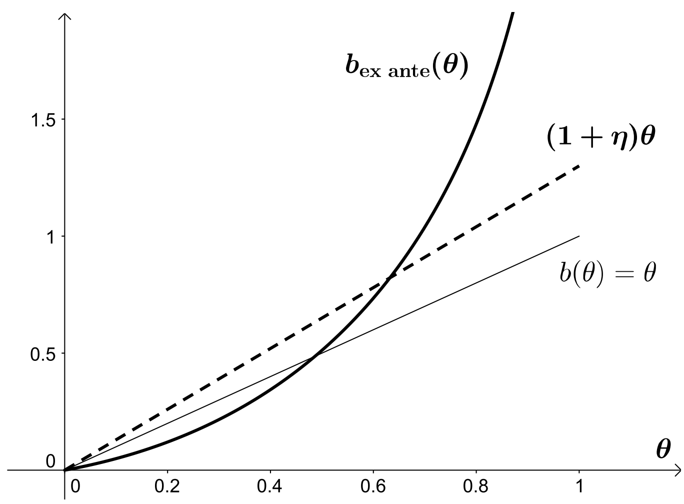

Figure 2.3: Ex-ante strategy and time-cons. constraint, $G(\theta) \sim U[0, 1], \eta = 0.3, \lambda = 4$

We see that low types ex ante may wish to underbid, while high types wish to strongly overbid. The intuition here is the same as in the Vickrey auction: bidders want to reduce expected gain-loss utility, and therefore try to reduce the uncertainty about winning. In particular high types would wish to insure with an aggressive bid against belief fluctuations during the auction process.

However, it is time-inconsistent to bid above $x = (1+\eta)\theta$. Even though a bidder with a high valuation would ex ante like to commit to an aggressive bidding strategy, at the time she has to do so, she is not any more willing to carry that action through: as the auction proceeds, the winning chances for the bidder gradually decline. Thus, she gradually becomes a low type with respect to the remaining auction, and therefore her initial strategy of overbidding becomes less appealing. Just one increment before the bidder's drop out, she perceives the remaining auction similarly as a Vickrey auction, where she has the lowest possible type. Hence, at that point in time, her optimal bidding strategy resembles that of the lowest type in the Vickrey auction, i.e. she bids no more than $x = (1+\eta)\theta$.

We have so far only considered constraints on equilibrium behavior at time 0 and at time x. It turns out that these are the binding constraints.

Lemma 2.4. *Consider a loss-averse bidder of type θ in the continuous English auction with one opponent. Let the opponent's drop-out price be distributed according to distribution F with nonzero density f on some positive support $[a, b]$. Then a strategy $x \in (a, b)$ is a PE if and only if*

1. $x \leq (1+\eta)\theta$;
2. *for any* $y \in [x, (1+\eta)\theta]$ *we have* $u_0(x, \theta, F|x) \geq u_0(y, \theta, F|x)$.

Theorem 2.2. *An increasing, almost everywhere differentiable function $b(\theta)$ is a symmetric equilibrium in the continuous English auction with two loss averse bidders if and only if for all θ*

1. $b(\theta) \leq (1+\eta)\theta$;
2. $b(\theta) \geq \min\left\{(1+\eta)\theta\ ;\ \big(1+\eta-\Lambda(1+\ln(1-G(\theta)))\big)\theta\ \right\}$.

Thus, any increasing smooth function in the the gray shaded area of Figure 2.4 constitutes a symmetric equilibrium.

The thick line indicates the preferred symmetric equilibrium (PPE). Point A, where the PPE hits the boundary of time consistent strategies can be easily determined:

$$\big(1+\eta-\Lambda(1+\ln(1-G(\theta)))\big)\theta = (1+\eta)\theta$$

if and only if $G(\theta) = 1 - 1/e \approx 0.632$.

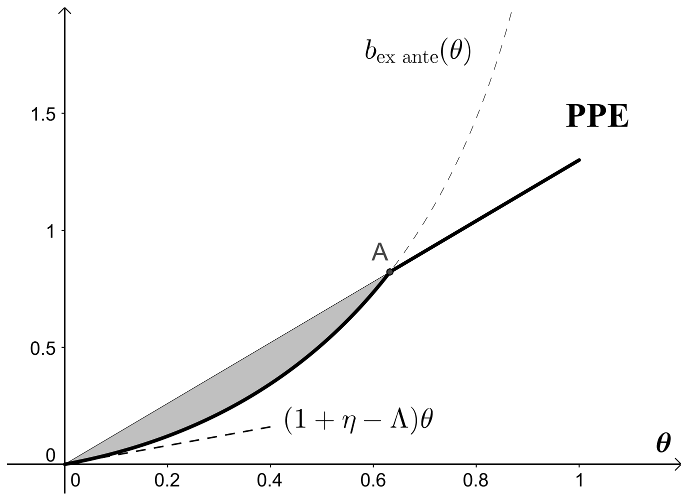

Figure 2.4: Equilibria in the English auction, $G(\theta) \sim U[0,1]$, $\eta = 0.3$, $\lambda = 4$

Note that the PPE is tangent to $(1+\eta-\Lambda)\theta$ at the lowest type. Hence there is underbidding for low types if and only if $\eta - \Lambda > 0$, thus if and only if $\lambda > 2$.

Corollary 2.3. *The symmetric PPE in the continuous English auction with two loss averse bidders is given by*

$$b_{PPE}(\theta) = \begin{cases} \big(1+\eta-\Lambda(1+\ln(1-G(\theta)))\big)\theta, & G(\theta) \leq 1-1/e, \\ (1+\eta)\theta, & G(\theta) > 1-1/e. \end{cases}$$

Low types underbid their intrinsic valuation θ in the PPE if and only if $\lambda > 2$.

2.5 Revenue Comparison

The equilibrium bidding function of an English auction with loss-averse bidders strongly depends on the question how quickly new information is absorbed in the reference point.

If the reference point consists of lagged beliefs, and the lag is sufficiently high, new information during the auction process will have no impact on bidders reference point. If values are private, there is therefore no impact of information gathered during the auction process. Each bidder will form her optimal decision with respect to the initial belief, and thus faces the same objective function as in the Vickrey auction—the strategic equivalence between English and Vickrey auction remains.

If bidders, however, update their reference point dynamically with respect to new information, loss-averse bidders bid at most $(1+\eta)\theta$.

Figure 2.5 shows the equilibrium bidding function for the Vickrey auction, $b_{\text{Vickrey}}(\theta)$, and the PPE of the English auction with dynamic reference point updating, $b_{\text{English}}(\theta)$. The shaded area indicates the potential other symmetric equilibria in the English auction, which are bounded by the line $(1+\eta)\theta$.

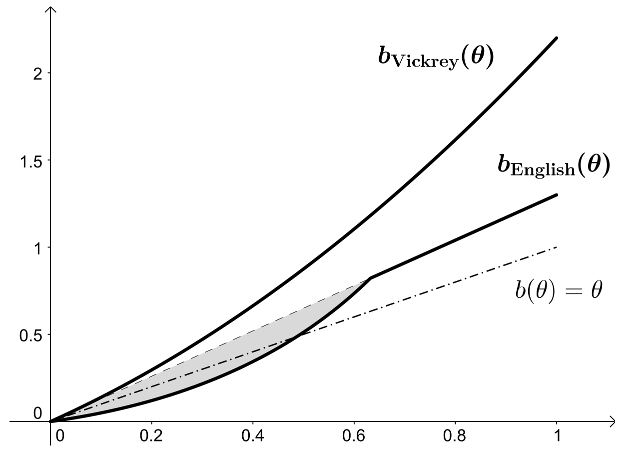

Figure 2.5: Equilibrium bidding function, $G(\theta) \sim U[0,1]$, $\eta = 0.3$, $\lambda = 4$

As we have seen in section 3, overbidding with respect to θ is moderate for low types and strong for high types in the Vickrey auction. We can see that $b_{\text{Vickrey}}(\theta)$ at the lowest type is tangent to $(1+\eta)\theta$—the upper bound of equilibria in the English auction. The intuition is that for low types the decision problem in both auction formats becomes increasingly similar: since bidders in the English auction only learn, whether there are opponents with lower valuation than their own, the information difference between the two auction formats at the time the bidder places her (maximal) bid is small for low types.

Since the bidding function in the Vickrey auction satisfies $b_{\text{Vickrey}}(\theta) > (1+\eta)\theta$ for all types $\theta > \theta^{\min}$, it is immediate that the Vickrey auction dominates the English auction with respect to revenue.

Theorem 2.3. *1. If bidders are loss averse and do not update their reference point during the auction process, the Vickrey auction and the English auction are strategically equivalent: for a given continuous belief on the maximal opponent bid, bidding b is a UPE in the Vickrey auction if and only if bidding up to b is a PE in the English auction.*

2. *If bidders are loss averse and update their reference point instantaneously during the auction process, equilibrium bids of the lowest type may coincide for both auction formats. For all other types, the Vickrey auction attains strictly higher revenue than the English auction.*

2.6 Extensions and Robustness

2.6.1 Loss Aversion in the Money Dimensions

We generalize the baseline model to the case where bidders are loss averse in both commodity dimensions—money and good.

The Vickrey Auction

The utility of a bidder of type θ who places a bid of b but has a reference point as if bidding b^* is given by

$$\begin{aligned}
u(b,\theta|b^*) :=& \sum_{k\in\{m,g\}} U(F^k(b,\theta,H)|F^k(b^*,\theta,H)) \\
=& \int_0^b \left(-s + \int_0^{b^*} \mu_m(t-s)dH(t) + \int_{b^*}^{\infty} \mu_m(-s)dH(t)\right) dH(s) \\
+& \int_b^{\infty} \left(\int_0^{b^*} \mu_m(t)dH(t) + \int_{b^*}^{\infty} \mu_m(0)dH(t)\right) dH(s) \\
+& \int_0^b \left(\theta + \int_0^{b^*} \mu_g(0)dH(t) + \int_{b^*}^{\infty} \mu_g(\theta)dH(t)\right) dH(s) \\
+& \int_b^{\infty} \left(\int_0^{b^*} \mu_g(-\theta)dH(t) + \int_{b^*}^{\infty} \mu_g(0)dH(t)\right) dH(s),
\end{aligned}$$

where H is again the distribution of the maximal opponent bid. The variable s corresponds to the realization of H, the variable t to the reference point. The first of the four summands corresponds to the utility in money if bidder i wins, the second if she loses. Similarly the third summand corresponds to utility in the good dimension if the auction is won, and summand four if the auction is lost.

In equilibrium the order statistic H is again endogenously determined by the opponents' equilibrium bids $b(\theta_{-i})$. Using the opponents' response functions, it is straightforward to calculate the symmetric equilibrium bidding function:

Theorem 2.4. *The unique symmetric increasing continuously differentiable UPE for n loss averse bidders in the Vickrey auction for commodities is given by*

$$
\begin{aligned}
b(\theta) =& \frac{1+\eta_g+\Lambda_g G^{n-1}(\theta)}{1+\lambda_m\eta_m}\theta \\
&+ \int_{\theta^{min}}^{\theta} \frac{\Lambda_m(1+\eta_g+\Lambda_g G^{n-1}(x))}{(1+\lambda_m\eta_m)^2} x \exp\left(\frac{\Lambda_m}{1+\lambda_m\eta_m}(G^{n-1}(\theta)-G^{n-1}(x))\right) dG(x).
\end{aligned}
$$

Note that

$$
b(\theta^{\min}) = \frac{1+\eta_g}{1+\lambda_m\eta_m}\theta^{\min},
$$

while for any $\theta > \theta^{\min}$

$$
b(\theta)) > \frac{1+\eta_g+\Lambda_g G^{n-1}(\theta)}{1+\lambda_m\eta_m}\theta > \frac{1+\eta_g}{1+\lambda_m\eta_m}\theta.
$$

In particular, for equally weighted loss aversion in both dimensions, low types underbid, while

$$
\begin{aligned}
b(\theta^{\max}) &> \frac{1+\eta+\Lambda G^{n-1}(\theta^{\max})}{1+\lambda\eta}\theta^{\max} \\
&= \frac{1+\eta+\Lambda}{1+\lambda\eta}\theta^{\max} \\
&= \theta^{\max}
\end{aligned}
$$

shows that high types overbid their intrinsic valuation. The intuition is that low types don't expect to win and try to avoid unexpected losses in the money dimension. In contrast, high types expect to win and try to avoid unexpected losses in the good dimension.

The English Auction

We avoid to fully classify the set of symmetric PE again, but rather straightforwardly prove that the revenue ranking between the two auction formats remains intact.[12] The following Lemma parallels Lemma 2.3.

Lemma 2.5. *Consider a loss-averse bidder of type θ in the continuous English auction with one opponent. Let the opponent's drop-out price be distributed according to distribution F with nonzero density f on some positive support $[a, b]$. Then,*

[12]The full derivation of the symmetric equilibrium bidding functions is available on request.

any time consistent bidding strategy $x \in (a, b)$ *satisfies*

$$x \leq \frac{1+\eta_g}{1+\lambda_m \eta_m}\theta.$$

Again, the bidders of high type ex ante like to commit excessive bids, but they know that the plan to bid above the threshold of $\frac{1+\eta_g}{1+\lambda_m\eta_m}\theta$ is time-inconsistent. Just one increment before they drop out, their belief to win and pay is virtually zero and—similarly to the lowest type in the Vickrey auction—they trade off the unexpected gain of the good against the unexpected loss in money, which may both occur with very small probability. If loss aversion is equally pronounced in both dimensions, then bidders underbid their intrinsic value θ, since losses weight stronger than gains.

Revenue Comparison

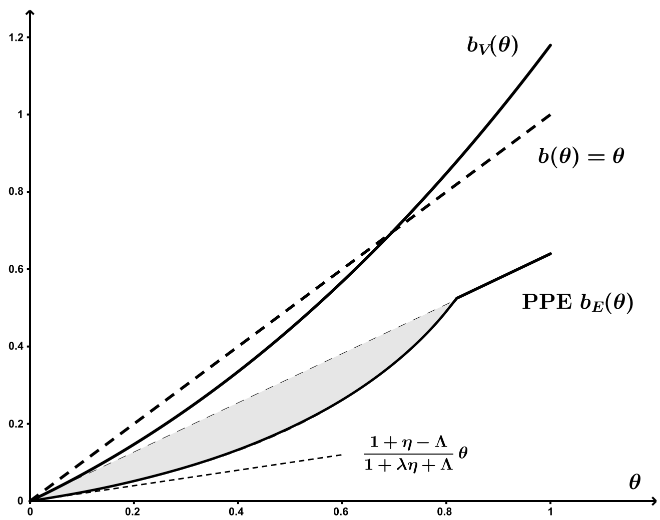

Figure 2.6: Equilibrium bidding functions, $G(\theta) \sim U[0,1]$, $\eta = 0.4$, $\lambda = 3$

Since in the Vickrey auction we have

$$b_{\text{Vickrey}}(\theta) \geq \frac{1+\eta_g}{1+\lambda_m \eta_m}\theta,$$

with equality only for $\theta^{\min}$, and in the English auction we have

$$b_{\text{English}}(\theta) \leq \frac{1+\eta_g}{1+\lambda_m \eta_m}\theta,$$

it is immediate that the Vickrey auction remains to dominate the English auction with respect to revenue.

Figure 2.6 shows the gray shaded area of potential equilibria in the English auctions, together with its PPE, and the equilibrium in the Vickrey auction. If loss aversion is equally pronounced in both dimensions, there is unambiguously underbidding in the English auction, while in the Vickrey auction low types underbid and high types overbid.

2.6.2 False Beliefs or Heterogeneous Preferences

So far we have assumed that all participating bidders are loss averse and hold rational beliefs over opponents' behavior. This is not a crucial assumption. Loss-averse bidders will bid higher in the Vickrey auction than in the English auction for *any* continuous belief with full support that they hold over opponents strategies.

Following the analysis of section 2.3, equation 2.1 in the proof of Theorem 2.1 states that for any such belief H the bidding function in the Vickrey auction is given by

$$b(\theta) = \big(1 + \eta(1 - H(b(\theta))) + \lambda\eta H(b(\theta))\big)\theta,$$

which shows that

$$b(\theta) > (1 + \eta)\theta$$

for all types, who win with positive probability. Contrary, in the English auction Lemma 2.3 shows that for any such belief

$$b(\theta) \leq (1 + \eta)\theta.$$

2.6.3 Reserve Price

A reserve price is a prominent tool in auctions to guarantee some minimum price. If agents are loss averse, a reserve price will also impact the bidding strategy of the bidders above the reserve price.

Consider a Vickrey auction with n loss-averse bidders. Since the implementation of a reserve price excludes low types from participation, an ex ante announcement of such would have a selection effect on bidders who participate. It would considerably change beliefs about the participating opponents' types. To abstract away from this effect, assume that a reserve price x is announced after the bidders committed to participate, and before bidders form their strategies. Bids below the reserve price remain feasible, but cannot win.

Proposition 2.4. *Let $b(\theta)$ be the equilibrium bidding function of n loss-averse bidders in the Vickrey auction without reserve price. If bidders are loss averse with respect to money, a public reserve price $\overline{x} > b(\theta^{min})$ increases the equilibrium bid of all bidders with $b(\theta) \geq \overline{x}$.*

Thus, if the object is sold, a reserve price increases revenues, even if it is not binding. To get the intuition for this result, note that the reserve price has no direct effect on the winning probability for bidders with $b(\theta) > \overline{x}$ in any symmetric increasing equilibrium. A reserve price has therefore no impact on loss aversion in the good dimension. However, the belief of paying less than $\overline{x}$ decreases. If bidders are loss averse with respect to money, high prices now induce less loss in the money dimension, with respect to expectations. This reduces expected gain-loss disutility from a high bid.

The same holds for similar reasons in the English auction with loss aversion in money, which we omit to prove here. In the English auction with loss aversion in the good dimension only, a reserve price has again no effect on equilibrium behavior.

Proposition 2.5. *Consider a continuous English auction with two bidders, who are loss averse in the good dimension. A reserve price of $\overline{x}$ has no effect on an equilibrium bidding function b for any type θ with $b(\theta) > \overline{x}$.*

2.6.4 Generalization to n bidders

In auctions where bidders face more than one opponent, the set of possible action plans becomes very large. Recall that an action plan prescribes a consistent action for any history and any future contingency at any time. While in the two bidder case the history is rather simple—either the opponent dropped out and the auction is over, or we are still in the auction process—with more bidders the individual decision at each time may in principle depend on the exact timing at which opponents dropped out in the past.

Since each decision must be sequentially optimal, given expectations about the future, one might hope to be able to restrict to Markov perfect equilibria, in the sense that at time t the individual type θ_i and the number of currently active bidders is a sufficient statistic for the optimal decision of bidder i. However, this is not the case. While the set of personal equilibria starting at time t can be determined without looking into the past, the specific equilibrium path will depend on the evolution of beliefs up to time t.

In order to deal with strategies contingent on histories, we define the following notation:

Definition 2.6. For any n-bidder auction, define for all $k \in \{0, ..., n-2\}$

$$H_k = \{(t_1, ..., t_k) | 0 \leq t_1 \leq, ..., \leq t_k\}$$

as the set of histories / future contingencies with k drop outs at the respective prices $t_1, ..., t_k$, with the convention $H_0 = \{\emptyset\}$.

With this notation, a complete action plan prescribes for each history and future contingency the price at which a bidder of type θ plans to drop out:

Definition 2.7. A pure strategy action plan prescribes a bidding strategy

$$b : \bigcup_{0 \leq k \leq n-2} H_k \times [\theta^{\min}, \theta^{\max}] \to \mathbb{R}_+,$$

with the restriction that if for any $(t_1, ..., t_k, \theta)$ we have

$$b(t_1, ..., t_k, \theta) > t_k,$$

The latter condition on the bidding function ensures that bidders cannot condition their drop out on events that happen after the drop out.

Again, we restrict attention to differentiable and increasing equilibrium bidding functions in the following sense:

Definition 2.8. A bidding strategy b in the English auction is differentiable and increasing if for all $(t_1, ..., t_k) \in \bigcup_{0 \leq k \leq n-2} H_k$ the function $b(t_1, ..., t_k, \theta)$ is differentiable and increasing in θ.

Example 2.1. Consider a continuous English auction with three loss-averse bidders. A complete strategy prescribes for every θ:

- A price $b(\theta)$ at which the bidder drops out if no opponent dropped out before
- For any opponent drop out at some price $t < b(\theta)$, a price $b(t, \theta)$ at which the bidder drops out in the subsequent two-bidder auction

The aim of the example is to illustrate, why the optimal strategy $b(t, \theta)$ for the two-bidder auction following the first drop out depends on t.Suppose that all three bidders bid according to the same symmetric equilibrium bidding strategy $(b(\theta), b(t, \theta))$. Let us focus on the decision problem of a bidder, whose valuation θ is sufficiently high, such that $b(\theta) = (1 + \eta)\theta$ were the only time-consistent strategy in the two-bidder English auction.

Suppose first that an opponent has a valuation of zero and drops out at $t = 0$. For the strategy $b(0, \theta)$ the bidder is now bound by the set of time-consistent strategies of the two-bidder auction, as outlined in Theorem 2.2. Since she has high beliefs to win, the only time-consistent strategy is $b(0,\theta) = (1+\eta)\theta$.

Next, we analyze optimal strategies $b(t,\theta)$ for t being smaller, but close to $b(\theta)$. Similar to the two-bidder auction, a bidder with a high winning probability would ex ante like to insure against belief fluctuations with an aggressive strategy. Any strategy for $b(t,\theta)$, however, must be time consistent in the sense that the bidder is willing to stick to it until t. Just before t the belief to win the auction has decreased considerably. The bidder trades off the expected gains from trade against the expected loss from news. The following Lemma states the expected loss at time t for the three bidder case.

Lemma 2.6. *Consider a continuous English auction with three loss-averse bidders. Assume all bidders follow a symmetric, differentiable, increasing bidding strategy* $(b(\theta), b(t,\theta))$*. Assume further that no bidder dropped out until* $t \in [b(\theta^{min}), b(\theta^{max})]$*. Let* $\theta(t)$ *be defined by* $b(\theta) = t$*. Then expected gain-loss utility at time* t *is given by*

$$L_t(\theta) = -\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(1-G(s))}{(1-G(\theta(t)))^2} \Bigg[\underbrace{\frac{G(\theta)-G(s)}{1-G(s)} - \left(\frac{G(\theta)-G(s)}{1-G(s)}\right)^2}_{A} \underbrace{- \ln\left(\frac{1-G(\theta)}{1-G(s)}\right) \frac{1-G(\theta)}{1-G(s)}}_{B} \Bigg] ds$$

The terms of $L_t(\theta)$ are easy to interpret. At time t the conditional marginal probability that the first drop out is of type s is given by $\frac{2g(s)(1-G(s))}{(1-G(\theta(t)))^2}$. In this case, the bidder would update the winning probability from $\left(\frac{G(\theta)-G(s)}{1-G(s)}\right)^2$ to $\frac{G(\theta)-G(s)}{1-G(s)}$ (term A). Further, term B shows the expected loss for the following 2-bidder auction, as calculated in Proposition 2.2.

Term A indicates an additional source of expected gain-loss disutility, compared to the two bidder auction: even if a bidder loses after all, beliefs to win don't necessarily gradually decline to zero, but might temporarily increase due to one opponent dropping out. This effect leads to more belief fluctuations and worsens bidder's trade-off between expected news disutility and expected gains from trade. As a result, it is no longer time consistent to bid up to $b(t,\theta) = (1-\eta)\theta$ for all t.

Corollary 2.4. *In any symmetric, increasing, differentiable equilibrium* $(b(\theta), b(t,\theta))$ *of the English auction with three loss-averse bidders, expected news disutility for any* $\theta \in (\theta^{min}, \theta^{max})$ *satisfies*

$$\lim_{t \to b(\theta)} \frac{L_t(\theta)}{\left(\frac{G(\theta)-G(\theta(t))}{1-G(\theta(t))}\right)^2} = -2\Lambda\theta.$$

If $b(t,\theta)$ *is continuous in* t*, then—by time-consistency—*

$$\lim_{t \to b(\theta)} b(t,\theta) \leq (1+\eta-\Lambda)\theta.$$

Since we have argued above that $b(0, \theta) = (1+\eta)\theta$, the corollary illustrates that bidding behavior $b(t, \theta)$ in general depends on opponents' drop-out history t.

Even if the sales price depends on all type realizations, it is immediate that for n bidders the revenue ranking between the two auction format remains: since bidders generically don't share the same valuation, in any symmetric continuous increasing equilibrium they will drop out of the auction consecutively, in order of their types. Eventually, with probability one, the two bidders with the highest valuation will end up in the two-bidder subgame. Here they are bound to the constraints on time-consistent behavior, as analyzed in section 2.4.2. In particular by Lemma 2.3, any time-consistent strategy for the two-bidder auction satisfies $b(\theta) \leq (1 + \eta)\theta$.

To summarize:

Corollary 2.5. *In a symmetric increasing equilibrium of the continuous English auction with n loss-averse bidders, the revenue may depend on all type realizations. For any opponent drop-out history, every bidder's maximal bid is bounded by $b(\theta) \leq (1 + \eta)\theta$. Thus, with n loss-averse bidders, the English auction remains to yield lower revenues than the Vickrey auction.*

Even if the auction outcome for many bidders is similar to the one for two bidders, it is worth noting that individual bidders obtain less utility, compared to two-bidder auctions with the same sales price. To see this, consider—hypothetically—that bidders could choose not to observe individual drop outs, but rather learn in each period, whether *any* opponent is still in the game. The auction would then subjectively resemble an English auction with two bidders, where the opponent's type is drawn from the first order-statistic over all opponents. The key difference is that information is fluctuating much less. As already mentioned earlier and stated in generality in Proposition 1 of Kőszegi and Rabin (2009), the collapse of multiple signals into one will always weakly decrease gain-loss disutility.

2.7 Conclusion

I studied the effects of expectation-based preferences in dynamic environments, comparing the dynamic English auction to the static Vickrey auction. If the reference point is static and doesn't respond to information, there is no strategic difference between the English auction and the Vickrey auction. If bidders update their reference point instantaneously with respect to new information, however, dynamic information in the English influences bidders endogenous preferences, and thus their bidding strategies. The classical strategic equivalence between the the two auction formats breaks down and the English auction attains strictly lower revenue than the Vickrey auction.

This difference highlights the importance of understanding the evolution of the reference point in dynamic environments. In particular, research about the speed of reference point adaptation with respect to new information is still in its infancy and deserves further study.

The non-equivalence of the two auction formats stands in sharp contrast to the revenue equivalence principles by Vickrey (1961) and Myerson (1981). Indeed, the powerful approach of mechanism design and the revelation principle relies on the assumption that agents' valuations are exogenously given and do not depend on the choice of mechanism. This assumption is violated if bidders have endogenous preferences that depend on expectations induced by the mechanism itself. In particular, if agents update their reference point with respect to new information in a multi-stage mechanism, such a mechanism cannot be replaced by a simple direct mechanism without changing agents' incentives. The failure of the revelation principle naturally leads to the question of optimal mechanism design in dynamic environments with expectation-based loss-averse agents. The study of optimal expectation management in these environments leaves an interesting field for future research.

2.8 Appendix to Chapter 2

Proof of Theorem 2.3.4. Suppose that all opponents bid according to some increasing, continuously differentiable bidding function $b(\theta)$. Since $G(\theta)$ is a distribution with strictly positive, continuous density g, it follows that the distribution of the maximal opponent bid, $H(x) = G^{n-1}(b^{-1}(x))$, is a differentiable distribution with positive, continuous density $h(x)$ on $[b(\theta^{\min}), b(\theta^{\max})]$ as well.

The bidding function $b(\theta)$ constitutes a UPE if and only if the utility function $u(x, \theta|b(\theta))$ attains its maximum at $x = b(\theta)$ for all θ. Differentiation with respect to x yields

$$\frac{\partial u(x, \theta|b(\theta))}{\partial x} = (\theta - x)h(x) + h(x)(1 - H(b(\theta))\mu(\theta) - h(x)H(b(\theta))\mu(-\theta).$$

By dividing by $h(x)$ and evaluating at $x = b(\theta)$ we obtion the first-order condition

$$0 = (\theta - b(\theta)) + (1 - H(b(\theta)))\eta\theta + H(b(\theta))\lambda\eta\theta.$$

Rearranging yields

$$b(\theta) = \big(1 + \eta(1 - H(b(\theta))) + \lambda\eta H(b(\theta))\big)\theta. \tag{2.1}$$

Using that $H(b(\theta)) = G^{n-1}(\theta)$ we obtain

$$b(\theta) = \big(1 + \eta(1 - G^{n-1}(\theta)) + \lambda\eta G^{n-1}(\theta)\big)\theta$$

as the unique equilibrium candidate. For sufficiency note first that

$$h(b(\theta)) = \frac{(G^{n-1})'(\theta)}{b'(\theta)} = \frac{(n-1)G^{n-2}(\theta)g(\theta)}{(1 + \eta(1 - G^{n-1}(\theta)) + \lambda\eta G^{n-1}(\theta)) + \Lambda(n-1)G^{n-2}(\theta)g(\theta)\theta}$$

is differentiable since $g(\theta)$ is differentiable. Now it is immediate that

$$\begin{aligned}\frac{\partial^2 u_i(x,\theta|b(\theta))}{(\partial x)^2}\Big|_{x=b(\theta)} &= -h(b(\theta)) + h'(b(\theta))\underbrace{(\theta - b(\theta)) + (1 - H(b(\theta))\mu(\theta) - H(b(\theta))\mu(-\theta)]}_{=0} \\ &< 0.\end{aligned}$$

□

Proof of Lemma 2.1. By the theorem of the integral over inverse functions, we have

$$\int_a^b F_i(x)dx = bF_i(b) - aF_i(a) - \int_0^1 c_{F_i}(p)dp = b - \int_0^1 c_{F_i}(p)dp.$$

Now, it is immediate that

$$\int_a^b (F_1(x) - F_2(x))dx = (b-b) - \int_0^1 c_{F_1}(p)dp + \int_0^1 c_{F_2}(p)dp = \int_0^1 (c_{F_2}(p) - c_{F_1}(p))dp.$$

□

Proof of Proposition 2.1. By applying Lemma 2.1, and using the fact that μ is piecewise linear, we can write

$$\begin{aligned}
\mathbb{E}(N(F_i|F) &= \Delta N(F_1|F) + (1-\Delta)N(F_2|F)\\
&= \Delta \int_0^1 \mu(c_{F_1}(p) - c_F(p))dp + (1-\Delta)\int_0^1 \mu(c_{F_2}(p) - c_F(p))dp\\
&= \Delta \int_a^b \mu(F(x) - F_1(x))dx + (1-\Delta)\int_a^b \mu(F(x) - F_2(x))dx\\
&= \Delta \int_a^b \mu(F(x) - F_1(x))dx + \int_a^b \mu((1-\Delta)F(x) - (1-\Delta)F_2(x))dx\\
&= \Delta \int_a^b \mu(F(x) - F_1(x))dx + \int_a^b \mu((1-\Delta)F(x) - (F(x) - \Delta F_1(x))dx\\
&= \Delta \int_a^b \mu(F(x) - F_1(x))dx + \int_a^b \mu(-\Delta F(x) + \Delta F_1(x))dx\\
&= \Delta \int_a^b \mu(F(x) - F_1(x))dx + \Delta \int_a^b \mu(-F(x) + F_1(x))dx\\
&= \Delta(-\lambda\eta + \eta)\int_a^b |F(x) - F_1(x)|dx\\
&= -\Delta\Lambda \int_a^b |F(x) - F_1(x)|dx\\
&= -\Delta\Lambda \int_0^1 |c_{F_1}(p) - c_F(p)|dp.
\end{aligned}$$

□

Proof of Proposition 2.2. Suppose the current clock price is t and the opponent hasn't dropped out yet. If the clock increases in increments of ε, then the conditional probability that the opponent drops out at the next increment is given by

$$\Delta_t := \frac{F(t+\varepsilon) - F(t)}{1 - F(t)}.$$

Given her strategy x and that the opponent hasn't dropped out at t, the bidder faces the conditional probability of $\frac{1-F(x)}{1-F(t)}$ to lose the auction. Thus, if F_t^x denotes the belief about payoffs in the good dimension at time t given strategy x, we have

$$F_t^x(z) = \begin{cases} \frac{1-F(x)}{1-F(t)}, & z < \theta,\\ 1, & z \geq \theta. \end{cases}$$

If the bidder wins in the next increment, the belief will update to

$$F_{t+\varepsilon}^x(z) = \begin{cases} 0, & z < \theta,\\ 1, & z \geq \theta. \end{cases}$$

According to Proposition 2.1, expected gain-loss utility of the increment from t to $t+\varepsilon$ is then given by

$$\mathbb{E}(N(F^x_{t+\varepsilon}|F^x_t) = -\Delta_t\Lambda \int |F^x_t(z) - F^x_{t+\varepsilon}(z)|dz = -\Delta_t\Lambda\frac{1-F(x)}{1-F(t)}\theta.$$

Now, the marginal loss at time t if ε goes to zero reads

$$\ell_t(x,\theta,F) = \lim_{\varepsilon\to 0}\frac{-\Delta_t\Lambda\frac{1-F(x)}{1-F(t)}\theta}{\varepsilon} = \frac{-f(t)}{(1-F(t))^2}(1-F(x))\Lambda\theta.$$

To calculate total expected gain-loss utility starting at time t, note that any information update at time $s>t$ is only informative and carries gain-loss utility if the opponent hasn't already dropped out between t and s, which holds true with the conditional probability $\frac{1-F(s)}{1-F(t)}$. Thus

$$\begin{aligned}
L_t(x,\theta,F) &= \lim_{\varepsilon\to 0}\sum_{i=0}^{\lfloor\frac{x-t}{\varepsilon}\rfloor-1} N(F^x_{t+(i+1)\varepsilon}|F^x_{t+i\varepsilon}) \\
&= \lim_{\varepsilon\to 0}\sum_{i=0}^{\lfloor\frac{x-t}{\varepsilon}\rfloor-1} -\frac{1-F(t+i\varepsilon)}{1-F(t)}\Delta_{t+i\varepsilon}\Lambda\frac{1-F(x)}{1-F(t+i\varepsilon)}\theta \\
&= \int_t^x \frac{-f(s)}{1-F(s)}\frac{1-F(x)}{1-F(t)}\Lambda\theta ds \\
&= (\ln(1-F(x)) - \ln(1-F(t)))\frac{1-F(x)}{1-F(t)}\Lambda\theta \\
&= \ln\left(\frac{1-F(x)}{1-F(t)}\right)\frac{1-F(x)}{1-F(t)}\Lambda\theta.
\end{aligned}$$

□

Proof of Lemma 2.2. At time t the winning probability is given by the probability that the opponent drops out between t and x, given he didn't drop out before t, thus $\frac{F(x)-F(t)}{1-F(t)}$. Thus, the update changes the probability of getting θ by

$$\frac{F(y)-F(t)}{1-F(t)} - \frac{F(x)-F(t)}{1-F(t)} = \frac{F(y)-F(x)}{1-F(t)}.$$

Hence,

$$N(F^y_t|F^x_t) = \mu\left(\frac{F(y)-F(x)}{1-F(t)}\theta\right) = \frac{\mu(F(y)-F(x))}{1-F(t)}\theta.$$

□

Proof of Lemma 2.3. The bidder does not want do deviate to a lower strategy at any time t, given plan x if and only if

$$u_t(y, \theta, F|x) \leq u_t(x, \theta, F|x)$$

for all $t \leq y \leq x$. In particular it is necessary that for all $t < x$ the derivative from the left satisfies

$$\begin{aligned} 0 &\leq \lim_{y \nearrow x} \frac{\partial u_t(y, \theta, F|x)}{\partial y} \\ &= \frac{f(x)}{1 - F(t)} \left(\theta - x + \lambda\eta\theta - \Lambda \left(1 + \ln \left(\frac{1 - F(x)}{1 - F(t)} \right) \right) \theta \right). \end{aligned}$$

This expression is well defined, since $F(t) < F(x) < 1$. Now, as t approaches x we get

$$\begin{aligned} 0 &\leq \lim_{t \to x} \frac{f(x)}{1 - F(t)} \left(\theta - x + \lambda\eta\theta - \Lambda \left(1 + \ln \left(\frac{1 - F(x)}{1 - F(t)} \right) \right) \theta \right) \\ &= \frac{f(x)}{1 - F(x)} \left(\theta - x + \lambda\eta\theta - \Lambda\theta \right). \end{aligned}$$

Since, by assumption, $f(x) > 0$, this means that necessarily

$$x \leq (1 + \lambda\eta - \Lambda)\theta = (1 + \eta)\theta.$$

□

Proof of Proposition 2.3. Given opponent's strategy F and bidder's type θ, a bid $b(\theta)$ is a personal equilibrium in the auction with commitment if and only if

$$u_0(y, \theta, F|b(\theta)) \leq u_0(b(\theta), \theta, F|b(\theta))$$

for all y. In particular, it is necessary that

$$\lim_{y \searrow b(\theta)} \frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y} \leq 0.$$

Since for $y > b(\theta)$ the utility at time zero reads

$$u_0(y, \theta, F|b(\theta)) = \int_0^y (\theta - s) dF(s) + \eta(F(y) - F(b(\theta)))\theta + \ln(1 - F(y))(1 - F(y))\Lambda\theta,$$

this necessary condition is equivalent to

$$f(b(\theta))\big(\theta - b(\theta) + \eta\theta - \Lambda(1 + \ln(1 - F(b(\theta))))\theta\big) \leq 0.$$

In any symmetric equilibrium, the opponent bids according to $b(\theta)$ as well, and therefore we have $F(b(\theta)) = G(\theta)$. From $g(\theta) = f(b(\theta))b'(\theta)$ and the restriction that b is increasing it follows that $f(b(\theta)) > 0$. Hence we have

$$b(\theta) \geq \big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta)))\big)\theta$$

for any equilibrium candidate. It remains to verify that

$$b(\theta) = \big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta)))\big)\theta$$

is a personal equilibrium, given opponent's response $b(\theta)$. For this it is sufficient to show that

$$\frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y} \leq 0$$

for all $y > b(\theta)$, and

$$\frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y} \geq 0$$

for all $y < b(\theta)$. Note that we can without loss of generality restrict to $y \in [b(\theta^{\min}), b(\theta^{\max})]$.

For any such y there exists some $\tilde{\theta}$ with $y = b(\tilde{\theta})$, since the bidding function is continuous.

Consider first $y > b(\theta)$, thus $\tilde{\theta} > \theta$. Then

$$\begin{aligned}
\frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y}\Big|_{y=b(\tilde{\theta})} &= f(b(\tilde{\theta}))\big(\theta - b(\tilde{\theta}) + \eta\theta - \Lambda(1 + \ln(1 - F(b(\tilde{\theta}))))\theta\big) \\
&< f(b(\tilde{\theta}))\big(\tilde{\theta} - b(\tilde{\theta}) + \eta\tilde{\theta} - \Lambda(1 + \ln(1 - F(b(\tilde{\theta}))))\tilde{\theta}\big) \\
&= \lim_{y \searrow b(\tilde{\theta})} \frac{\partial u_0(y, \tilde{\theta}, F|b(\theta))}{\partial y} \\
&= 0.
\end{aligned}$$

Similarly, for $y < b(\theta)$, thus $\tilde{\theta} < \theta$ we have

$$\begin{aligned}
\frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y}\Big|_{y=b(\tilde{\theta})} &= f(b(\tilde{\theta}))\big(\theta - b(\tilde{\theta}) + \lambda\eta\theta - \Lambda(1 + \ln(1 - F(b(\tilde{\theta}))))\theta\big) \\
&> f(b(\tilde{\theta}))\big(\tilde{\theta} - b(\tilde{\theta}) + \eta\tilde{\theta} - \Lambda(1 + \ln(1 - F(b(\tilde{\theta}))))\tilde{\theta}\big) \\
&= \lim_{y \searrow b(\tilde{\theta})} \frac{\partial u_0(y, \tilde{\theta}, F|b(\theta))}{\partial y} \\
&= 0.
\end{aligned}$$

□

Proof of Lemma 2.4. Consider a bidding strategy x.

Claim 1: If and only if $x \leq (1+\eta)\theta$, it is at no time $t < x$ profitable to deviate to a lower strategy $y \in [t, x)$.
Proof: the "only if" has been proved in Lemma 2.3. For the "if", assume that $x \leq (1+\eta)\theta$. Consider a deviation at some time $t < x$ from x to $y \in [t, x)$. We first look at the change in expected gain-loss disutility: term A can be interpreted as the change due to different expectations at each time between t and y, while term B is forgone gain-loss disutility, since the auction necessarily ends at y:

$$
\begin{aligned}
&L_t(y,\theta,F) - L_t(x,\theta,F)\\
=&\Lambda\theta\left(\ln\left(\frac{1-F(y)}{1-F(t)}\right)\frac{1-F(y)}{1-F(t)} - \ln\left(\frac{1-F(x)}{1-F(t)}\right)\frac{1-F(x)}{1-F(t)}\right)\\
=&\Lambda\theta\left(\int_t^y \frac{-f(s)}{1-F(s)}ds\frac{1-F(y)}{1-F(t)} - \int_t^x \frac{-f(s)}{1-F(s)}ds\frac{1-F(x)}{1-F(t)}\right)\\
=&\Lambda\theta\left(\int_t^y \frac{-f(s)}{1-F(s)}ds\frac{1-F(y)}{1-F(t)} - \int_t^y \frac{-f(s)}{1-F(s)}ds\frac{1-F(x)}{1-F(t)} - \int_y^x \frac{-f(s)}{1-F(s)}ds\frac{1-F(x)}{1-F(t)}\right)\\
=&\Lambda\theta\left(\underbrace{\int_t^y \frac{-f(s)}{1-F(s)}ds\frac{F(x)-F(y)}{1-F(t)}}_{A} - \underbrace{\int_y^x \frac{-f(s)}{1-F(s)}ds\frac{1-F(x)}{1-F(t)}}_{B}\right)\\
\leq&\Lambda\theta\int_y^x \frac{f(s)}{1-F(s)}ds\frac{1-F(x)}{1-F(t)}\\
<&\Lambda\theta\int_y^x f(s)ds\frac{1-F(x)}{(1-F(x))(1-F(t))}\\
=&\Lambda\theta\frac{F(x)-F(y)}{1-F(t)}
\end{aligned}
$$

Now we have

$$
\begin{aligned}
&u_t(y,\theta,F|x) - u_t(x,\theta,F|x)\\
=&\frac{1}{1-F(t)}\left(-\int_y^x (\theta - s)dF(s) + \mu(F(y)-F(x))\theta + \Lambda\theta(F(x)-F(y))\right)\\
<&\frac{F(x)-F(y)}{1-F(t)}\left(-\theta + x - \lambda\eta\theta + \Lambda\theta\right)\\
=&\frac{F(x)-F(y)}{1-F(t)}\left(-(1+\eta)\theta + x\right)\\
\leq&0.
\end{aligned}
$$

Thus, there is no profitable deviation to $y < x$ at any time, which concludes the proof of claim 1.

Claim 1 directly shows the necessity of 1. for any PE. Certainly, 2. is necessary as well.

Claim 2: If it is not profitable to deviate to a strategy $y > x$ at time $t = 0$, then it is not profitable at any time $t \leq x$.
Proof: It is not profitable to deviate to a strategy $y > x$ at time t if and only if

$$0 \geq u_t(y,\theta,F|x) - u_t(x,\theta,F|x)$$

Now,

$$
\begin{aligned}
&u_t(y,\theta,F|x) - u_t(x,\theta,F|x)\\
=&\frac{1}{1-F(t)}\left(\int_x^y (\theta - s)dF(s) + \mu(F(y)-F(x))\theta\right)\\
&+\Lambda\theta\left(\frac{1-F(y)}{1-F(t)}\ln\left(\frac{1-F(y)}{1-F(t)}\right) - \frac{1-F(x)}{1-F(t)}\ln\left(\frac{1-F(x)}{1-F(t)}\right)\right)\\
=&\frac{1}{1-F(t)}\left(\int_x^y (\theta - s)dF(s) + \mu(F(y)-F(x))\theta \; ...\right.\\
&\left. ... + \Lambda\theta((1-F(y))\ln(1-F(y)) - (1-F(x))\ln(1-F(x)) + (F(y)-F(x))\ln(1-F(t)))\right).
\end{aligned}
$$

Note that the expression in the big brackets is decreasing in t. Thus, if it is negative for $t = 0$, then it is as well negative for all $t > 0$. Hence, if

$$0 \geq u_0(y,\theta,F|x) - u_0(x,\theta,F|x)$$

then

$$0 \geq u_t(y, \theta, F|x) - u_t(x, \theta, F|x)$$

for all $t > 0$, which concludes the proof of claim 2.

Now we are ready to show sufficiency: assume 1. and 2. hold. Then by claim 1 it can't be profitable to deviate to a lower strategy at any time. To show that there is no profitable deviation to a higher strategy, take any time-consistent strategy $y \geq x$. By claim 1 this necessarily means $y \in [x, (1+\eta)\theta]$. From 2. it follows that $u_0(x, \theta, F|x) \geq u_0(y, \theta, F|x)$. Then, by claim 2, the agent does not want to deviate to a higher strategy at any time, and x is indeed a PE. □

Proof of Theorem 2.2. Take some increasing equilibrium function. By Lemma 2.4, it satisfies $b(\theta) \leq (1+\eta)\theta$ for all $\theta \in (\theta^{\min}, \theta^{\max})$. If $b(\theta) < (1+\eta)\theta$ for some θ, then—again by Lemma 2.4—any $y \in [x, (1+\eta)\theta]$ satisfies $u_0(x, \theta, F|x) \geq u_0(y, \theta, F|x)$. This means that

$$\lim_{y \searrow x} \frac{\partial u_0(y, \theta, F|x)}{\partial y} \leq 0,$$

which—as we have seen in the proof of Proposition 2.3—straightforwardly solves to

$$b(\theta) \geq \big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta)))\big)\theta$$

in equilibrium. This shows that any increasing equilibrium satisfies 1. and 2. for all $\theta \in (\theta^{\min}, \theta^{\max})$. By continuity it also holds for all $\theta \in [\theta^{\min}, \theta^{\max}]$. Conversely, assume that $b(\theta)$ satisfies 1. and 2. By Lemma 2.4 it only remains to show that for any

$$y \in [b(\theta), (1+\eta)\theta]$$

we have

$$u_0(b(\theta), \theta, F|b(\theta)) \geq u_0(y, \theta, F|b(\theta)).$$

This condition is trivially satisfied for any θ with $b(\theta) = (1+\eta)\theta$. Consider therefore θ with $b(\theta) < (1+\eta)\theta$. It suffices to show that

$$\frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y} \leq 0$$

for all $y \in [b(\theta), (1+\eta)\theta]$. Let $\tilde{y}$ be any of such y. Since

$$b(\theta^{\max}) = (1+\eta)\theta^{\max} > (1+\eta)\theta \geq \tilde{y} \geq b(\theta),$$

and b is continuous, there exists some $\tilde{\theta} \geq \theta$ with $b(\tilde{\theta}) = \tilde{y}$. Now,

$$\begin{aligned}
\frac{\partial u_0(y, \theta, F|b(\theta))}{\partial y}\Big|_{y=\tilde{y}} &= [(1+\eta)\theta - \tilde{y} - \Lambda\theta(1 + \ln(1 - F(\tilde{y})))]f(\tilde{y}) \\
&= [\underbrace{(1 + \eta - \Lambda(1 + \ln(1 - F(b(\tilde{\theta})))))}_{>0}\theta - b(\tilde{\theta})]f(b(\tilde{\theta})) \\
&\leq [\underbrace{(1 + \eta - \Lambda(1 + \ln(1 - F(b(\tilde{\theta})))))\tilde{\theta}}_{\leq b(\tilde{\theta})} - b(\tilde{\theta})]f(b(\tilde{\theta})) \\
&\leq 0.
\end{aligned}$$

□

Proof of Corollary 2.3. We have

$$\big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta)))\big)\theta \leq (1+\eta)\theta$$

if and only if $-(1 + \ln(1 - G(\theta))) \leq 0$, which is equivalent to $G(\theta) \leq 1 - 1/e$. Therefore, by Theorem 2.2, a fuction $b(\theta)$ is a symmetric equilibrium if and only if

- $b(\theta) \in \big[(1 + \eta - \Lambda(1 + \ln(1 - G(\theta))))\theta, (1+\eta)\theta\big]$ for $G(\theta) \leq 1 - 1/e$, and
- $b(\theta) = (1+\eta)\theta$ for $G(\theta) > 1 - 1/e$.

We determine the utility maximizing equilibrium on the interval where $G(\theta) \leq 1 - 1/e$. Bidder's expected utility of a bid x is

$$\begin{aligned}
u_0(x, \theta, F|x) &= \int_0^x (\theta - s)dF(s) + L_t(x, \theta, F) \\
&= \int_0^x (\theta - s)dF(s) + \Lambda\theta \ln(1 - F(x))(1 - F(x)).
\end{aligned}$$

Thus, for any $x \geq \big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta)))\big)\theta$

$$\begin{aligned}
\frac{\partial u_0(x, \theta, F|x)}{\partial x} &= (\theta - x)f(x) - \Lambda\theta(1 + \ln(1 - F(x)))f(x) \\
&\leq (\theta - (1 + \eta - \Lambda(1 + \ln(1 - G(\theta))))\theta)f(x) - \Lambda\theta f(x) \\
&\leq (\theta - (1 + \eta - \Lambda)\theta)f(x) - \Lambda\theta f(x) \\
&= -f(x)\eta\theta \\
&< 0.
\end{aligned}$$

This shows that the lowest x among all equilibrium strategies yields the highest utility.

Finally, since for the PPE

$$b(\theta^{\min}) = \big(1 + \eta - \Lambda(1 + \ln(1 - G(\theta^{\min})))\big)\theta^{\min} = (1 + \eta - \Lambda)\theta^{\min},$$

there is underbidding for low types in the PPE if and only if

$$0 > \eta - \Lambda = 2\eta - \lambda\eta,$$

hence if and only if $\lambda > 2$. □

Proof of Theorem 2.3. For (1) we show that without interim update the equilibrium concepts of the static UPE and the dynamic PE coincide. Given type θ and a continuous belief H on the maximal opponent bid, bidding (up to) b induces the same payoff belief (and therefore reference point) $F^k(b, \theta, H)$ for $k \in \{money, good\} = \{m, g\}$ in the Vickrey and the English auction. Consider a bidder in the English auction who plans to bid up to b but deviates during the auction process, such that the final payoff in dimension $k \in \{m, g\}$ is distributed according to F. If there is no interim updating during the auction, the bidder updates her reference point only once when the auction is terminated. Integrating the utility in dimension k for each possible auction outcome yields expected utility of

$$\begin{aligned}
U_{\text{English}}(F|F^k(b, \theta, H)) &= \int (x + N(\mathbb{1}_{[x,\infty)}|F^k(b, \theta, H)))dF(x) \\
&= \int \left(x + \int_0^1 \mu(x - c_F(p))dp \right) dF(x) \\
&= \int \left(x + \int_{-\infty}^{\infty} \mu(x - c_F(F(s)))dF(s) \right) dF(x) \\
&= \int \left(x + \int_{-\infty}^{\infty} \mu(x - s)dF(s) \right) dF(x) \\
&= U_{\text{Vickrey}}(F|F^k(b, \theta, H)).
\end{aligned}$$

Thus, equally for the UPE concept in the Vickrey auction and the PE concept in the English auction, an action b is an equilibrium if and only if for all distributions (F^m, F^g) that are induced by a deviation strategy we have

$$\sum_{k \in \{m,g\}} U(F^k(b, \theta, H)|F^k(b, \theta, H)) \geq \sum_{k \in \{m,g\}} U(F^k|F^k(b, \theta, H)).$$

The subtle difference lies in the fact that a bidder in the Vickrey auction is constrained to deviations $\hat{b} \in \mathbb{R}^+$, while a bidder in the English auction with multiple opponents can use complex history dependent deviation strategies, leading to are larger set of potential price distributions than in the Vickrey auction. Clearly, if action b is optimal with respect to all possible deviations in the English auction, it is in particular optimal with respect to deviations to all history-independent strategy $\hat{b} \in \mathbb{R}^+$. Thus, if bidding up to b is a PE in the English auction, then bidding b is a UPE in the Vickrey auction. For the converse, assume that b is a UPE in the Vickrey auction and let (F^m, F^g) be the payoff distribution of some deviation strategy in the English auction. Since H is continuous, there is some $\hat{b}$ such that $F^g(\hat{b}, \theta, H) = F^g$. Further, since strategy $\hat{b}$ wins the auction if and only if the maximal opponent strategy is below $\hat{b}$, it is the most cost effective strategy that wins with probability $H(\hat{b})$. Thus the distribution F^m induces weakly higher costs than $F^m(\hat{b}, \theta, H)$ in the sense of first-order stochastic dominance. It follows that

$$U(F^m | F^k(b, \theta, H)) \leq U(F^m(\hat{b}, \theta, H) | F^k(b, \theta, H)) \leq U(F^m(b, \theta, H) | F^k(b, \theta, H)),$$

and since consequently

$$\sum_{k \in \{m,g\}} U(F^k(b, \theta, H) | F^k(b, \theta, H)) \geq \sum_{k \in \{m,g\}} U(F^k | F^k(b, \theta, H)),$$

the strategy b is a PE in the English auction.

For (2) note that by Theorem 2.1 the equilibrium bidding function for the Vickrey auction is given by

$$b_{\text{Vickrey}}(\theta) = (1 + \eta + \Lambda G^{n-1}(\theta))\theta,$$

whereas any equilibrium bidding function in the English auction with instantaneous reference point updating by Lemma 2.4 satisfies

$$b_{\text{English}}(\theta) \leq (1 + \eta)\theta.$$

Since, by assumption, $G^{n-1}(\theta)$ is strictly increasing, we have $G^{n-1}(\theta) > 0$ for all $\theta > \theta^{\min}$, and the claim follows. □

Proof of Theorem 2.1. The structure of the proof is similar to Lange and Ratan

(2010). The utility function can be simplified to

$$\begin{aligned}u(b,\theta|b^*) = &\int_0^b (\theta - s)dH(s)\\ &+ \int_0^b \int_0^{b^*} \mu_m(t-s)dH(t)dH(s) + (1-H(b^*))\int_0^b \mu_m(-s)dH(s)\\ &+ (1-H(b))\int_0^{b^*} \mu_m(t)dH(t)\\ &+ H(b)(1-H(b^*))\mu_g(\theta) + H(b^*)(1-H(b))\mu_g(-\theta).\end{aligned}$$

Suppose that all opponents bid according to some increasing, continuously differentiable bidding function $b(\theta)$. Since $G(\theta)$ is a distribution with strictly positive, continuous density g, distribution of the maximal opponent bid $H(x) = G^{n-1}(b^{-1}(x))$ is a differentiable distribution with positive, continuous density $h(x)$ on $[b(\theta^{\min}), b(\theta^{\max})]$ as well. The bidding function $b(\theta)$ constitutes a UPE if and only if the utility function $u(x,\theta|b(\theta))$ attains a maximum at $x = b(\theta)$ for all θ. Differentiation of the utility function with respect to x yields

$$\begin{aligned}\frac{\partial u(x,\theta|b(\theta))}{\partial x} =&(\theta - x)h(x) + \int_0^{b(\theta)} \mu_m(t-x)h(x)dH(t) + (1-H(b(\theta)))\mu_m(-x)h(x)\\ &- h(x)\int_0^{b(\theta)} \mu_m(t)dH(t) + h(x)(1-H(b(\theta)))\eta_g\theta + h(x)H(b(\theta))\lambda_g\eta_g\theta.\end{aligned}$$

By dividing by $h(x)$ and evaluating at $x = b(\theta)$, we obtain the first-order condition

$$\begin{aligned}0 \overset{!}{=}&(\theta - b(\theta)) + \int_0^{b(\theta)} \mu_m(t-b(\theta))dH(t) + (1-H(b(\theta))\mu_m(-b(\theta))\\ &- \int_0^{b(\theta)} \mu_m(t)dH(t) + (1-H(b(\theta)))\eta_g\theta + H(b(\theta))\lambda_g\eta_g\theta\\ =&(\theta - b(\theta)) - \lambda_m\eta_m\int_0^{b(\theta)} (b(\theta)-t)dH(t) + (1-H(b(\theta)))(-\lambda_g\eta_g b(\theta))\\ &- \eta_m\int_0^{b(\theta)} tdH(t) + (1-H(b(\theta)))\eta_g\theta + H(b(\theta))\lambda_g\eta_g\theta,\end{aligned}$$

which simplifies to

$$0 = (1+\eta_g)\theta - (1+\lambda_m\eta_m)b(\theta) + \Lambda_m\int_0^{b(\theta)} tdH(t) + \Lambda_g H(b(\theta))\theta. \tag{2.2}$$

Using that $H(b(\theta)) = G^{n-1}(\theta)$ we can rewrite this equation to

$$0 = (1+\eta_g)\theta - (1+\lambda_m\eta_m)b(\theta) + \Lambda_m\int_0^{\theta} b(s)dG^{n-1}(s) + \Lambda_g G^{n-1}(\theta)\theta.$$

Differentiation with respect to θ yields

$$0 = (1+\eta_g) - (1+\lambda_m\eta_m)b'(\theta) + \Lambda_m b(\theta)(G^{n-1})'(\theta) + \Lambda_g(G^{n-1}(\theta)\theta)'.$$

The rearranged equation

$$b'(\theta) = \frac{\Lambda_m (G^{n-1})'(\theta)}{1+\lambda_m\eta_m} b(\theta) + \frac{1+\eta_g+\Lambda_g(\theta G^{n-1}(\theta))'}{1+\lambda_m\eta_m}$$

is a first-order linear differential equation, which solves to

$$b(\theta) = \exp\left(\frac{\Lambda_m}{1+\lambda_m\eta_m}G^{n-1}(\theta)\right)\left(\int_0^\theta \frac{1+\eta_g+\Lambda_g(xG^{n-1}(x))'}{1+\lambda_m\eta_m}\exp\left(-\frac{\Lambda_m}{1+\lambda_m\eta_m}G^{n-1}(x)\right)dx + C\right),$$

where C is the constant of integration. Since $G(x) = 0$ for $x \leq \theta^{min}$, we have

$$b(\theta^{\min}) = \exp(0)\left(\int_0^{\theta^{\min}} \frac{1+\eta_g}{1+\lambda_m\eta_m}\exp(0)dx + C\right) = \frac{1+\eta_g}{1+\lambda_m\eta_m}\theta^{\min} + C.$$

To determine C, we insert $\theta^{\min}$ into equation (2.2) and obtain that

$$0 = (\theta^{\min} - b(\theta^{\min})) + (-\lambda_m\eta_m b(\theta^{\min})) + \eta_g\theta^{\min},$$

or equivalently

$$b(\theta^{\min}) = \frac{1+\eta_g}{1+\lambda_m\eta_m}\theta^{\min},$$

which shows that $C = 0$. Now we can use partial integration in order to rewrite the solution into

$$b(\theta) = \frac{1+\eta_g+\Lambda_g G^{n-1}(\theta)}{1+\lambda_m\eta_m}\theta + \int_0^\theta \frac{\Lambda_m(1+\eta_g+\Lambda_g G^{n-1}(x))}{(1+\lambda_m\eta_m)^2} x \exp\left(\frac{\Lambda_m}{1+\lambda_m\eta_m}(G^{n-1}(\theta) - G^{n-1}(x))\right) dG(x).$$

Since $G(x) = 0$ for all $x \leq \theta^{\min}$, we finally have

$$b(\theta) = \frac{1+\eta_g+\Lambda_g G^{n-1}(\theta)}{1+\lambda_m\eta_m}\theta + \int_{\theta^{\min}}^\theta \frac{\Lambda_m(1+\eta_g+\Lambda_g G^{n-1}(x))}{(1+\lambda_m\eta_m)^2} x \exp\left(\frac{\Lambda_m}{1+\lambda_m\eta_m}(G^{n-1}(\theta) - G^{n-1}(x))\right) dG(x).$$

For sufficiency note first that $b'(\theta)$ is differentiable, since $g(\theta)$ is—by assumption—differentiable. It follows that

$$h(b(\theta)) = \frac{(G^{n-1})'(\theta)}{b'(\theta)}$$

is differentiable as well. Now it is immediate that

$$\begin{aligned}
&\left.\frac{\partial^2 u(x,\theta|b(\theta))}{(\partial x)^2}\right|_{x=b(\theta)}\\
=&\left.\frac{\partial}{\partial x}\left(h(x)\frac{\partial u(x,\theta|b(\theta)/\partial x)}{h(x)}\right)\right|_{x=b(\theta)}\\
=&h'(b(\theta))\underbrace{\left.\left(\frac{\partial u(x,\theta|b(\theta)/\partial x)}{h(x)}\right)\right|_{x=b(\theta)}}_{=0}\\
&+h(b(\theta))\underbrace{\left[-1+\int_0^{b(\theta)}-\lambda_m\eta_m dH(t)-\lambda_m\eta_m(1-H(b(\theta)))\right]}_{<0}\\
<&\ 0.
\end{aligned}$$

□

Proof of Lemma 2.5. Assume the clock increases in increments of ε and the bidder plans to bid up to $x \in (a, b)$. Assume the clock price is $x - \varepsilon$, and the opponent has not dropped out yet. We analyze bidders incentives to bid at x given her plan to do so.

Let $\Delta = \Delta(\varepsilon) = \frac{F(x)-F(x-\varepsilon)}{1-F(x-\varepsilon)}$ be the probability that the opponent drops out at x, given he is still in at $x - \varepsilon$. This means the bidder beliefs to win the auction and get a payoff of $(\theta, -(x - \varepsilon))$ with probability Δ. If the bidder bids at x she receives a utility of

$$\begin{aligned}
u(x,\theta,F|x) =&\underbrace{\Delta(\theta-(x-\varepsilon))}_{\text{classical utility}}+\Delta\underbrace{(1-\Delta)(\eta_g\theta-\lambda_m\eta_m(x-\varepsilon))}_{\text{gain-loss of winning the auction}}\\
&+(1-\Delta)\underbrace{\Delta(-\lambda_g\eta_g\theta+\eta_m(x-\varepsilon))}_{\text{gain-loss of losing the auction}}.
\end{aligned}$$

If she drops out before bidding x, she receives

$$u(x-\varepsilon,\theta,F|x) = \underbrace{\Delta(-\lambda_g\eta_g\theta+\eta_m(x-\varepsilon))}_{\text{gain-loss of losing the auction}}.$$

If bidding up to x is time consistent, then

$$u(x,\theta,F|x) \geq u(x-\varepsilon,\theta,F|x).$$

This is equivalent to

$$\Delta[\theta - (x - \varepsilon) + (1 - \Delta)(\eta_g\theta - \lambda_m\eta_m(x - \varepsilon)) - \Delta(-\lambda_g\eta_g\theta + \eta_m(x - \varepsilon))] \geq 0.$$

Since F has a positive density, we have $\Delta > 0$, and it follows

$$(1 + \eta_g)\theta - (1 + \lambda_m\eta_m)(x - \varepsilon) + \Delta(\Lambda_g\theta + \Lambda_m(x - \varepsilon)) \geq 0.$$

Since F has no atoms, $\lim_{\varepsilon\to 0} \Delta(\varepsilon) = 0$. Thus, in the limit as the increment size goes to zero, we obtain

$$(1 + \eta_g)\theta - (1 + \lambda_m\eta_m)x \geq 0,$$

or equivalently

$$x \leq \frac{1 + \eta_g}{1 + \lambda_m\eta_m}\theta.$$

□

Proof of Proposition 2.4. I sketch the main steps of the proof. If a bidder wins the auction, he has to pay $\max\{b, \overline{x}\}$ with b being the maximal opponent bid. Given opponents' strategies, let $H_{RP}(b)$ be the distribution of the maximal opponent bid with reserve price $\overline{x}$. By replacing s with $\max\{s, \overline{x}\}$ and t with $\max\{t, \overline{x}\}$ in the utility function in section 2.6.1, the utility of a bidder of type θ who bids b with a reference point as if bidding b^* is

$$\begin{aligned}
u(b, \theta | b^*) = &\int_0^b \left(-\max\{s, \overline{x}\} + \int_0^{b^*} \mu_m(\max\{t, \overline{x}\} - \max\{s, \overline{x}\}) dH_{RP}(t) \right) dH_{RP}(s) \\
&+ \int_0^b \int_{b^*}^{\infty} \mu_m(-\max\{s, \overline{x}\}) dH_{RP}(t) dH_{RP}(s) \\
&+ \int_b^{\infty} \left(\int_0^{b^*} \mu_m(\max\{t, \overline{x}\}) dH_{RP}(t) + \int_{b^*}^{\infty} \mu_m(0) dH_{RP}(t) \right) dH_{RP}(s) \\
&+ \int_0^b \left(\theta + \int_0^{b^*} \mu_g(0) dH_{RP}(t) + \int_{b^*}^{\infty} \mu_g(\theta) dH_{RP}(t) \right) dH_{RP}(s) \\
&+ \int_b^{\infty} \left(\int_0^{b^*} \mu_g(-\theta) dH_{RP}(t) + \int_{b^*}^{\infty} \mu_g(0) dH_{RP}(t) \right) dH_{RP}(s).
\end{aligned}$$

Following the derivation of the necessary condition for a symmetric increasing equilibrium[13] in the proof of Theorem 2.4 with this modified utility function, we obtain

[13]The proof of existence of a symmetric increasing continuous equilibrium bidding function $b_{RP}(\theta)$ with reserve price, and its uniqueness for θ with $b(\theta) \geq \overline{x}$ is omitted. It is a modification of Proof 2.4.

for all θ with $b_{RP}(\theta) \geq \overline{x}$ the following modification of equation (2.2):

$$0 = (1+\eta_g)\theta - (1+\lambda_m\eta_m)b_{RP}(\theta) + \Lambda_g H_{RP}(b_{RP}(\theta))\theta + \Lambda_m \int_0^{b_{RP}(\theta)} y dH_{RP}(y) + H_{RP}(\overline{x})\overline{x}.$$

Rearranging yields

$$b_{RP}(\theta) = \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\theta + \Lambda_g H_{RP}(b_{RP}(\theta)) + \Lambda_m \int_{\overline{x}}^{b_{RP}(\theta)} y dH_{RP}(y) + H_{RP}(\overline{x})\overline{x}\right).$$

Let $\overline{\theta}$ be defined by $b(\overline{\theta}) = \overline{x}$. We need to show that $b(\theta) < b_{RP}(\theta)$ for any $\theta \geq \overline{\theta}$. Assume otherwise, and let $\tilde{\theta} = \min\{\theta \in [\overline{\theta}, \theta^{\max}] | b(\theta) \geq b_{RP}(\theta)\}$. The minimum exists by continuity of b and b_{RP}. Now we have

$$\begin{aligned}
b(\tilde{\theta}) &= \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\tilde{\theta} + \Lambda_g H(b(\tilde{\theta})) + \Lambda_m \int_0^{b(\tilde{\theta})} y dH(y)\right) \\
&< \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\tilde{\theta} + \Lambda_g H(b(\tilde{\theta})) + \Lambda_m \int_{b(\overline{\theta})}^{b(\tilde{\theta})} y dH(y) + H(\overline{x})\overline{x}\right) \\
&= \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\tilde{\theta} + \Lambda_g G^{n-1}(\tilde{\theta}) + \Lambda_m \int_{\overline{\theta}}^{\tilde{\theta}} b(s) dG^{n-1}(s) + G^{n-1}(\overline{\theta})\overline{x}\right) \\
&\leq \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\tilde{\theta} + \Lambda_g G^{n-1}(\tilde{\theta}) + \Lambda_m \int_{\overline{\theta}}^{\tilde{\theta}} b_{RP}(s) dG^{n-1}(s) + G^{n-1}(\overline{\theta})\overline{x}\right) \\
&= \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\tilde{\theta} + \Lambda_g H_{RP}(b_{RP}(\tilde{\theta})) + \Lambda_m \int_{b_{RP}(\overline{\theta})}^{b_{RP}(\tilde{\theta})} y dH_{RP}(y) + H_{RP}(b_{RP}(\overline{\theta}))\overline{x}\right) \\
&\leq \frac{1}{1+\lambda_m\eta_m}\left((1+\eta_g)\tilde{\theta} + \Lambda_g H_{RP}(b_{RP}(\tilde{\theta})) + \Lambda_m \int_{\overline{x}}^{b_{RP}(\tilde{\theta})} y dH_{RP}(y) + H_{RP}(\overline{x})\overline{x}\right) \\
&= b_{RP}(\tilde{\theta}),
\end{aligned}$$

a contradiction. □

Proof of Proposition 2.5. For any given opponent strategy distribution F, the implementation of a reserve price $\overline{x}$ is perceived by the bidder as if playing against a distribution

$$F_{RP}(z) = \begin{cases} 0, & z < \overline{x}, \\ F(z), & z \geq \overline{x}. \end{cases}$$

In particular $F_{RP}(z) = F(z)$ for all $z \geq \overline{x}$. Following Lemma 2.4, a strategy $x > \overline{x}$ is a PE if and only if

1. $x \leq (1+\eta)\theta$

2. For any $y \in [x, (1+\eta)\theta]$ we have $u_0(x, \theta, F|x) \geq u_0(y, \theta, F|x)$.

Since for any $y, x > \overline{x}$ we have

$$u_0(y, \theta, F_{RP}|x) = u_0(y, \theta, F|x),$$

these conditions remain unchanged under a reserve price of $\overline{x}$. Therefore, the set of symmetric equilibria for two loss-averse bidders remains unchanged as well. □

Proof of Lemma 2.6. From the perspective of a representative bidder, we denote with $F(x)$ the distribution of prices, at which a particular opponent drops out, i.e. $F(b(\theta)) = G(\theta)$. Similarly we denote with $F_t(x)$ the distribution of drop-out prices of the remaining opponent, given the other opponent drops out at t. Since the remaining opponent j didn't drop out until t, his type θ_j necessarily satisfies $\theta_j > \theta(t)$, and therefore

$$F_t(b(t,\theta)) = Prob(\theta_j \leq \theta | \theta_j > \theta(t)) = \frac{G(\theta) - G(\theta(t))}{1 - G(\theta(t))}.$$

If we denote with $L_{2,t}$ expected gain-loss utility in the two-bidder subgame following an opponent's drop out at price t, then by Proposition 2.2

$$\begin{aligned} L_{2,t}(\theta) &= \ln\left(\frac{1 - F_t(b(t,\theta))}{1 - F_t(t)}\right) \frac{1 - F_t(b(t,\theta))}{1 - F_t(t)} \Lambda\theta \\ &= \ln(1 - F_t(b(t,\theta))(1 - F_t(b(t,\theta))\Lambda\theta \\ &= \ln\left(\frac{1 - G(\theta)}{1 - G(\theta(t))}\right) \frac{1 - G(\theta)}{1 - G(\theta(t))}. \end{aligned}$$

For the 3-bidder auction leading to the first drop out, consider first price increments of ε. Suppose the clock is at price s and both opponents are still remaining. Since we restrict to symmetric increasing bidding functions, a bidder of type θ wins the auction if and only if both opponents have a type lower that θ. Given that they didn't drop out until s, this holds true with probability $\left(\frac{G(\theta) - G(\theta(s))}{1 - G(\theta(s))}\right)^2$.

The probability that a particular opponent j drops out at the next increment is

$$\Delta(s) = \frac{F(s+\varepsilon) - F(s)}{1 - F(s)}.$$

At the next increment $s + \varepsilon$ there are three possibilities:

- With probability $(\Delta(s))^2$ both opponents drop out. The bidder wins with certainty, which induces a gain of $\left(1 - \left(\frac{G(\theta) - G(\theta(s))}{1 - G(\theta(s))}\right)^2\right)\eta\theta$.

- With probability $2\Delta(s)(1-\Delta(s))$ exactly one opponent drops out. The bidder updates her belief to win, which induces a gain of $\left(\frac{G(\theta)-G(\theta(s+\varepsilon))}{1-G(\theta(s+\varepsilon))} - \left(\frac{G(\theta)-G(\theta(s))}{1-G(\theta(s))}\right)^2\right)\eta\theta$.

- With probability $(1-\Delta(s))^2$ no opponent drops out, which induces a loss of $\left(\left(\frac{G(\theta)-G(\theta(s+\varepsilon))}{1-G(\theta(s+\varepsilon))}\right)^2 - \left(\frac{G(\theta)-G(\theta(s))}{1-G(\theta(s))}\right)^2\right)\lambda\eta\theta$.

Since F is continuous, $\Delta(s)$ approaches zero, as the increment size goes to zero. Therefore, in the limit for the continuous English auction, the probability that both opponents drop out at the same time is of second order and has no impact on expected gain-loss utility. Applying Proposition 2.1, expected gain-loss utility in the increment from s to $s+\varepsilon$ for small ε with both opponents being active approaches

$$L_{s+\varepsilon}(\theta)-L_s(\theta) = -2\Delta(s)(1-\Delta(s))\left(\frac{G(\theta)-G(\theta(s+\varepsilon))}{1-G(\theta(s+\varepsilon))} - \left(\frac{G(\theta)-G(\theta(s))}{1-G(\theta(s))}\right)^2\right)\Lambda\theta.$$

As the increment size goes to zero, in the limit the marginal expected gain-loss utility with both opponents being active at time s is given by

$$\ell(s)(\theta) = \frac{-2f(s)}{1-F(s)}\left(\frac{G(\theta)-G(\theta(s))}{1-G(\theta(s))} - \left(\frac{G(\theta)-G(\theta(s))}{1-G(\theta(s))}\right)^2\right)\Lambda\theta.$$

At time t, the probability that time $s>t$ is reached without at least one opponent drop out is $\left(\frac{1-F(s)}{1-F(t)}\right)^2$. Consequently the marginal probability of a drop out at s—which triggers the 2-bidder auction with expected loss $L_{2,s}$—is

$$\frac{\partial}{\partial s}\left(\frac{(1-F(s))^2}{(1-F(t))^2}\right) = \frac{2f(s)(1-F(s))}{(1-F(t))^2}.$$

Putting the two sources of gain-loss utility together and integrating over s yields

$$\begin{aligned} L_t(\theta) &= \int_t^{b(\theta)}\left(\left(\frac{1-F(s)}{1-F(t)}\right)^2\ell(s) + \frac{2f(s)(1-F(s))}{(1-F(t))^2}L_{2,s}(\theta)\right)ds \\ &= -\Lambda\theta\int_t^{b(\theta)}\frac{2f(s)(1-F(s))}{(1-F(t))^2}\left(\frac{G(\theta)-G(\theta(s))}{1-G((\theta(s))} - (\frac{G(\theta)-G(\theta(s))}{1-G(\theta(s))})^2\right)ds \\ &\quad + \Lambda\theta\int_t^{b(\theta)}\frac{2f(s)(1-F(s))}{(1-F(t))^2}\ln\left(\frac{1-G(\theta)}{1-G(\theta(s))}\right)\frac{1-G(\theta)}{1-G(\theta(s))}ds \end{aligned}$$

Since $F(s) = G(\theta(s))$ and consequently $f(s) = g(\theta(s))/b'(\theta(s))$, integration by substitution yields

$$L_t(\theta) = -\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(1-G(s))}{(1-G(\theta(t)))^2} \left[\frac{G(\theta)-G(s)}{1-G(s)} - \left(\frac{G(\theta)-G(s)}{1-G(s)} \right)^2 - \ln \left(\frac{1-G(\theta)}{1-G(s)} \right) \frac{1-G(\theta)}{1-G(s)} \right] ds$$

□

Proof of Corollary 2.4. Define

$$\delta(s) = \frac{G(\theta)-G(s)}{1-G(s)}$$

Since for $\theta < \theta^{\max}$ we have $\delta(s) < 1$, and we can use the power series of the logarithm to rewrite

$$L_t(\theta) = -\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(1-G(s))}{(1-G(\theta(t)))^2} \left[\delta(s) - (\delta(s))^2 - (-\delta(s) - \frac{\delta(s)^2}{2} - \frac{\delta(s)^3}{3} ...)(1-\delta(s)) \right] ds$$

Since $\lim_{s \to \theta} \delta(s) = 0$, we have

$$\begin{aligned}
&\lim_{t \to b(\theta)} \frac{L_t(\theta)}{\left(\frac{G(\theta)-G(\theta(t))}{1-G(\theta(t))} \right)^2} \\
&= \lim_{t \to b(\theta)} -\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(1-G(s))}{(G(\theta)-G(\theta(t))^2} \left[\delta(s) - (\delta(s))^2 - (-\delta(s) - \frac{\delta(s)^2}{2} ...)(1-\delta(s)) \right] ds \\
&= \lim_{(\theta(t) \to \theta} -\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(1-G(s))}{(G(\theta)-G(\theta(t))^2} \left[\delta(s) - (\delta(s))^2 - (-\delta(s) - \frac{\delta(s)^2}{2} ...)(1-\delta(s)) \right] ds \\
&= \lim_{\theta(t) \to \theta} -\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(1-G(s))}{(G(\theta)-G(\theta(t))^2} 2\delta(s) ds \\
&= \lim_{\theta(t) \to \theta} -2\Lambda\theta \int_{\theta(t)}^{\theta} \frac{2g(s)(G(\theta)-G(s))}{(G(\theta)-G(\theta(t))^2} ds \\
&= \lim_{\theta(t) \to \theta} -2\Lambda\theta \left[\frac{-(G(\theta)-G(s))^2}{(G(\theta)-G(\theta(t))^2} \right]_{\theta(t)}^{\theta} \\
&= \lim_{\theta(t) \to \theta} -2\Lambda\theta \\
&= -2\Lambda\theta
\end{aligned}$$

Now, since $b(t,\theta)$ is continuous in t, $\lim_{t \to b(\theta)} b(t,\theta)$ exists. We prove the threshold of time-consistent behavior for $(\theta^{\min}, \theta^{\max})$ by contradiction. For the boundaries it follows by continuity. Assume that there is some $\overline{\theta} \in (\theta^{\min}, \theta^{\max})$ with

$$\lim_{t \to b(\overline{\theta})} b(t, \overline{\theta}) > (1 + \eta - \Lambda)\overline{\theta}.$$

Since $b(t,\theta)$ is continuous there is some $\hat{t} < b(\overline{\theta})$ and $\hat{\theta} \in [\theta(\hat{t}), \overline{\theta}]$, such that

$$b(t,\theta) > (1+\eta-\Lambda)\overline{\theta}$$

for all $t \in [\hat{t}, b(\overline{\theta})]$, $\theta \in [\hat{\theta}, \overline{\theta}]$. This implies that the sales price for the good exceeds $(1+\eta-\Lambda)\overline{\theta}$ if no bidder drops out until $\hat{t}$. If $b(t,\theta)$ is a time-consistent strategy, then at time $\hat{t}$ a bidder of type $\overline{\theta}$ must weakly prefer this strategy to an instantaneous drop out. Since at time $\hat{t}$ her belief to win is $\left(\frac{G(\overline{\theta})-G(\theta(\hat{t}))}{1-G(\theta(\hat{t}))}\right)^2$, this condition reads

$$-\lambda\eta\overline{\theta}\left(\frac{G(\overline{\theta})-G(\theta(\hat{t}))}{1-G(\theta(\hat{t}))}\right)^2 < \left(\frac{G(\overline{\theta})-G(\theta(\hat{t}))}{1-G(\theta(\hat{t}))}\right)^2 (\overline{\theta}-(1+\eta-\Lambda)\overline{\theta}) + L_{\hat{t}}(\theta),$$

with strict inequality since the price strictly exceeds $(1+\eta-\Lambda)\overline{\theta}$. This is equivalent to

$$L_{\hat{t}}(\overline{\theta}) > -2\Lambda\overline{\theta}\left(\frac{G(\overline{\theta})-G(\theta(\hat{t}))}{1-G(\theta(\hat{t}))}\right)^2,$$

a contradiction for $\hat{t}$ sufficiently close to $b(\overline{\theta})$. □

Chapter 3

Persuasion against Self-Control Problems

This Chapter is based on von Wangenheim (2018).

3.1 Introduction

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

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

I draw a connection between the dynamic consumption model with quasi-hyperbolic discounting consumers by Carrillo and Mariotti (2000) and the optimal information design approach by Kamenica and Gentzkow (2011), and derive a consumer-optimal information signal about the consumers' risk type. The derived optimal signal consists of a simple binary signal, which displays whether or not the risk is below some threshold. It is therefore consistent with commonly observed recommendations by experts as well as certified product labels which require products to satisfy legally defined limits on harmful substances.

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

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

I follow the normative approach suggested by O'Donoghue and Rabin (2002, 2003) in defining consumer welfare as ex-ante consumer utility before facing consumption decisions.[1] As a social planner therefore takes the perspective of the consumer at time 0, the model can equivalently be regarded as a model of self-persuasion, in which not the regulator but the consumer himself strategically acquires information some time before he faces the consumption decision.

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

First, I analyze the consumer-optimal signal when the consumer has no access to further information before making the consumption decisions. As the incentives of a welfare maximizing planner are aligned with the consumer whenever he does not face an instantaneous consumption decision, this seems appropriate where information is not immediately available at reasonable cost. For instance, it seems very implausible that a consumer would first conduct a medical test on his risk type, whenever he instantaneously faces the opportunity to smoke.

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

There are numerous examples where institutions implicitly use cut-off signals in form of recommendations, guidelines or definitions, when consumers themselves fail

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

to thoroughly interpret data. Definitions of obesity, limits for responsible drinking, classification of medical risk groups are only some examples. Moreover, certification in form of labels—such as the European Union eco label and energy label—are commonly based on predefined thresholds.

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

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

The remainder of this chapter is structured as follows: In Section 3.2, I discuss the related literature. In Section 3.3, I introduce the model. I start Section 3.4 with the benchmark of full information, before I analyze the case of full information control by the planner, and the case of costless consumer learning. Section 3.5 concludes this chapter. All proofs are relegated to the appendix.

3.2 Related Literature

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

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

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

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

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

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

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

3.3 The Model

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

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

The probability θ of a realization of the externality is unknown to the consumer. It is distributed according to some prior distribution with cdf F_0 and continuous, positive density $f_0(\theta)$ on support $[0, 1]$.[3] However, at time 0 a regulator on behalf of the consumer may design an information signal about θ, which the consumer observes at no cost. In the beginning of any subsequent period the consumer may acquire further information about θ. Details on the informational process are explained after the characterization of the intrapersonal conflict.

Intertemporal payoffs and intrapersonal conflict

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

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

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

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

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

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

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

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

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

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

Information Provision and Learning

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

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

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

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

In many circumstances it seems plausible that the regulator has exclusive control over information in the sense that it is too costly for the agent to acquire additional information himself whenever he faces a consumption decision.[4]

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

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

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

Equilibrium Concept

In each period $t \geq 1$, a strategy for the consumer consists of a learning decision in form of a signal S_t, and the consumption decision $x_t \in \{0,1\}$, based on the posterior derived from updating with respect to the signal. Since the payoff relevant information at the beginning period t is captured by belief F_t, and is independent of time, it is natural to focus on Markov strategies.

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

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

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

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

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

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

Assumption 3.1. The consumer breaks any tie in favor of the regulator: Whenever the consumer at any time is indifferent between preferred signal structures or the two consumption decision, he takes the decision that maximizes utility of self-0.

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

3.4 Analysis

Let

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

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

Assumption 3.2. If the consumer knows with certainty that the externality realizes he prefers to abstain, i.e.

$$\beta C > 1.$$

The full learning benchmark

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

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

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

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

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

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

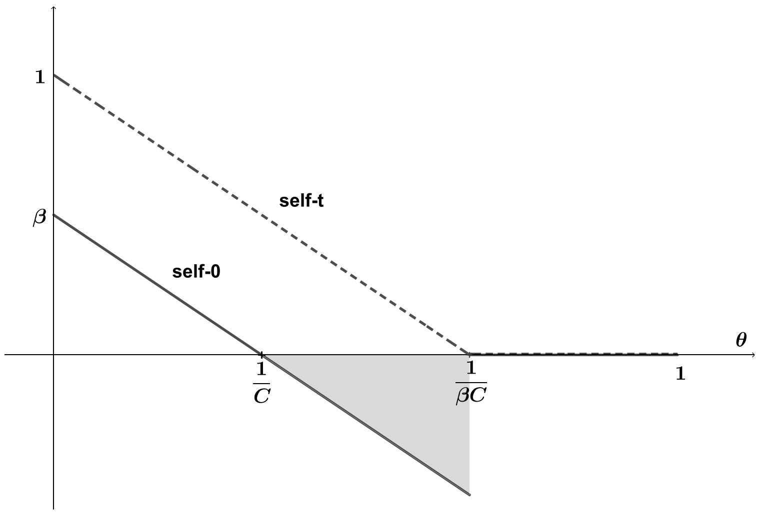

Figure 3.1: Welfare loss under complete information

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

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

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

Compared to the first-best utility

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

the commitment problem induces a welfare loss of

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

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

Full Information Control

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

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

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

Proposition 3.1. *If the regulator has full control over information, the welfare maximizing signal structure is described by a threshold $y \in \left[\frac{1}{C}, \frac{1}{\beta C}\right)$, and the signal*

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

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

1. *If*

$$\mathbb{E}\left[\theta \middle| \theta > \frac{1}{C}\right] \geq \frac{1}{\beta C}$$

then $y = \frac{1}{C}$, and the signal induces first best welfare.

2. If

$$\mathbb{E}\left[\theta \middle| \theta > \frac{1}{C}\right] < \frac{1}{\beta C}$$

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

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

The optimal signal can be interpreted as an incentive compatible recommendation by the regulator to the consumer. Given consumer's risk type, the regulator recommends whether to consume or abstain, and the consumer wishes to follow the recommendation.

There are numerous real-world examples where regulators give recommendations in form of consumption thresholds. For instance, many governments have adopted guidelines that define thresholds for responsible alcohol consumption (Kalinowski and Humphreys (2016)). the world health organization (WHO) defines a body mass index of $25kg/m^2$ as the cut-off point for overweight.[5]

Moreover, labels are a common tool to provide consumption recommendation on the basis of thresholds. For instance, in 2006 the Food Standard Agency (FSA) in the UK introduced a traffic light rating system for food nutrition values. Products with sugar or saturated fats above certain thresholds are highlighted with a red light to display the health risk of the product. Further, the organic food label by the European Union defines minimum requirements for food to be labelled as organic.[6] Proposition 3.1 describes how to choose the respective thresholds in order to induce the welfare maximizing consumption decision for consumers.

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

[5] `http://apps.who.int/iris/bitstream/10665/37003/1/WHO_TRS_854.pdf`

[6] (EC) No. 889/2008 and (EC) No. 834/2007

Lemma 3.1. *If all consumers have to take the same consumption decision rule* $x(\theta) \in \{0,1\}$*, then the optimal decision rule from the perspective of self-t for* $t \geq 1$ *is*

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

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

Following the argument of Proposition 3.1 with this objective, one immediately obtains

Corollary 3.1. *If self-t for* $t \geq 1$ *has full control over information provided by the regulator, the welfare maximizing signal structure is described by a threshold* $y \in \left[\delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}, \frac{1}{\beta C}\right)$*, and the signal*

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

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

1. *If*
$$\mathbb{E}\left[\theta \middle| \theta > \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}\right] \geq \frac{1}{\beta C}$$
then $y = \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}$*, and the signal implements self-t's preferred consumption decision rule.*

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

Costless Learning

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

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

Definition 3.4. Consider a Markov perfect equilibrium. An information signal S_F is relevant, if acquiring S_F when the belief is F at any time t induces different expected utility to self-t or self-0 than acquiring no information.

Lemma 3.2. *On equilibrium path in a Markov perfect equilibrium the consumer will not acquire relevant information signals at any time $t > 1$.*

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

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

Since no incarnation wants to know less than his successor, each incarnation will acquire full information himself whenever the successor would do so. As an immediate consequence, the time-independent strategy of always acquiring full information regardless of the current belief forms a Markov perfect equilibrium.[7]

Corollary 3.2. *For any subgame starting at any time $t \geq 0$ the Markov strategy of a full information signal*

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

for all beliefs F constitutes a Markov perfect equilibrium.

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

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

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

Lemma 3.3. *For any $t \geq 0$, the Markov perfect equilibrium of the subgame starting time $t+1$ which minimizes the utility of self-t is given by the full information equilibrium, where each self-s for $s \geq t+1$ for every belief chooses a full information signal.*

The equilibrium in Proposition 3.1 where the regulator has full control over information can be sustained with costless consumer learning if the one time gain from acquiring a preferable information structure does not exceed the loss from being stuck in the full information equilibrium afterwards. Proposition 3.2 gives precise conditions when this is the case.

Proposition 3.2.

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

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

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

2. *The cut-off y and the consumption decisions coincide with those in Proposition 3.1, where the regulator has full control over information, if and only if the optimal cut-off in Proposition 3.1 satisfies*

$$\mathbb{E}\left[\theta \middle| \theta \in \left[y, \frac{1}{\beta C}\right]\right] \geq \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C}. \tag{3.1}$$

3. *Otherwise, if condition (3.1) is not satisfied then y is uniquely determined by*

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

   Welfare is strictly higher than under full information and strictly higher than in the case of Corollary 3.1, where the consumer at time $t = 1$ chooses the signal structure.

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

On equilibrium path, the consumer does not exert learning at any t, and consumes according to signal S whenever $\theta < y$. If some consumer incarnation deviates

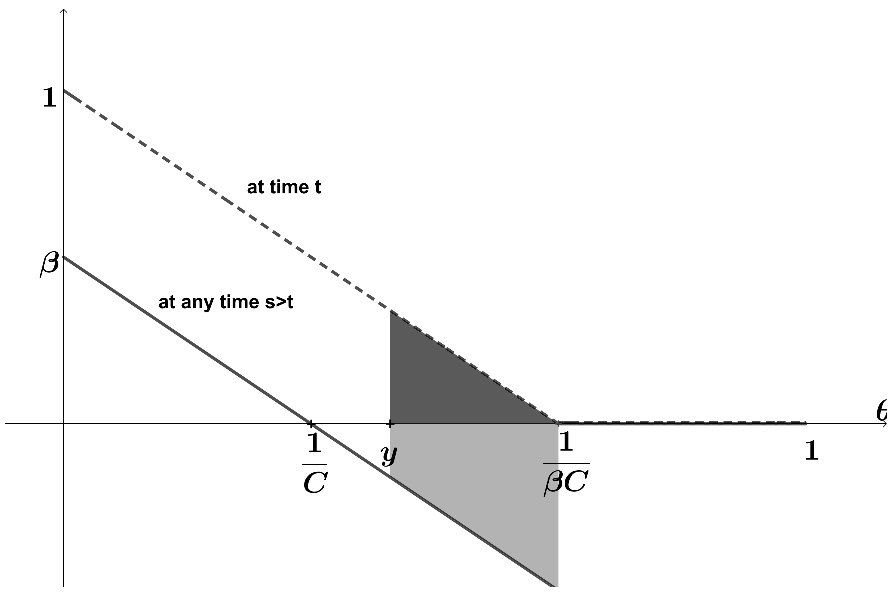

Figure 3.2: Gain and Loss of a Deviation from Self-t's Perspective

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

While there are many other Markov perfect equilibria for the regulator's signal in period 0 (including the full information equilibrium), the described equilibrium is not only welfare maximizing, but also arises naturally in this context. The information signal in period 0 can be regarded as an incentive compatible recommendation of a default option by the regulator: All consumer incarnations find it optimal to take the information as given and base their consumption decision on the implied recommendation.[8]

Next, we look at sufficient conditions for the distribution to satisfy condition

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

3.1 for the optimal cut-off y, so that the outcome under full information control and in the PMPE under costless learning coincide.

Recall from Proposition 3.1 that the optimal threshold y under full information control satisfies $y \in [\frac{1}{C}, \frac{1}{\beta C}]$. Consequently, $\mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]]$ is in $[\frac{1}{C}, \frac{1}{\beta C}]$ with the exact value depending on where the prior distribution on $[\frac{1}{C}, \frac{1}{\beta C}]$ has most of its mass. Intuitively, in order to satisfy

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

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

Corollary 3.3. *If the prior is nondecreasing on $[\frac{1}{C}, \frac{1}{\beta C}]$, and $\delta \geq \frac{1}{2}$, then Condition 3.1 in Proposition 3.2 is satisfied and the PMPE with costless learning coincides with the equilibrium with full informational control as described in Proposition 3.1.*

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

More generally, a sufficiently high δ always relaxes Condition 3.1, made precise in the following Corollary.[9]

Corollary 3.4. *For any prior distribution there exists a $\overline{\delta} < 1$ such that for all $\delta \in [\overline{\delta}, 1]$ Condition 3.1 in Proposition 3.2 is satisfied and the PMPE with costless learning coincides with the equilibrium with full informational control as described in Proposition 3.1.*

3.5 Conclusion

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

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

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

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

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

3.6 Appendix

Proof of Proposition 3.1. Let S be the signal structure chosen by the regulator, let $G(\theta, s)$ be its joint distribution on $[0,1] \times S$. Since the consumer has no access to further information at $t \geq 1$, a strategy for the consumer consists of a sequentially rational consumption decision $x(s) \in \{0,1\}$ for all signal realizations $s \in S$. Sequential utility maximization requires

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

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

$$U_0((S,G)) = \frac{\beta\delta}{1-\delta} \int_{(\theta,s)\in[0,1]\times S} x(s)(1 - C\theta) dG(\theta, s). \tag{3.2}$$

This problem is a classical persuasion problem as defined in Kamenica and Gentzkow (2011) with continuous state space $[0,1]$. In Proposition 3 of their Web Appendix they show that an optimal signal exists for such persuasion problems. Further, according to their Proposition 1, we can restrict to information signals $S = \{s_1, s_2\}$, where realization s_1 induces consumption, while realization s_2 induces abstention.[10] Let from now $S = \{s_1, s_2\}$ with distribution $G(\theta, s)$ be an optimal F_0-consistent signal structure.

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

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

Since

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

we have

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

Consequently, the consumer abstains for $\theta > y$ under the cut-off signal, whenever he abstains for $S = s_2$ under signal S. Analogously, one can show that since $\mathbb{E}[\theta|S = s_1] > \mathbb{E}[\theta|\theta \leq y]$, the consumer consumes for $\theta \leq y$ under the cut-off signal whenever he consumes for $S = s_1$ under signal S. Plugging this decision rule into the regulator's objective (3.2), and using $\mathbb{E}[\theta|S = s_1] > \mathbb{E}[\theta|\theta \leq y]$, we see that the regulator's utility under the cut-off signal

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

exceeds her utility under S.

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

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

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

Proof of Lemma 3.1. Utility of self-t from consumption with risk type θ is

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

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

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

thus if and only if

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

□

Proof of Corollary 3.1. The proof is identical to the proof of Proposition 3.1, where we replace 'regulator' with 'self-1' and the regulator's objective $U_0((S, G))$ by self-1's objective

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

In particular, the optimal signal is a cut-off signal, which aims to induce the most possible abstention for types $\theta \leq (1-\delta)\frac{1}{\beta C} + \delta\frac{1}{C}$ as calculated in Lemma 3.1. Consequently, the optimal cut-off y satisfies

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

and the result follows. □

Proof of Lemma 3.2. We start by defining the collapse of two signal structures. Let S be an F-consistent signal structure with distribution G. For some signal realization $s_k \in S$ let $\tilde{F} = G(\cdot|s_k)$ be the posterior distribution. Further, let $\tilde{S}$ be an $\tilde{F}$-consistent signal structure with distribution $\tilde{G}$. The collapse of S and $\tilde{S}$ is the signal structure with signal space $\overline{S} = (S\backslash\{s_k\}) \sqcup \tilde{S}$ and a joint distribution on $([0,1] \times \overline{S}, \mathcal{B}([0,1] \times \overline{S}))$ defined via

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

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

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

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

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

Suppose now S_{F_t} is relevant for $t > 1$. As self-t is indifferent between S_{F_t} and the trivial signal, this implies that self-0 is not. Since by Assumption 3.1 self-t always chooses self-0's preferred action whenever he is indifferent between his preferred action, self-0 strictly prefers self-t to acquire signal S_{F_t} rather than the trivial signal. This implies he prefers future incarnations to have the distribution of posteriors from S_{F_t} rather than belief F_t. In particular, self-0 prefers self-$t-1$ to acquire the collapse of $S_{F_{t-1}}$ with S_{F_t} rather than his equilibrium choice $S_{F_{t-1}}$. Since self-$t-1$ is indifferent between the two, but chooses $S_{F_{t-1}}$, Assumption 3.1 is violated, a contradiction.

□

Proof of Lemma 3.3. By Corollary 3.2, the full information strategy for the subgame starting at $t+1$ is a Markov perfect equilibrium. Take any other Markov perfect equilibrium of the subgame. Lemma 3.2 states that for the game starting at $t = 1$ with belief F_1 there is only information acquisition at $t = 1$ in equilibrium. By renaming the time index it is immediate that for any subgame starting at $t+1$ with belief F_{t+1} there is only information acquisition at time $t + 1$ in equilibrium. Consequently, the consumption decision is identical at all times starting at $t + 1$. Let $S_{F_{t+1}}$ with distribution G be the information signal at $t + 1$ and $x(s)$ be the consumption decision for signal realization s. Then the expected utility for self-$t+1$ generated by his own consumption decision under signal $S_{F_{t+1}}$ is

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

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

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

Call accordingly

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

and

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

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

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

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

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

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

□

Proof of Proposition 3.2. By Lemma 3.2 we can restrict to Markov equilibria with no information acquisition on equilibrium path at any time $t \geq 1$. Such an equilibrium induces the same consumption decisions for all incarnations of the consumer.

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

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

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

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

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

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

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

Next, we show that a solution to the relaxed problem exists, if and only if it exists in the class of cut-off signals. Let S be a non-cutoff signal and let again $y \in [0, 1]$ be such that $\mathbb{P}(S = s_2) = \mathbb{P}(\theta > y)$. We showed in the proof of Proposition 3.1 that the cut-off signal which displays s_1 if and only if $\theta \leq y$ improves the regulator's utility compared to S and does not change consumer's consumption decision. Moreover this cut-off signal improves self-1's utility: The cut-off signal changes the consumer's action from abstention to consumption whenever $S = s_2$ and $\theta \leq y$. It changes the consumer's action from consumption to abstention whenever $S = s_1$ and $\theta > y$. Since

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

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

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

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

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

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

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

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

or differently if and only if

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

Rearraning this condition yields

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

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

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

which is equivalent to

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

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

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

Since the optimal cut-off from Proposition 3.1 where the consumer cannot acquire information satisfies

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

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

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

which shows 2.

Otherwise, if

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

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

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

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

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

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

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

In particular, this implies that

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

where the right hand side is by Corollary 3.1 the self-1 preferred cut-off. Thus, abstention is higher on $[\frac{1}{C}, \frac{1}{\beta C}]$ than under the self-1 preferred signal from Corollary 3.1 and induces higher welfare. □

Proof of Corollary 3.4. By assumption, the prior has a continuous and strictly positive density f_0, thus attains its minimum $f_0^{\min}$ and its maximum $f_0^{\max}$ on $[0, 1]$. Let $\overline{\delta} = \frac{f_0^{\max}}{f_0^{\max} + f_0^{\min}}$. Rearranging yields

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

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

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

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

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

□

Proof of Corollary 3.3. If the prior is nondecreasing on $[\frac{1}{C}, \frac{1}{\beta C}]$, then the conditional distribution of the prior on $[\frac{1}{C}, \frac{1}{\beta C}]$ first order stochastically dominates the uniform distribution on $[\frac{1}{C}, \frac{1}{\beta C}]$. Hence, for all $\delta \geq \frac{1}{2}$ we have

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

□

Bibliography

Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference points and effort provision. *The American Economic Review*, 101(2):470–492.

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

Balestrieri, F., Izmalkov, S., and Leao, J. (2015). The market for surprises: selling substitute goods through lotteries.

Baron, D. P. and Besanko, D. (1984). Regulation and information in a continuing relationship. *Information Economics and policy*, 1(3):267–302.

Battaglini, M. (2005). Long-term contracting with markovian consumers. *The American Economic Review*, 95(3):637–658.

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

Benartzi, S., Thaler, R. H., et al. (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics*, 110(1):73–92.

Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *The American Economic Review*, 105(3):921–957.

Bergemann, D. and Pesendorfer, M. (2007). Information structures in optimal auctions. *Journal of Economic Theory*, 137(1):580–609.

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

Camerer, C. (1995). Individual decision making. In Kagel, J. H. and Roth, A. E., editors, *The Handbook of Experimental Economics*. Princeton University Press.

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

Courty, P. and Li, H. (2000). Sequential screening. *The Review of Economic Studies*, 67(4):697–717.

Crémer, J., Khalil, F., and Rochet, J.-C. (1998). Contracts and productive information gathering. *Games and Economic Behavior*, 25(2):174–193.

Dato, S., Grunewald, A., Müller, D., and Strack, P. (2017). Expectation-based loss aversion and strategic interaction.

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

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

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

Ehrhart, K. and Ott, M. (2014). Reference-dependent bidding in dynamic auctions.

Eisenhuth, R. (2012). Reference dependent mechanism design. *Job Market Paper*.

Eisenhuth, R. and Ewers, M. (2010). Auctions with loss averse bidders.

Eliaz, K. and Spiegler, R. (2014). Reference dependence and labor market fluctuations. *NBER macroeconomics annual*, 28(1):159–200.

Ericson, K. M. M. and Fuster, A. (2011). Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments. *The Quarterly Journal of Economics*, 126(4):1879–1907.

Ericson, K. M. M. and Fuster, A. (2014). The endowment effect. *Annu. Rev. Econ.*, 6(1):555–579.

Eső, P. and Szentes, B. (2007). Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74(3):705–731.

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

Gill, D. and Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *The American Economic Review*, 102(1):469–503.

Green, C. E. and Lomanno, M. V. (2012). *Distribution channel analysis: A guide for hotels*. HSMAI Foundation.

Grunert, K. G., Fernández-Celemín, L., Wills, J. M., genannt Bonsmann, S. S., and Nureeva, L. (2010). Use and understanding of nutrition information on food labels in six european countries. *Journal of Public Health*, 18(3):261–277.

Hawley, K. L., Roberto, C. A., Bragg, M. A., Liu, P. J., Schwartz, M. B., and Brownell, K. D. (2013). The science on front-of-package food labels. *Public health nutrition*, 16(3):430–439.

Heffetz, O. and List, J. A. (2014). Is the endowment effect an expectations effect? *Journal of the European Economic Association*, 12(5):1396–1422.

Heidhues, P. and Kőszegi, B. (2008). Competition and price variation when consumers are loss averse. *The American Economic Review*, 98(4):1245–1268.

Heidhues, P. and Kőszegi, B. (2014). Regular prices and sales. *Theoretical Economics*, 9(1):217–251.

Herweg, F., Müller, D., and Weinschenk, P. (2010). Binary payment schemes: Moral hazard and loss aversion. *The American Economic Review*, 100(5):2451–2477.

Hoffmann, F. and Inderst, R. (2011). Pre-sale information. *Journal of Economic Theory*, 146(6):2333–2355.

Hoffmann, F., Inderst, R., and Turlo, S. (2017). Regulating cancellation rights with consumer experimentation.

Horowitz, J. K. and McConnell, K. E. (2002). A review of wta/wtp studies. *Journal of Environmental Economics and Management*, 44(3):426–447.

Kagel, J. H. (1995). Auctions: A survey of experimental research. In Kagel, J. H. and Roth, A. E., editors, *The Handbook of Experimental Economics*. Princeton University Press.

Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1990). Experimental tests of the endowment effect and the coase theorem. *Journal of political Economy*, 98(6):1325–1348.

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

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

Kessler, A. S. (1998). The value of ignorance. *The Rand Journal of Economics*, 29(2):339–354.

Kőszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165.

Kőszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. *The American Economic Review*, 97(4):1047–1073.

Kőszegi, B. and Rabin, M. (2009). Reference-dependent consumption plans. *The American Economic Review*, 99(3):909–936.

Krähmer, D. and Strausz, R. (2011). Optimal procurement contracts with pre-project planning. *The Review of Economic Studies*, 78(3):1015–1041.

Krähmer, D. and Strausz, R. (2015). Optimal sales contracts with withdrawal rights. *The Review of Economic Studies*, 82(2):762–790.

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

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

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

Lange, A. and Ratan, A. (2010). Multi-dimensional reference-dependent preferences in sealed-bid auctions–how (most) laboratory experiments differ from the field. *Games and Economic Behavior*, 68(2):634–645.

Lewis, T. R. and Sappington, D. E. (1994). Supplying information to facilitate price discrimination. *International Economic Review*, pages 309–327.

Lewis, T. R. and Sappington, D. E. (1997). Information management in incentive problems. *Journal of political Economy*, 105(4):796–821.

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

Li, H. and Shi, X. (2017). Discriminatory information disclosure. *American Economic Review*, 107(11):3363–85.

Lucking-Reiley, D. (1999). Using field experiments to test equivalence between auction formats: Magic on the internet. *The American Economic Review*, 89(5):1063–1080.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1):58–73.

Myerson, R. B. (1986). Multistage games with communication. *Econometrica*, 54(2):323–358.

Nocke, V., Peitz, M., and Rosar, F. (2011). Advance-purchase discounts as a price discrimination device. *Journal of Economic Theory*, 146(1):141–162.

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

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

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

Pagel, M. (2016). Expectations-based reference-dependent preferences and asset pricing. *Journal of the European Economic Association*, 14(2):468–514.

Pagel, M. (2017). Expectations-based reference-dependent life-cycle consumption. *The Review of Economic Studies*, 84(2):885–934.

Pavan, A., Segal, I., and Toikka, J. (2014). Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653.

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

Roesler, A.-K. and Szentes, B. (2017). Buyer-optimal learning and monopoly pricing. *The American Economic Review*, 107(7):2072–80.

Rosato, A. (2014). Loss aversion in sequential auctions: Endogenous interdependence, informational externalities and the" afternoon effect".

Schindler, J. (2003). *Auctions with interdependent valuations. Theoretical and empirical analysis, in particular of internet auctions.* PhD thesis, WU Vienna University of Economics and Business.

Shapiro, D. and Shi, X. (2008). Market segmentation: The role of opaque travel agencies. *Journal of Economics & Management Strategy*, 17(4):803–837.

Shiryaev, A. N. (1996). *Probability, volume 95 of Graduate texts in mathematics.* Springer-Verlag, New York.

Shogren, J. F., Shin, S. Y., Hayes, D. J., and Kliebenstein, J. B. (1994). Resolving differences in willingness to pay and willingness to accept. *The American Economic Review*, 84(1):255–270.

Simon, L. K. and Stinchcombe, M. B. (1989). Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214.

Smith, A. (2012). Lagged beliefs and reference-dependent preferences.

Strahilevitz, M. A. and Loewenstein, G. (1998). The effect of ownership history on the valuation of objects. *Journal of Consumer Research*, 25(3):276–289.

Szalay, D. (2009). Contracts with endogenous information. *Games and Economic Behavior*, 65(2):586–625.

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

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

Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 106(4):1039–1061.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37.

Viscusi, W. K. (1978). A note on "lemons" markets with quality certification. *The Bell Journal of Economics*, 9(1):277–279.

von Wangenheim, J. (2017a). Consumer-optimal information design.

von Wangenheim, J. (2017b). English versus vickrey auction with loss averse bidders.

von Wangenheim, J. (2018). Persuasion against self-control problems.

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

Selbständigkeitserklärung

Für diese Dissertation habe ich keine anderen Hilfsmittel außer der angeführten Literatur benutzt.
Ich bezeuge durch meine Unterschrift, dass meine Angaben über die bei der Abfassung meiner Dissertation benutzten Hilfsmittel, über die mir zuteil gewordene Hilfe sowie über frühere Begutachtungen meiner Dissertation in jeder Hinsicht der Wahrheit entsprechen.

Berlin, 22. März 2018

Jonas von Wangenheim