

DISCUSSION PAPER SERIES

DP14063
(v. 2)

**LEMONS AND PEACHES: A (ROBUST)
MULTI-STAGE BUYING MECHANISM
WITH MULTIPLE APPLICATIONS**

Hans Gersbach, Akaki Mamageishvili and Oriol
Tejada

INDUSTRIAL ORGANIZATION



LEMONS AND PEACHES: A (ROBUST) MULTI-STAGE BUYING MECHANISM WITH MULTIPLE APPLICATIONS

Hans Gersbach, Akaki Mamageishvili and Oriol Tejada

Discussion Paper DP14063
First Published 17 October 2019
This Revision 22 July 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Hans Gersbach, Akaki Mamageishvili and Oriol Tejada

LEMONS AND PEACHES: A (ROBUST) MULTI-STAGE BUYING MECHANISM WITH MULTIPLE APPLICATIONS

Abstract

We introduce a four-stage, multi-price buying mechanism, which can be used by a (big) buyer to separate low-quality sellers - called "lemon" owners - from high-quality sellers - called "peach" owners. With a partition of sellers, the buyer obtains the commodities from the "peach" owners at a price that matches the willingness to sell. By contrast, "lemon" owners are trapped into selling their items at a low, or even negligible, price. Our mechanism is robust for several extensions of our baseline setup, offers applications for market makers and regulators, and may be used by interest groups in politics.

JEL Classification: C72, D4, D7, D72, D82, D86

Keywords: Lemons market - Partition - Signaling - Commitment - Market maker - Vote-buying - Lobbying - Decoy ballots

Hans Gersbach - hgersbach@ethz.ch
ETH Zurich and CEPR

Akaki Mamageishvili - amamageishvili@ethz.ch
ETH Zurich

Oriol Tejada - toriol@ethz.ch
ETH Zurich

Acknowledgements

We would like to thank Salvador Barberà, David Basin, Afsoon Ebrahimi, Georgy Egorov, Lara Schmid, Larry Samuelson, participants of the 2018 Lisbon Game Theory Meeting, the 15th Conference on Web and Internet Economics, and the Astute Modeling Seminar at ETH Zurich for valuable comments. All errors are ours.

Lemons and Peaches: A (Robust) Multi-stage Buying Mechanism with Multiple Applications*

Hans Gersbach

CER-ETH – Center of Economic
Research at ETH Zurich and CEPR
Zürichbergstrasse 18
8092 Zurich, Switzerland
hgersbach@ethz.ch

Akaki Mamageishvili

CER-ETH – Center of Economic
Research at ETH Zurich
Zürichbergstrasse 18
8092 Zurich, Switzerland
amamageishvili@ethz.ch

Oriol Tejada

CER-ETH – Center of Economic
Research at ETH Zurich
Zürichbergstrasse 18
8092 Zurich, Switzerland
toriol@ethz.ch

This version: July 2020

Abstract

We introduce a four-stage, multi-price buying mechanism, which can be used by a (big) buyer to separate low-quality sellers—called “lemon” owners—from high-quality sellers—called “peach” owners. With a partition of sellers, the buyer obtains the commodities from the “peach” owners at a price that matches the willingness to sell. By contrast, “lemon” owners are trapped into selling their items at a low, or even negligible, price. Our mechanism is robust for several extensions of our baseline setup, offers applications for market makers and regulators, and may be used by interest groups in politics.

Keywords: lemons market; partition; signaling; commitment; market maker; vote-buying; lobbying; decoy ballots

JEL Classification: C72, D4, D7, D72, D82, D86

*We would like to thank Salvador Barberà, David Basin, Afsoon Ebrahimi, Georgy Egorov, Lara Schmid, Larry Samuelson, participants of the 2018 Lisbon Game Theory Meeting, the 15th Conference on Web and Internet Economics, and the Astute Modeling Seminar at ETH Zurich for valuable comments. All errors are ours.

1 Introduction

Motivation and goal

Dating back at least to Akerlof (1970), we know that markets may not function well when they are plagued with adverse selection. When sellers know the quality of their goods, but buyers only have statistical information about this quality in the entire population, in particular, markets may break down or only low-quality goods—usually called *lemons*—may be traded. In this paper, we explore the so-called *Lemons Problem* when the set of sellers is partitioned in a finite number of groups and there is statistical information about the sellers’ distribution of types for each of these groups. Can a seller partition solve the Lemons Problem for the buyer(s)?

Model and main insights

To elaborate, suppose that an agent—the *buyer*—wants to buy a given number of items (for example, used cars) of high quality, and that she can buy them from several individuals—the *sellers*—each of whom owns one item. For the sake of clarity, throughout the paper, we use “she” to refer to the buyer and “he” to refer to any seller. Items can be of high quality or low quality, and the quality of a particular item is private information of the seller. Sellers’ valuations of their own commodity are exogenously given—more on this in Section 3—, are equal across types, and are higher for high-quality items than for low-quality items. The buyer, for her part, values low-quality and high-quality items at least as much as the sellers, and she has a budget for buying (high-quality) items.

We analyze the buyer’s problem under the assumptions that (i) the pool of sellers can be divided into several *groups* according to some characteristic other than item quality (and item valuation), say location or age, and that (ii) the buyer has information about the (relative) number of high-quality items across groups. For instance, the set of all used cars could be partitioned in a way that each group consists of all used cars whose owners live in a particular region or city. A seller partition may arise exogenously but can also be (further) configured endogenously by the buyer. For large (continuous) populations, our results apply for any random partition of the sellers, provided that the ex ante type distribution for the entire population of sellers is known.

The existence of a partition of sellers allows the buyer to employ a (non-direct) mechanism—which we call a *Devil’s Menu*—consisting of several price offers and allocation rules with which she can separate high-quality sellers from low-quality sellers. The payment each seller is promised depends on his own behavior but is also contingent on the aggregate behavior of all the partition groups. This creates downward competition between low-quality sellers across the different partition groups which induces all the sellers to report their type truthfully. The buyer then purchases all (or some) high-quality goods from a predetermined number of groups,

each at a price equal to the willingness to sell of the high-quality sellers. Some low-quality items also have to be bought by the buyer, yet at a lower, or even negligible, price—this is further discussed below. The extent of the surplus that the buyer can extract from the sellers then depends on the number of free parameters used in the Devil’s Menu (more parameters, larger surplus extraction) and on the extent to which items can be bought sequentially (more sequential voting, larger surplus extraction). In the absence of a seller partition, a Devil’s Menu becomes largely ineffective.

A Devil’s Menu: Mechanism description

A Devil’s Menu is a four-stage procedure. In Stage 1, the buyer offers the sellers the possibility to send one of two messages, called h (*high*) and l (*low*). These messages can be used by the sellers to reveal their quality type. As a key feature, the price that each seller is offered in exchange for his item depends on the message sent and on the (final) *status* of the group to which he belongs. The latter is determined in Stage 3. Accordingly, each message leads to one of two possible prices and thus the procedure involves four prices in total.

In Stage 2, sellers send one message at most—they can also choose *not* to participate in the mechanism. We stress that sellers do not know which price will ultimately prevail when they send the message.

In Stage 3, the status is determined for all groups. A group has the status *selected* if the ratio of the number of h messages to the number of high-quality sellers is among the set of the lowest q ratios for all groups, with ties being broken by fair randomization. Parameter q must be strictly lower than the number of groups upon which the pool of sellers is divided. Then, the prices offered to sellers for each message sent are set in accordance with the following properties. The price offered to sellers who sent h and whose group is selected is equal, or slightly above, the valuation and thus willingness to sell of the high-quality buyers. The price offered to sellers who sent h and whose group is not selected is also positive, but (arbitrarily) small. The prices offered to sellers who sent l are low, no matter whether their group is selected or not, but are higher than the price offered to sellers who sent h and whose group is not selected.

In Stage 4, *all* sellers can decide whether they want to sell their item at the price determined in Stage 3. Because sellers can walk away after sending their message, a Devil’s Menu is not a direct mechanism. It also means that our mechanism does not require any commitment power on the part of the sellers. For her part, the buyer must be committed at least to the allocation rule and to buying the items at the prices announced—more on this below.

Assuming that no weakly dominated strategy is played, the high-quality sellers’ behavior under a Devil’s Menu is very simple: they send message h and agree to sell their goods only if the price eventually offered to them is at least as high as their valuation. By contrast, a Devil’s Menu puts low-quality sellers into the following dilemma: Sending message h is more attractive, as long as their group is selected. However, by sending such a message, these sellers increase

their group's ratio of h messages to high-quality sellers, thereby raising the risk that the group will not be selected. For non-selected groups, by contrast, prices are very low regardless of the message sent, albeit marginally higher if message l was sent than if message h was sent. This means, in particular, that low-quality sellers should send message l if they anticipate that their group will not be selected with high probability.

Main (technical) results

We show that by setting the four prices of a Devil's Menu appropriately, the buyer induces a unique equilibrium of the underlying (Bayesian) game, in which *all* sellers report their type truthfully. That is, sellers with a high-quality item send message h , while sellers with a low-quality item send message l . Moreover, sellers sell their item to the buyer at the prevailing price, with the exception of high-quality sellers whose group has not been selected. In the unique equilibrium, all groups therefore have the same ratio of h messages to high-quality items, namely one. This implies that there is the same positive probability that any group is selected. Then, the four-price scheme ensures that the price that low-quality sellers expect to be offered in exchange for their item if they send message l is higher than the certain payoff associated with sending message h in a non-selected group. This first result constitutes the *weak form* of a Devil's Menu. It guarantees that the price low-quality sellers receive in exchange for their item is significantly lower than the valuation of high-quality sellers, although it is not negligible. The mechanism can be adjusted easily so that the buyer only buys a subset of the items whose sellers sent message h ; the rest of such sellers are offered zero.

We prove three further results about a Devil's Menu. First, running this mechanism sequentially and targeting all, or some, high-quality items of only one group each time reduces the budget the buyer spends, since it lowers the price eventually paid for low-quality items. Second, this latter price can, in fact, be set arbitrarily low, thereby eliminating almost all superfluous expenditures. The buyer's budget can then be used almost entirely for high-quality items. This is the *strong form* of a Devil's Menu. There are two variants of a Devil's Menu that ensure this: on the one hand, the price associated with message h in a non-selected group can be made arbitrarily low—up to the point where only three different prices are actually offered—although this entails the possibility that additional equilibria may exist. On the other hand, a more sophisticated six-price Devil's Menu may be used, which again ensures uniqueness of the equilibrium targeted. This latter variant requires that the buyer offers two different prices to sellers who sent message l and belong to selected groups, based on the groups' interim status. Finally, although in the description of all variants of a Devil's Menu we impose that the buyer buys all low-quality items, this is not necessary. With six prices, the suggested mechanism can be easily modified so that the buyer is committed to buying the commodity only from those sellers who sent message l and belong to a selected group with some exogenously given positive probability, which can be arbitrarily low (and, in particular, need not be equal to one). This enables the buyer to acquire much fewer low-quality items.

Contribution

The variant of the classical Lemons Problem that we explore comprises multiple sellers and one (big) buyer with a budget for high-quality items and is augmented by one crucial feature: seller partition.¹ Our main insight is that this feature resolves the Lemons Problem for the buyer because it enables her to use a Devil’s Menu. This is an important result that advances knowledge about markets in general. Yet its relevance extends to other domains. For instance, our results indicate a number of ways in which a collective decision taken by some (political) body through voting can be manipulated by a third party even if the latter’s pocket is not very deep.

Next, suppose momentarily that the buyer can commit to a single execution of the mechanism, that there is certainty about the total number of high-quality sellers in the population, and that there are no sellers who want to sabotage the market. Under these three circumstances, there is a much simpler mechanism than a Devil’s Menu that allows the buyer to acquire the desired number of high-quality items at a price equal to the sellers’ willingness to sell without incurring any additional expenses, even in the absence of seller partition. However, this is a knife-edge result for this simple mechanism. By contrast, the above three conditions are not needed for the buyer’s full (or partial) surplus extraction from the sellers via a Devil’s Menu. This allows us to conclude that the latter is a robust mechanism. In our baseline model, we proceed with strong commitment power, no uncertainty about the type distribution, and absence of saboteurs. Doing so makes our results more transparent and helps us to later assess how a Devil’s Menu can deal with the complications when we relax such conditions.

Applications

We offer several applications of our results. First, while the buyer may use the acquired items for herself as consumption goods, the fact that a Devil’s Menu allows the buyer to elicit the quality of the items could alternatively lead her to adopt a market maker role and resell the items of known quality—see Section 7.1. If the (benevolent) buyer could credibly announce the quality of goods—say, because she has the means to make the original h and l messages available to participants in the resale market—, she could then eliminate the Lemons Problem for any potential (small) buyer. Of course, a malevolent buyer could also abuse her market power for activities that are undesirable from a social perspective. In Section 7.2.2, we discuss a noteworthy example of this scenario in a voting context.

Second, a Devil’s Menu can be applied when there is another agent beyond the buyer and the sellers—say, a central authority—that can verify the quality of an item upon request and is committed to buying any item from any seller who wants to sell at a given price. This pins down the valuation of all sellers and makes such reservation prices equal across types. The

¹Since the buyer wants to buy a (possibly very large) number of items even if she exhausts her budget, we can refer to her as a *big buyer*.

scenario analyzed in Section 7.2.2 is also a good example of this application of our results.

Third, the market setup we analyze, with one buyer and multiple sellers, is equivalent from a formal perspective to a setup in which a political body made up of several members or voters must decide through voting which of two alternatives to implement, and in which there is an external *adversary* who wants to buy some votes. By doing so, the adversary—who can be a lobbying group or a special interest group—seeks to influence the political body’s decision. For the application of our results, it suffices to identify the sellers with the voters and the buyer with the adversary (see Section 7.2.1). Our main insight is that by using a Devil’s Menu, the adversary can capture the political body—say, a committee or an entire polity—relatively cheaply without overspending due to the fact that individual voter preferences are private.

Fourth and last, a Devil’s Menu opens up new possibilities for the use of large, albeit anonymized, datasets about sellers of particular items. Typically, such datasets can be obtained from service providers. These providers are not allowed by law to sell information about any particular individual seller type, but as a general rule, they can sell aggregated data about groups of potential sellers, e.g., the share of high-quality items. The spread of mechanisms based on a Devil’s Menu could foster the further development of such datasets.

Organization of the paper

The remainder of the paper is organized as follows. In Section 2 we review the literature most closely related to our paper. In Section 3 we present our baseline model. In Section 4 we introduce a Devil’s Menu with four prices and show its weak form. In Section 5 we analyze alternative forms of a Devil’s Menu, including two variants of its strong form. In Section 6 we show the robustness of a Devil’s Menu when the number of high-quality items in each group is not known, the buyer cannot commit to a single execution of a mechanism, and there are market saboteurs. In Section 7 we discuss various applications of a Devil’s Menu. Section 8 concludes. The proofs are in the Appendix.

2 Connection to the Literature

The two main approaches explored in the literature to alleviate the Lemons Problem are buyer screening and seller signaling, depending on who moves (first). In both areas, the literature is very extensive; we refer to Riley (2001) for an early survey on the topic. As to classic papers, Stiglitz (1975) studied screening under different institutions, while Spence (1973) showed how high-quality sellers can distinguish themselves from low-quality sellers by using carefully designed costly signals to signal their own quality. Related to our paper, Kim (2012) (see also Fang, 2001; Mailath et al., 2000) showed that low-quality sellers may have incentives to separate themselves endogenously from high-quality sellers in order to become attractive to

multiple buyers, by eliminating these buyer’s uncertainty.² This partially resolves the asymmetric information problem and allows the realization of trade gains, while offering a rationale for the co-existence of multiple market places.³ We contribute to this literature by showing how general partitions of the pool of sellers based on characteristics other than quality can help to solve the lemons problem for a *single* buyer, even if she does not have deep pockets and regardless of how low the proportion of high-quality items is in the population. It suffices for her to use a Devil’s Menu.

The latter insights are valid even if there are saboteurs and/or the buyers’ commitment power is limited and her information about the number of high-type sellers is scant. If *none* of these assumptions holds, a much simpler mechanism achieves the same goal as the Devil’s Menu without using the seller partition. This simple mechanism is reminiscent of classical one-principal-multiple-agents mechanisms such as those considered by Crémer and McLean (1988) in the context of auctions and by Piketty (1993) in the context of taxation. Complementary to these papers, we show that the principal (the buyer) can extract all the surplus from the agents (the sellers) by having the individual payoffs to the latter depend on the actions chosen by all agents in a particular way. To be able to do so, the principal must have some information about the agents’ type distribution. We differ insofar that our results do not rely on the assumption that the principal must know the exact number of agents (sellers) of each type as in Piketty (1993), nor is it necessary for the principal to know the degree to which the (interim) agents’ beliefs about the type distributions are correlated as in Crémer and McLean (1988).⁴ Provided the buyer knows the number of high-type agents in each group relative to the others, what drives surplus extraction by the buyer in our model is the existence of different groups of sellers. These groups are needed to induce downward competition between them via a Devil’s Menu. We also differ from the above papers in that, by considering several variants of a Devil’s Menu, we can relate the level of surplus extraction to the number of free parameters of the menu price and to the extent to which items can be bought sequentially.

It is often the case that beyond statistical information about item quality in the market, buyers and/or sellers might have access to other type of information, such as seller partition, post-purchase repair rate (Peterson and Schneider, 2017), the sellers’ past trading history (Kim, 2017), quality certification (Elfenbein et al., 2015), or the existence of middlemen (Biglaiser and Friedman, 1994). Because our mechanism can be used by a market maker to reopen a market, we also contribute to the particular strand of literature that examines how to restore efficiency in markets for lemons (see e.g Daley and Green, 2012, for a dynamic setup with news arrival). A host of different public policy interventions, including government purchases, are

²Mailath et al. (2000) also considers a seller (worker) characteristic that is not payoff-relevant, but their dynamic setup focuses on two-sided search and is thus very different from ours.

³Adverse selection has also been studied in social networks (see e.g Lin et al., 2013).

⁴From a technical viewpoint, we also differ from Crémer and McLean (1988) and Piketty (1993) insofar that we deal with uncertainty about the ex post type distribution by focusing neither on dominant strategy equilibria nor on Bayesian equilibria, but rather on ex post Nash equilibria.

discussed in Moreno and Wooders (2016) for general lemons markets. Restricting opportunities for trade after an initial round of trade is shown to be optimal in Fuchs and Skrzypacz (2015). There is an extensive literature on lobbying and vote-buying, and in general about the problem how a collective decision taken through voting can be affected by monetary payments (bribes) made by a third party (see e.g. Buzard and Saiegh, 2016; Dekel et al., 2008; Felgenhauer and Grüner, 2008; Finan and Schechter, 2012; Groseclose and Snyder, 1996; Le Breton and Salanié, 2003; Louis-Sidois and Musolf, 2020; Schneider, 2014). The paper that is closest to ours is Dal Bo (2007). Both his paper and ours show that political bodies may be vulnerable to external influence and that acquiring this influence can be done very cheaply by such a third party (also called *adversary*). The key feature is that through payment promises that are contingent on individual and aggregate voter behavior the adversary can create competition among the voters that allows her to capture the political body and dictate the outcome. The difference with Dal Bo (2007) is that our mechanism is more robust—against lack of commitment, lack of information, and presence of saboteurs—and that the main source of inefficiency for the adversary is that she needs to tell apart the voters who share her same preferences from those who do not.

In the paper we also discuss how new technologies such as blockchain and smart contracts offer a way of implementing our mechanism in real-world environments. We are not the only ones to suggest the possibility that technological change can provide new means to alleviate the asymmetric information problem in a lemons market (see e.g. Aoyagi and Adachi, 2018; Cong and He, 2019; O’Dair and Owen, 2019; Zavolokina et al., 2019). We contribute to this strand of literature by providing and analyzing a new buying mechanism.

Finally, our analysis can be interpreted as an exercise in mechanism design without commitment to a single application of our mechanism. This theme has been explored in other contexts by Skreta (2006, 2015), who studied mechanism design problems where the designer cannot commit to *not* repeating the mechanism from a dynamic perspective. In the absence of full commitment power, it is well known that the revelation principle cannot be generically applied (see e.g. Freixas et al., 1985; Hart and Tirole, 1988) and mechanisms that are not direct—such as a Devil’s Menu—may have to be considered.⁵ From the perspective of the mechanism-design literature that focuses on the role of commitment, our contribution is to show that lack of buyer commitment power can be compensated with seller partition to resolve the adverse selection problem faced by the buyer. From the general perspective of mechanism design, we also add to the well-know results on VCG mechanisms (Clarke, 1971; Groves, 1973; Vickrey, 1961) by introducing a (non-direct) mechanism that induces all sellers to reveal their type truthfully, and minimizes the transfers needed to attain the optimal solution for the designer (*viz.* the

⁵We refer to Bó and Hakimov (2020) for a recent paper on sequential revelation mechanisms.

buyer).⁶ To be effective, the mechanism we propose requires the buyer to have additional, non-payoff-related information about the sellers (at least for finite populations of sellers).

3 The Baseline Model

3.1 A problem with many sellers and one buyer

There is a finite set of risk-neutral sellers denoted by N . Each seller $i \in N$ owns *one* item and may be of one of two types: if $t_i = H$, he has a *high-quality item*; if $t_i = L$, he has a *low-quality item*.⁷ The former sellers are henceforth called *high-quality sellers*, the latter *low-quality sellers*. Types are privately known. For simplicity, we assume that the sellers' valuation for a high-quality item is $V > 0$, with the seller's valuation for a low-quality item being normalized to 0. This assumption suffices to generate an adverse selection problem.⁸ The utility of seller $i \in N$ when selling an item of quality $v_i \in \{0, V\}$ and receiving a monetary transfer t_i is⁹

$$u_i(v_i, t_i) = -v_i + t_i.$$

In Section 6.3 we consider the possibility that one or a few of the sellers owning a low-type commodity want to sabotage the market. Considering this does not affect the results we obtain for our baseline setup.

There is also a *buyer*, henceforth denoted by B , with a budget $b > 0$. Her (quasi-linear) utility depends on the number of high-quality items she has acquired, m_h , and the total transfers t spent on acquiring items—either of high or low quality—as follows:

$$u_B(m_h, t) = \alpha \cdot \min\{m_h, m\} - t. \tag{1}$$

In equation (1), m is a fixed positive number that is lower than the total number of high-quality sellers and $\alpha > 0$ denotes the importance of high-quality items relative to money. We assume that b and α are sufficiently large so that we can think of B 's preferences as being lexicographic first in the number of high-quality items bought (up to m), and second in the budget that is left after buying these (and other) items. Such an assumption is not needed for the definition of our mechanism, and the mechanism would still work under less stringent assumptions.¹⁰ Hence,

⁶For mechanism design approaches that could be used for solving special cases of the Lemons Problem, see Börgers (2015).

⁷As is common in the literature, we assume that item quality is exogenously given. We refer to Kawai (2014) for a model in which quality can be endogenously determined by the seller.

⁸Valuations are obtained from the utility sellers derive from consuming or using the good. They represent the amount of money sellers would need to receive in exchange for the good so that their utility remains the same.

⁹It suffices that $u_i(v_i, t_i)$ is decreasing in v_i and increasing in t_i .

¹⁰The analysis becomes, however, much more cumbersome if the buyer varies the number of items she wants to buy with small variations of the expected expenditures.

a necessary condition for B 's utility to be maximized is that she buys m high-quality items, provided that she does not violate her budget constraint. In the standard literature on adverse selection typically assumed that the (risk-neutral) buyer(s) has a maximum willingness-to-pay for each quality type and that, conditional on the prior information available, she is not willing to trade, despite the fact that she would do so if she knew the true quality. Instead, we consider a buyer with large willingness to pay for m high-quality items, who is only constrained by her budget.¹¹

Of course, for any item it must be that the buyer's valuation is at least as much as the valuation of a seller owning such an item. This occurs when the buyer's consumption level is higher than that of the sellers or when as a result of the buyer acquiring sufficiently many (good-quality) items, the buyer can change the outside market structure that determines the valuations of all participants in the (new) market. For instance, a big buyer can have access to further markets or affect the ratio of high-quality items to good-quality items in the population—say, by not dumping some low-quality items into the market—in a way that all (remaining) goods are traded regardless of quality. This is further discussed in Section 7.1.

3.2 The buyer's problem: first-best and second-best solutions

To attain the above goal of buying m high-quality items, buyer B must devise a (potentially stochastic) mechanism, denoted by M , to buy commodities from the sellers at some stipulated prices. Recall that for the execution(s) of the mechanism, buyer B has budget b , whose lower bound is given below—see (6). The budget b may represent the buyer's entire budget or just the part she has committed to for the execution of a given mechanism.¹² The optimality problem faced by B can be therefore formulated as follows:

$$\begin{aligned} \max_M \mathbb{E}[u_B(m_h, t)|M] \\ \text{s.t. } t \leq b. \end{aligned} \tag{2}$$

In our setup, finding a mechanism M that yields a maximizer of (2) is not feasible. This is because seller participation cannot be enforced and because high-quality sellers will never sell their commodities at a price lower than V . Thus, the (*constrained*) *first-best solution* from the

¹¹One could include more buyers in the model. Yet, as long as one of the buyers is much bigger (i.e. has a much larger budget) than the other buyers combined, our results would not change substantially and only some small inefficiencies would arise. Our focus on market makers—see e.g. Section 7.1—provides a rationale for the assumption that there is one (big) buyer, which in turns makes the analysis more transparent and less cumbersome. We refer to Muthoo and Mutuswami (2011) for a paper on the market for lemons with one buyer and several sellers. Other papers in the literature investigate the other polar case in which there is only one seller (see e.g. Daley and Green, 2012). The case of one buyer is a reasonable assumption if our setup is used to model political influence—see Section 7.2.1.

¹²For instance, committing a certain amount of money for the execution of a mechanism can be easily done through smart contracts (see Section 7.3).

perspective of the buyer is characterized by $m_h = m$ and $t = m \cdot V$, which we write as

$$z^* := (m, m \cdot V). \quad (3)$$

In the first-best solution, the buyer acquires m high-quality (and possibly some low-quality) commodities at the willingness to sell by the sellers holding such items or, equivalently, at the competitive price in the market for both types of commodities. While z^* can be attained in some variants of a Devil's Menu, we also analyze different variants of our mechanism that yield

$$z^{**} := (m, x),$$

with $m \cdot V < t \leq b$. Any solution z^{**} is called *second-best*. For simplicity, we initiate our analysis with the assumption that the buyer can commit to a single execution of a mechanism. Later, in Section 6.2, we show that such an assumption is not needed for our results.

3.3 Seller partition

A key ingredient of our model is that the set N can be partitioned into $\bar{k} > 1$ groups $N_1, \dots, N_{\bar{k}}$, with $n_k^H > 0$ and $n_k^L > 0$ denoting the number of sellers of high type and low type within N_k , respectively, for all $k \in \{1, \dots, \bar{k}\}$. For now, we assume that $(n_k^H, n_k^L)_{k=1}^{\bar{k}}$ is known to the buyer. In Section 6.1 we show that this assumption can be relaxed without affecting our results.

For any set S , it is convenient to have s denote its cardinality. Then, we let $n_k := n_k^H + n_k^L$, for all $k \in \{1, \dots, \bar{k}\}$, and $n := n_1 + \dots + n_{\bar{k}}$ denote the total number of sellers in each of the groups and in the entire population, respectively. Additionally, we let $n^L := n_1^L + \dots + n_{\bar{k}}^L$ and $n^H := n - n^L$ denote the total number of high-quality and low-quality sellers in the whole population, respectively.

3.4 Recasting the buyer's goal

Finally, assume momentarily that

$$n_k^H = n^H / \bar{k} \text{ for all } k \in \{1, \dots, \bar{k}\}, \quad (4)$$

and

$$m = q \cdot (n^H / \bar{k}), \quad (5)$$

for some $q < \bar{k}$. That is, all groups have the same number of high-quality items, and the buyer is interested in buying all high-quality items from exactly q subgroups. Unless stated otherwise, we proceed on the simplifying assumption that B 's goal is to devise a mechanism that allows her to buy high-quality commodities of exactly q groups, at a minimal cost. Under (4) and (5),

this follows naturally from (1) and (2). When the number of high-quality items varies across groups, such an assumption makes exposition simpler but is not crucial for our results. The mechanism(s) we present can easily be adjusted to the general case in which the seller wants to buy a lower number of items from q groups. It suffices for buyer B to add a final stage to the mechanism(s) we consider, in which she will effectively buy the items from the selected groups at the prevailing prices only for certain sellers, who are chosen by randomization. Adding this stage does not change the equilibrium behavior of sellers, no matter their type.

4 A Devil’s Menu

In this section, we introduce and analyze a mechanism—called a *Devil’s Menu*—that relies on the partition of the set of sellers in different groups. Before we start, note that to buy the high-quality items from any q groups, a first attempt based on brute force approach would be for buyer B to try and buy all commodities at a price $V + \varepsilon$, for some $\varepsilon > 0$ arbitrarily small. Doing so is feasible if B ’s budget, denoted by b , is large enough. It suffices to assume that

$$b \geq \max_{S \subseteq \{1, \dots, k\}, |S|=q} \left[(V + \varepsilon) \cdot \sum_{k \in S} n_k \right]. \quad (6)$$

The right-hand side in the above inequality is the amount the buyer would have to spend if she bought *all* items from an arbitrary selection of q groups. To do this, she would offer to pay $V + \varepsilon$ to all sellers in these groups in exchange for their items. Because the price offered to the sellers would be higher than their item valuation, regardless of the type of item they hold, they would *all* accept the transaction. Note that in equation (6) we use the maximum over all possible subsets of groups, since the mechanisms we introduce do not allow buyer B to choose q specific groups.

It is clear that the brute-force attack is very inefficient, since the buyer has to buy all items, of high quality and of low quality, at the same high price, namely $V + \varepsilon$. In particular, a second-best solution z^{**} is attainable through this brute-force attack only if budget b is very large. In what follows, we show that buyer B can achieve her goal more efficiently—i.e., better second-best solutions z^{**} , or even the first-best solution z^* , can be attained—if she uses a Devil’s Menu.

We start by describing the mechanism that guarantees that the buyer can buy all high-quality items in q groups, each at a price $V + \varepsilon$, without having to pay this same amount for all low-quality items in these groups. As a tie-breaking rule, we assume henceforth that the individual seller does not offer his item to the buyer if his valuation of the item is equal to the price offered by the buyer. The mechanism works as follows:

1. Buyer B offers all sellers the possibility to send two messages (if at all): message h or

message l . Depending on the ratio of messages h to high-quality sellers across groups, buyer B selects q groups, as formulated in Step 3 below. Moreover, B offers the following prices for items:

	message h	message l
group selected	$p_1^{se} = V + \varepsilon$	$p_2^{se} = \theta$
group not selected	$p_1^{ns} = \varepsilon$	$p_2^{ns} = 2\varepsilon$

Table 1: A Devil’s Menu with four prices—depending on the message sent (column) and the final status of the seller’s group (row).

In this four-price procedure, $\varepsilon > 0$ is assumed to be arbitrarily small, while θ satisfies $2\varepsilon < \theta \leq V - \varepsilon$.¹³

- Each seller sends one of the two messages or none. For each $k \in \{1, \dots, \bar{k}\}$, we let m_k be the number of sellers from N_k who sent message h , and we set $\rho_k := m_k/n_k^H$.¹⁴
- Assume w.l.o.g that $\rho_1 \leq \rho_2 \leq \dots \leq \rho_{\bar{k}}$ and let $\rho := \rho_q$. Then define the following three sets of groups:

$$C = \{N_k | k \in \{1, \dots, \bar{k}\}, \rho_k < \rho\},$$

$$T = \{N_k | k \in \{1, \dots, \bar{k}\}, \rho_k = \rho\},$$

$$O = \{N_k | k \in \{1, \dots, \bar{k}\}, \rho_k > \rho\}.$$

C, T, O are abbreviations for *chosen*, *tied*, and *omitted*, respectively. Accordingly, we have $\bar{k} = c + t + o$ and $t \geq \max\{1, q - c\}$. All groups in C are *selected* by buyer B . From the set T , $q - c$ groups are chosen by fair randomization, so that their *final status* is also selected. That is, a group in T has a probability $\frac{q-c}{t}$ of being selected by the mechanism—and then belonging to T is its *interim status*. The final status of the remaining groups is *non-selected*.

- Once all groups have been assigned their final status—and hence the price offered to each applicant has been determined—, these sellers decide whether to sell their item at the prevailing price or not. As to buyer B , she is obliged to buy.

A *Devil’s Menu* is a procedure involving four prices in which sellers sequentially commit, first, to a subset of prices for the particular message sent, and, second, to the possibility of selling

¹³The reason why we require $\theta \leq V - \varepsilon$ and not only $\theta < V$ becomes clear in Section 5.3.

¹⁴Our results also hold if we define $\rho := m_k - n_k^H$.

their item once the final price has been determined for the message type they sent. These final prices depend on the choices made by *all* the sellers. In turn, the buyer commits to the correct execution of the mechanism—in Section 7.3, we discuss ways how the buyer can commit to the protocol underlying the mechanism. It is useful to display the (expected) payoff matrices of sellers. Since a low-quality seller has zero reservation price and hence he always sells his item in Step 4, his expected payoff is given in Table 2.

	C	T	O
h	$V + \varepsilon$	$\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon$	ε
l	θ	$\frac{q-c}{t} \cdot \theta + \frac{t-q+c}{t} \cdot 2\varepsilon$	2ε

Table 2: Expected payoff of a low-quality seller in a Devil’s Menu with four prices—depending on the message sent (row) and the final status of the seller’s group (column).

High-quality sellers do not sell their item if their group is not chosen, nor they sell their item if they have sent message l . In either case, they are offered a price below their reservation price. By not selling, they end up with a payoff of V . Their expected payoff is given in Table 3.

	C	T	O
h	$V + \varepsilon$	$\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot V$	V
l	V	V	V

Table 3: The expected payoff of a high-quality seller in a Devil’s Menu with four prices—depending on the message sent (row) and the final status of the seller’s group (column).

4.1 The weak form of a Devil’s Menu

With the four-price Devil’s Menu in place, and taking the optimal choices of Step 4 discussed above into account, we can define a (simultaneous-move) Bayesian game, which we denote by \mathcal{G}_4 . The player set is N , and for each seller i , the strategy set is $S_i = \{l, h, \emptyset\}$ and the type set is $T_i = \{L, H\}$. The payoff matrices are given by Tables 2 and 3, with the selection of groups determined in accordance with Step 3 of the Devil’s Menu. The following observations follow immediately: First, for each seller, to abstain and not to send a message is weakly dominated by sending message h (for high-quality sellers) and sending message l (for low-quality sellers). We assume that no player plays a weakly dominated strategy, so that all sellers accept to

participate in a Devil’s Menu and send a message.¹⁵ This implies that, effectively, we can consider $S_i = \{l, h\}$ for each seller i . Second, for a high-quality seller, sending message h weakly dominates sending message l . After these observations, we obtain our first main result. For the remainder of the paper and unless specified otherwise, a Bayesian Nash equilibrium in which weakly dominated strategies are eliminated is simply called an *equilibrium*.

Theorem 1. [Weak form of a Devil’s Menu] Suppose $\theta \geq \frac{q}{k} \cdot V + 2\varepsilon$. Then the strategy profile σ^* in which all low-quality sellers send message l and all high-quality sellers send message h is the only equilibrium of game \mathcal{G}_4 .

Proof. See Appendix. □

The rationale for the functioning of a Devil’s Menu with four prices as described in Theorem 1 is clear. By offering different prices—whose realizations depend on the sellers’ behavior, both at group level and aggregate level—, low-quality sellers are trapped. These sellers would like to send message h and simultaneously ensure that their group is selected by the buyer. Sending a message h , however, reduces their group’s chances to be selected. As a consequence, a Devil’s Menu generates downward pressure on $\rho_1, \dots, \rho_{\bar{k}}$, which results in only high-quality sellers sending message h in all groups. In particular, only high-quality sellers of the selected groups are paid the high price $V + \varepsilon$, with low-quality sellers in these same groups obtaining θ instead. In turn, low-quality sellers in non-selected groups obtain 2ε . High-quality sellers in non-selected groups do not sell their items. To sum up, the budget that the buyer pays in equilibrium is never higher than the following bound:

$$\max_{S \subseteq \{1, \dots, \bar{k}\}, s=q} \left[(V + \varepsilon) \cdot \sum_{k \in S} n_k^H + \theta \cdot \sum_{k \in S} n_k^L + 2\varepsilon \cdot \sum_{k \notin S} n_k^L \right] =: \bar{b}. \quad (7)$$

The buyer chooses the lowest possible value of θ to minimize her costs, so she sets θ equal to $\frac{q}{k}V + 2\varepsilon$.¹⁶ Compared to a buying strategy based on brute force, the buyer can therefore extract from each low-quality seller a surplus that is (approximately) equal to

$$V - \frac{q}{k} \cdot V = \frac{\bar{k} - q}{\bar{k}} \cdot V,$$

if we set $\theta = (q/\bar{k}) \cdot V$ and $\varepsilon \approx 0$. That is, the higher the ratio $\frac{\bar{k} - q}{\bar{k}}$, the higher surplus extraction. In particular, if the latter ratio is close to one, surplus extraction is close to, but falls short of, being full. To sum up, a second-best solution z^{**} is obtained, which yields a higher utility for buyer B than the buying strategy based on brute force.

¹⁵Looking for Bayesian Nash equilibria in pure, weakly undominated strategies of this simultaneous-move game is equivalent to looking for perfect Bayesian equilibria in pure strategies of the dynamic game defined by the four-step mechanism described in Section 4, in which no seller plays a weakly dominated strategy.

¹⁶If $q = \bar{k}$, the result of Theorem 1 is still valid. Nevertheless, it does not yield any budget savings for the buyer compared to a brute-force attack in which *all* sellers are paid V .

In our analysis we have proceeded with the assumption that the buyer wants to buy all high-quality items of a certain number of groups and has a budget that allows her to do so at a unit price that equals the willingness to sell of the sellers holding such items. As mentioned in Section 3, alternative objectives for the buyer could also be achieved by slight modifications of a Devil’s Menu with four prices. One alternative is to consider that the buyer might want to buy only a certain share of the high-quality items from any of the q groups she is targeting. A Devil’s Menu with four prices can still be applied in such cases, with the following variation: As a last step of the mechanism, some of the sellers who sent message h in each group are selected randomly and are offered the price prescribed by Table 1 in exchange for their items. The remaining sellers who sent message h are offered zero instead. Under the same assumptions as in Theorem 1, such a mechanism ensures that the buyer can achieve her new objective.¹⁷ The reason is that this variation does not alter the downward pressure on the ratio of messages h sent to sellers holding high-quality items in each group, and again, only high-quality sellers send message h . Among them, however, only those will sell their items, who belong to selected groups and are eventually offered the positive price—as prescribed by Table 1.¹⁸ The same considerations hold for the remaining forms of a Devil’s Menu, which are introduced in Section 5.

4.2 Critical and non-critical assumptions

Next, we discuss some further critical and less critical assumptions of a Devil’s Menu, which also apply to the mechanisms that are introduced subsequently. We defer to Section 6 the analysis of the role of uncertainty about the number of high-quality items in each partition group, of limited buyer commitment, and of the presence of one or a few sellers who want to sabotage the market.

The first result in the present section is that the bound \bar{b} is not only sufficient to buy the high-quality items from q groups on the equilibrium path, it would also be sufficient to match expenditures if a single low-quality seller deviated from the unique equilibrium. Theorem 1 and the way a Devil’s Menu is constructed imply the next result.

Corollary 1. *The budget bound \bar{b} is equal to, or higher than, the expenditures faced by the buyer when one low-quality seller deviates and sends message h .*

This corollary follows directly from the observation that a deviation by one low-quality seller turns his group into a non-selected group with probability one. Accordingly, this seller cannot

¹⁷From a mathematical perspective, this is equivalent to considering the valuation for high-quality items to be $p \cdot V$, where p is the probability that a seller who sent message h will eventually be offered a positive price and not just zero.

¹⁸Remarkably, the correct execution of this modified mechanism requires a budget for on- and off-equilibrium behavior that is lower than the one specified in equation (7), yet with another objective.

generate a budget problem for the buyer.¹⁹ This means that to handle off-equilibrium threats by one single seller, the buyer does not need to have deep pockets. We further note that to guarantee that the mechanism can be run—and that the equilibrium described in Theorem 1 is unique—, it suffices for the buyer’s budget b to satisfy inequality (6) and for this to be common knowledge.

Another assumption of our model is that sellers make their applications simultaneously. However, this is not crucial for our results. Indeed, at any point in time where some sellers have already sent their messages as prescribed by Theorem 1, the best response of any remaining seller is to send the message in the same way. That is, the unique equilibrium outcome of Theorem 1 also arises as the unique equilibrium of any dynamic version of game \mathcal{G}_4 . Whether the game is dynamic or not, computing the optimal strategy is an easy task for any seller, so that the requirements to behave rationally are neither very demanding nor unrealistic. This adds to the plausibility of a Devil’s Menu as a way to separate low-quality sellers from high-quality sellers in real-world applications.

It is also important to mention that the mechanism and our previous analysis can both be accommodated to the case of a continuum of sellers, as long as the shares of high-quality sellers in all groups is common knowledge. Following the same logic contained in our proofs—see the Appendix—, it can be shown that a set of sellers of non-zero measure will not deviate from sending a truthful message about their types. This means that our mechanism(s) can screen seller types if we have finitely many sellers and exact common information about how many of them are of high quality in each group, and also if we have a continuum of sellers and the ratio of high-quality sellers to low-quality sellers in each group is known for the distribution from which types are drawn. For a continuum of sellers, a suitable version of the law of large numbers ensures that the ex post ratio of high-quality items to low-quality items in any element of a finite random partition of the set of sellers coincides with the ex ante ratio for the entire population. Hence, if the latter ratio is known, so are all the former ratios. In such cases, it is sufficient that objects are characterized only by quality.

Finally, we have confined ourselves to a simple model with only two seller types (or only two seller valuations). The construction of the Devil’s Menu can nonetheless be extended using the logic behind the two-type case to screen more than two seller types. With T types, in particular, this would require either $2T$ prices.

¹⁹Deviations by a larger number of low-quality sellers would not cause a budget problem as long as the total number of groups where individual low-quality sellers deviate by sending message h is smaller than or equal to $\bar{k} - q$. This is, in particular, the case when low-quality sellers can only collude if they belong to the same group.

5 Buying Low-Quality Items Cheaper

In the preceding section, we have shown that the buyer can buy all high-quality items of q groups, but she cannot avoid paying the price θ to all low-quality sellers in those groups. In this section, we outline three ways to lower the price.

5.1 Sequential item buying

First, the buyer may proceed sequentially and buy only the high-quality items of one group at a particular point in time. Specifically, suppose that there are q points in time and the buyer commits to buying the high-quality items of only one group at each date.²⁰ We prove the following result:

Proposition 1. *Let $\theta \geq \frac{1}{k-q+1} \cdot V + 2\varepsilon$ and suppose that the buyer uses a Devil's Menu with four prices sequentially and buys the high-quality items of one group at each date. Then, in any perfect Bayesian equilibrium in which sellers choose actions that are not weakly dominated, all sellers act as prescribed by the strategy profile σ^* , according to which all low-quality sellers send message l and all high-quality sellers send message h .*

Proof. See Appendix. □

One can easily verify that

$$\frac{q}{k} > \frac{1}{k-q+1},$$

so buying objects from one group at a time reduces the expenditures on low-quality items compared to the case in which objects are bought from all groups at once. This is because the downward pressure put on groups decreases if we increase the number of groups the buyer must buy the goods from.

5.2 A strong form of a Devil's Menu: Four prices

There are at least two alternative versions of a Devil's Menu that require only negligible payments for low-quality items, and thus they yield the first-best solution z^* . Both alternatives—which are manifestations of the so-called *Strong Form of a Devil's Menu*—enable the buyer to use nearly her entire budget on high-quality items. Recall that b is merely the budget the buyer is committed to for the realization of the mechanism, and may only represent a small share of her entire budget.

²⁰Our analysis can be easily generalized to the case where the buyer commits to buying from more than one group, yet from fewer than q groups, at each date.

The first alternative follows the same mechanism as the Devil’s Menu outlined in Section 4, but the price setting is slightly different. Specifically, the price scheme is as follows:

	message h	message l
group selected	$p_1^{se} = V + \varepsilon$	$p_2^{se} = 2\varepsilon$
group not selected	$p_1^{ns} = \varepsilon$	$p_2^{ns} = 2\varepsilon$

Table 4: A strong form of a Devil’s Menu with four prices—depending on the message sent (column) and the final status of the seller’s group (row).

Note that by setting $\theta = 2\varepsilon$, Table 4 follows from Table 1 and that, moreover, only three different prices are actually offered. The corresponding simultaneous-move game is denoted by $\widehat{\mathcal{G}}_4$. We obtain the following result:

Proposition 2. *The strategy profile σ^* , in which all low-quality sellers send message l and all high-quality sellers send message h , is an equilibrium of game $\widehat{\mathcal{G}}_4$.*

Proof. See Appendix. □

Proposition 2 has important consequences. Since ε can be chosen arbitrarily small, the budget of the buyer can be used almost entirely for buying high-quality items in the equilibrium described in Proposition 2. Total expenditures are then bounded by

$$\max_{S \subseteq \{1, \dots, \bar{k}\}, s=q} (V + \varepsilon) \cdot \sum_{k \in S} n_k^H + 2\varepsilon \cdot \sum_{k \in \{1, \dots, \bar{k}\}} n_k^L.$$

One possible drawback of this strong form of a Devil’s Menu is that game $\widehat{\mathcal{G}}_4$ may have other equilibria in which some low-quality sellers also send message h . To eliminate these other equilibria, one has to introduce further refinements in the price offering. This constitutes the second part of the strong form of a Devil’s Menu.

5.3 A strong form of a Devil’s Menu: Six prices

An alternative way of reducing the cost of buying low-quality items almost entirely is to enlarge the menu of prices. The price scheme of the second strong form of a Devil’s Menu is as follows:

	message h	message l
group in set C (and selected)	$V + \varepsilon$	$V - \varepsilon$
group in set T and selected	$V + \varepsilon$	θ
group not selected	ε	2ε

Table 5: A strong form of the Devil’s Menu with six prices—depending on the message sent (column) and the interim and final status of the seller’s group (row).

In this six-price mechanism, $\varepsilon > 0$ is arbitrarily small. It also suffices to choose θ such that $\theta \geq 2\varepsilon$. In addition, note that in Table 5 only five different prices are actually offered. The corresponding simultaneous-move game is denoted by \mathcal{G}_6 . We obtain the following result:

Theorem 2. [A strong form of a Devil’s Menu] Suppose that $\theta \geq 2\varepsilon$. Then the strategy profile σ^* in which all low-quality sellers send message l and all high-quality sellers send message h is the only equilibrium of game \mathcal{G}_6 .

Proof. See Appendix. □

A Devil’s Menu with six prices eliminates almost all expenditures on low-quality items and again generates more buying power for high-quality items. Indeed, by taking $\theta = 2\varepsilon$, total expenditures in equilibrium are now bounded by

$$\max_{S \subseteq \{1, \dots, \bar{k}\}, s=q} (V + \varepsilon) \cdot \sum_{k \in S} n_k^H + 2\varepsilon \cdot \sum_{k \in \{1, \dots, \bar{k}\}} n_k^L.$$

In particular, the buyer extracts all surplus from low-quality sellers.

It is important to point out that this second strong form of a Devil’s Menu rests decisively on the assumption that the buyer can credibly offer two different prices for selected groups, at least for the sellers who send message l . The two prices discriminate among groups based on their interim status. On the one hand, the group may belong to C —and hence its interim status is immediately selected. On the other hand, the group may belong to T and be selected only after it has been chosen by fair randomization.²¹ This property can be built into the algorithms executed by smart contracts—see Section 7.3—, but it may be less accepted by sellers.

Finally, we note that a six-price Devil’s Menu exhibits one additional desirable property. Since $\varepsilon > 0$ can be made arbitrarily small without affecting the result, one could simply conceive of this parameter as originating in the following way: To those sellers who have sent message h

²¹It may happen that set T contains only one group, in which case this group is selected with certainty. Sellers of this group are thus offered lower prices than sellers of all other selected groups.

and belong to groups who have not been selected, buyer B is committed to paying a fixed price with some (arbitrarily low) positive probability. This does not change the sellers' behavior on and off the equilibrium path and results in the buyer buying much fewer low-quality items.

6 Robustness: A Simple Mechanism as Benchmark

In the baseline model analyzed in the previous sections we have assumed that (i) the buyer knows the number of high-quality items in each group, (ii) the buyer can commit to a single execution of any mechanism of her choice, and (iii) all sellers derive utility only if they are able to sell their items above their reservation price. For such a setup, we have shown that a Devil's Menu can yield the first-best outcome for the buyer by relying on the partition of the set of sellers to induce downward competition between groups that prompts truth-telling behavior among all sellers. Yet, under the above three assumptions, the buyer could buy all required high-quality items m at the minimum total cost without using the seller partition through a much simpler mechanism, called the *simple mechanism*, that specifies the following course of actions:

1. The buyer announces the possibility to send one of two messages, h and l .
2. Each seller sends one of the two possible messages, if at all.
3. With sellers being able to decide whether to sell their item or not at the prevailing price—which is determined in this same stage—and the buyer being obliged to accept the transaction at the request of the sellers, payoffs are realized according to the following rule:
 - If more than n^H sellers sent message h , then each of them is offered 0 in exchange for his item. In turn, all sellers who sent message l are offered ε , also in exchange for their items. Recall that n^H denotes the total number of high-quality sellers in the population.
 - If n^H sellers at most sent message h , then m of them are chosen randomly. They are then offered $V + \varepsilon$ in exchange for their item. Sellers who sent message h and are not selected are offered ε in exchange for their items. Finally, all sellers who sent message l are offered ε , also in exchange for their items. We assume that $\varepsilon > 0$ is arbitrarily small.

It is easy to see that for the Bayesian game induced by the simple mechanism, all high-quality sellers send message h and all low-quality sellers send message l , provided that sellers do not play weakly dominated strategies and such strategies are removed iteratively. Moreover, the

expenditures used to buy low-quality items are negligible. Hence the first-best solution z^* (approximately) obtains as an outcome of the mechanism.

Nevertheless, the following assumptions are needed for the above result about the simple mechanism: (i) the buyer has full commitment power, (ii) there is no uncertainty about the number of high-quality items in each group (and, hence, in the entire population), and (iii) there are no saboteurs among the sellers owning low-type items are necessary. By contrast, a Devil's Menu still yields the same desirable outcome for the buyer even if assumptions (i)–(iii) are relaxed because, unlike the simple mechanism, it uses the seller partition. For our robustness analysis, we focus on the four-price Devil's Menu of Theorem 1. Yet, a discussion of other variants of a Devil's Menu can be done along the same lines.

6.1 Uncertainty about the type distribution

Consider now that for each group N_k , with $k \in \{1, \dots, \bar{k}\}$, the (positive) number n_k^H indicating how many items of high-quality there are in group N_k is random. Furthermore, assume that for each pair $k, j \in \{1, \dots, \bar{k}\}$,

$$n_k^H - n_j^H = \mathbb{E} [n_k^H] - \mathbb{E} [n_j^H]. \quad (8)$$

That is, while the number of high-quality items is unknown in each group, the difference of these numbers across groups is common knowledge. Alternatively, for each $k \in \{1, \dots, \bar{k}\}$ we can envision that

$$n_k^H = m_k^H + X, \quad (9)$$

where m_k^H is some fixed positive integer that depends on k , and X is some random variable with support on the non-negative integers that is independent of k . One possibility is for random variable X to capture any exogenous shocks that might render some goods valueless. In our setup, conditions (8) and (9) are equivalent.

Imposing (8), and hence (9), means that the total number of high-quality items—and, hence, its proportion to low-quality items—is uncertain. When X has always value zero, the setup reduces to the one considered in the previous sections. As in these sections, we assume here that the goal of the buyer is to buy the high-quality items of any q groups.²²

Next, for each pair of groups N_k, N_j , define $k \succ j$ if and only if

$$m_k - m_j < \mathbb{E} [n_k^H] - \mathbb{E} [n_j^H]. \quad (10)$$

Let also $k \succeq j$ if and only if $\neg(j \succ k)$. One can easily see that \succeq is transitive and complete,

²²The number of high-quality items is now stochastic, so we assume that the buyer commits herself to buying an unknown number of goods. As mentioned at the end of Section 4.1, a Devil's Menu can always be modified so that the buyer is committed to buying only a fixed number of these goods.

so that it defines a total order of the set $\{1, \dots, \bar{k}\}$ or, equivalently, a total order of the set composed of the elements of the seller partition. Each variant of a Devil's Menu that has been defined in the previous sections can then be adjusted to this setup if groups are ordered according to \succeq instead of using parameters $\rho_1, \dots, \rho_{\bar{k}}$.

We obtain the following result:

Proposition 3. *Suppose that $\theta \geq \frac{q}{k} \cdot V + 2\varepsilon$ and that no seller plays a weakly dominated strategy. Then, in the unique ex-post equilibrium of game \mathcal{G}_4 all sellers report their type truthfully.*

Proof. See Appendix. □

That is, at least as far as ex post equilibria in which no seller plays a weakly dominated strategy are concerned, a Devil's Menu yields the same outcome when the buyer knows the *absolute* number of high-quality items in each group and when she only knows the *relative* number high-quality items in each group compared to the other groups of the seller partition.

What about the performance of the simple mechanism when the total number of high-quality items is unknown? To keep the comparison with Proposition 3, we focus on ex post equilibria in which sellers play no weakly dominated strategy. By construction, the simple mechanism must specify a threshold for the number of goods that the buyer will consider to be of high quality, of which she will buy some (in a random way). If the threshold is higher than the actual number of high-quality items in the population, the buyer will end up buying some low-quality goods and paying $V + \varepsilon$ for them with high probability. If the threshold is lower than the actual number of high-quality items in the population, the buyer will buy no (high-quality) items. Hence, unlike for a Devil's Menu, the simple mechanism is not robust if we introduce uncertainty about the total number of high-quality sellers in the population.

Some final remarks are in order with regard to Proposition 3. First, information about the (absolute) number of sellers owning low-quality items is immaterial for outcomes. This means, in turn, that information about the ratio of high-quality items to low-quality items both at the general population level and at the group level is also irrelevant for equilibrium behavior. Second, while random variable X is the same for all groups, the beliefs about the probability distribution of this random variable for different sellers and for the buyer may differ in any arbitrary way without affecting equilibrium behavior. Third, (8) specifies *sufficient* conditions about the uncertainty regarding the number of high-quality items in each group, but the result of Proposition 3 is likely to hold more generally. This is because what drives this result is the possibility for the buyer to set a *relative* reference point for groups based on the number of sellers claiming to be of high-quality and (some estimate of) the actual number of high-quality sellers that induces downward competition across groups by punishing the groups where more members claim to be of high type. The investigation of *necessary* conditions is left for the future.

6.2 Limited buyer commitment

Next, suppose that the buyer *cannot* commit to a single execution of a selected mechanism of her choice but only to the rules of any instance in which such a mechanism is run. This means, in particular, that the buyer will repeat the execution of the mechanism as long as it is beneficial for her to do so. What are the implications of this type of limited buyer commitment?

To answer this question, we consider the following (very stylized) dynamic game, denoted by \mathcal{DSG} , which takes place over two periods, indexed by $t = 1, 2$, for a given mechanism M selected by the buyer and a set of outcomes (messages, trades, and payments) Z :

1. In period $t = 1$, mechanism M takes place. If the output or mechanism M belongs to Z , the game ends. Otherwise, no trades and payments occur and the game moves to period $t = 2$.
2. In period $t = 2$, mechanism M takes place again and the game ends.

We make two simplifying assumptions, which are nonetheless unnecessary for the main thrust of our results. First, set Z consists of the outcomes for which, conditional on the fact that no seller plays weakly dominated strategies, the buyer believes with probability at least p that her expenses are used almost entirely to buy q high-quality. We assume that $p \in [0, 1]$ is common knowledge. Note that the case where $p = 1$ is one possibility. Then, if the outcome belongs to set Z , it is immaterial whether the buyer could have committed beforehand to not running mechanism M again after such an outcome. The reason is that we assume it to be common knowledge that in such cases, the buyer will be (sufficiently) satisfied with the outcome and will refrain from running the mechanism one more time.²³ By contrast, if the outcome does *not* belong to set Z , it is not credible that the buyer will not want to run the mechanism once more. Second, we proceed under the assumption that the *same* mechanism must be in place in both periods. This means that the menu of prices offered to sellers must be the same in both periods.²⁴

Because we consider a dynamic setting, beyond the quality of the good he owns, any seller i must be characterized by his discount factor, denoted by δ_i . We assume that the prior distribution

²³It suffices for it to be common knowledge after which outputs of the mechanism the buyer will have no incentive to run a mechanism again and thus her commitment power is not relevant. In particular, we do not need to bother with whether or not the buyer can learn the quality right after acquiring an object, as similar outcomes can be obtained in both cases.

²⁴If the value of ε for the simple mechanism can be chosen to be lower in period $t = 2$ than in period $t = 1$, one could consider that some sellers derive utility from an amount x of money that is equal to 0 if $x = 0$, to some constant κ_1 if $0 < x < V'$, and to some constant κ_2 if $x \geq V'$, where $\frac{q}{k} \cdot V < V' < V$. Utilities of this type, which are almost flat for low values of money, are intended to capture random behavior (see e.g. Moscati and Tubaro, 2011). The reason is that an agent with such an utility is indifferent between many outcomes as they all yield the same utility, and hence s/he might choose actions that would never be chosen by agents with strictly increasing utility.

of types (namely, of discount factors) is such that for any seller i ,

$$\delta_i = \begin{cases} 1 & \text{if } i = j, \\ \delta & \text{otherwise,} \end{cases}$$

where $\delta < 1$ and j denotes one particular seller. The main insights we obtain in this section extend to more general distributions of discount factors. It suffices that the number of low-quality sellers with discount factor one is not large. For her part, the buyer has discount factor δ . A discount factor equal to one captures random behavior in the sense that all lotteries involving the same outcome in period $t = 1$ with some probability \tilde{p} and in period $t = 2$ with probability $1 - \tilde{p}$, with $\tilde{p} \in [0, 1]$, yield the same utility.

For the analysis of (the games corresponding to) the above dynamic situation, we focus on perfect Bayesian equilibria in which agents cannot choose actions that are stage-dominated; we just call them equilibria. The only source of commitment for the buyer then stems from her publicly known goal of reaching an outcome of set Z and from the fact that she will execute a mechanism at most twice. This suffices to capture the idea of buyer's limited commitment, since all sellers anticipate that they will move to period $t = 2$ if the buyer deems it necessary to do so because she has not attained her goal in period $t = 1$. Under full commitment, by contrast, the buyer can ensure that there is only one period in which mechanism M will be (entirely) run and that all the sellers will believe so. In such cases, the underlying game is equivalent to the one-period model we have analyzed in the previous sections (for the Devil's Menu) and above in the present section (for the simple mechanism). In either case, the buyer buys with probability one the intended number of high-quality items at (approximately) the minimum expenses described in Section 4, and thus derives utility (approximately) equal to $u_B(z^*)$. Recall that z^* has been defined in (3). Clearly, it must be the case that $z^* \in Z$.

Then, one can show the following result in the case of the simple mechanism.

Proposition 4. *Let $\delta \in [0, 1]$. If the simple mechanism is considered, there is an equilibrium of game \mathcal{DSG} in which*

- (i) *The low-quality seller j with discount factor one sends message h in period $t = 1$,*
- (ii) *The buyer acquires m high-quality items in period $t = 2$, and derives utility approximately equal to $\delta \cdot u_B(z^*)$.*

The above result is clearly suboptimal from the buyer's perspective, as well as from the perspective of all sellers different from j , particularly if δ is much lower than one. This means that lack of commitment power can have severe consequences for the buyer if she uses the simple mechanism. The proof of Proposition 4 is almost straightforward: in period $t = 2$, the last period, the buyer acquires the desired number of high-quality items and each low-quality seller i

(including j) sends message l and derives utility $\delta_i \cdot \varepsilon$. This follows from Theorem 1. Since $\delta_j = 1$, seller j is indifferent between receiving a monetary payment ε in period $t = 1$ and in period $t = 2$. Because all high-quality sellers send message h in every period, seller j cannot obtain a payoff that is higher than ε .

Consider now the Devil’s Menu with four prices of Theorem 1. We can show the following result:

Proposition 5. *Suppose that the mechanism of game \mathcal{G}_4 is considered for game \mathcal{DSG} and assume that*

$$\delta < \frac{\bar{k}}{q \cdot (\bar{k} - 1 + q)}. \quad (11)$$

Then, in any equilibrium of game \mathcal{DSG} the buyer acquires the desired number of high-quality items in period $t = 1$ and derives utility (approximately) equal to $u(z^)$.*

Proof. See Appendix. □

That is, a Devil’s Menu ensures second-best outcomes for the buyer even under limited commitment, provided that the discount factor is (moderately) low as required by Condition (11).²⁵ Why is this the case for a Devil’s Menu and not for the simple mechanism? The reason is that the (off-equilibrium path) threat that any low-quality seller can impose if he contributes to an excessive number of h messages in one group—by sending a message h himself—becomes ineffective under a Devil’s Menu, since the buyer will simply buy the high-quality items from another group. That is, the possibility to induce downward group competition is a substitute for the buyer’s lack of commitment power, as it helps to resolve her adverse selection problem.

6.3 Presence of saboteurs

In this section, we assume that there is one low-quality seller, say j , who derives utility $-u_B$, where u_B denotes the buyer’s utility.²⁶ This means that seller j ’s (main) goal is to prevent the buyer from acquiring the desired number of high-quality items, so we call seller j a *saboteur*. There are two main possibilities. First, assume that neither the buyer nor the other sellers know that such a seller exists. Then, our previous analysis shows that (i) under the simple mechanism, the buyer will not acquire any object, and (ii) under a Devil’s Menu (with four prices), the buyer will acquire the desired objects and spend exactly as in Theorem 1. The difference between the simple mechanism and a Devil’s menu stems from the fact that, in the latter, seller j has no means to harm the buyer, since sending message h only affects his group not being selected.

²⁵The bound given by Condition (11) is not tight, so the result in Proposition 5 holds more generally. A low discount factor ensures that sellers different from j will not risk that the game moves into period $t = 2$.

²⁶Our analysis can be easily extended to the case where there are saboteurs in at most $\bar{k} - 1 - q$ groups of the partition.

Second, assume that it is common knowledge that seller j exists, although his identity is private information. On the one hand, under the simple mechanism, the buyer can acquire all high-quality objects but she will also need to acquire j 's item, as she cannot tell the quality difference between seller j 's item and that of all high-type sellers. It suffices for her to commit to buying $m + 1$ items. On the other hand, under a Devil's Menu (with four prices) the buyer will again acquire exactly the desired objects and spend exactly as in Theorem 1.²⁷

7 Applications and Implementation

In this section, we first discuss two potential applications of a Devil's Menu. Then we discuss what the real-world implementation of a Devil's Menu could look like.

7.1 Market makers and regulations

For our analysis, we have focused on a buyer with a large willingness to buy items of high quality (up to some number), who faces a budget constraint. We have then shown that a crucial feature of a Devil's Menu is that it allows the buyer to learn the quality of the items she is buying. This property opens up a variety of applications of the mechanism.

For instance, an agent—whether benevolent or not—with a sufficiently large budget could use a Devil's Menu to adopt a market maker role in a market for lemons. After acquiring a certain amount of high- and low-quality commodities by using a Devil's Menu, this agent could resell these items to other buyers with prices differentiated according to quality. This action could incidentally restore the well-functioning of a market that was previously plagued with adverse selection, say by affecting (the beliefs on) the ratio of high-quality items to low-quality items that remain in the market. The market-maker agent could signal the quality of her acquired goods to buyers in the resale market if she could credibly disclose the messages h or l she obtained from the sellers in the original market by applying the Devil's Menu. This credibility could be achieved by several means.

One option would be to run the Devil's Menu on a publicly accessible domain such as a blockchain so that everybody could observe the messages h or l sent by the sellers. Again, if a low-quality seller sends message h in a Devil's Menu, his group will not be chosen for buying and reselling. Hence, public observation of messages h and l are sufficient to organize the market.

Absent the possibility to learn the messages sent in the Devil's Menu, another option would be that the buyers in the resale market could learn the share of high- and low-quality goods

²⁷One can easily verify that the arguments in Theorem 1 carry over to this setup. It suffices to note that seller j 's group will belong to set O and that $\frac{q-c}{k-c-1} \leq \frac{q}{k}$.

the market maker had bought, possibly with the help of Big Data approaches. Suppose this was possible, and the market maker bought both item types in equal amounts and then offered two baskets of goods of equal size in the resale market, each one (allegedly) composed of items of each type and associated with different prices. Then it would be credible that one basket contains only high-quality items and the other only low-quality items. This is because switching items from one group to the other would not be profitable for the market maker since doing so would not affect her revenues.

The possibility to learn the quality of certain goods in certain markets through a Devil’s Menu opens up opportunities not only for market makers but also for improving the regulations that govern the functioning of the markets. For instance, a government agency with a large budget could run a Devil’s Menu on a large scale by buying up many items from a subset of groups, and organize liquid resale markets to restore the functioning of markets plagued with adverse selection.

7.2 Application to collective decisions

The (formal) setup we have analyzed can also be used to model political influence over collective decisions that are made through voting, by, say, a polity or a committee.²⁸ It suffices to identify the buyer with an outside party—which we call the *adversary*—that wishes to influence the voting outcome by buying some votes, and to identify the sellers with the voters. The latter can be citizens or committee members.

To apply our results, we assume that both the adversary’s and the voters’ reservation price are fully determined exogenously (e.g. as in Laine, 1977; Lalley and Weyl, 2018; Mueller, 1973). This is a reasonable assumption in some cases—e.g. it is often imposed in committee voting (see e.g. Dal Bo, 2007). It also sidesteps the complex underlying voting process where the valuations of voters for their votes are determined through pivotality from the citizens’ intrinsic valuation of the right to vote. This modelling assumption enables us to focus on the adversary’s attempts to buy votes. Our setup is the most basic one featuring such a property. From this perspective it is also natural to assume, as we do in Section 3.1, that the adversary is interested in buying a number of votes that suffices for her to be pivotal for the outcome.

Our setup considers one adversary and, as we have seen, features a one-principal, many-agents problem. There is a variety of situations with one-sided special interest (see e.g. Leaver and Makris, 2006). This means that our assumption that there is one adversary is not restrictive from a real-world perspective.

To build on our previous analysis, we also need to consider that voters can be partitioned in

²⁸The equivalence holds not only for the majority rule applied to the entire population but also for more general voting rules such as higher majority thresholds (different from unanimity) and multi-district, first-past-the-post elections.

groups. Most political systems divide the citizenry in several districts. For each of these, polls provide aggregate information about the percentage of citizens in each district preferring each alternative, but not what alternative particular individuals in this district favor. In the case of committees, say parliaments, aggregate information of this type can also be available if, for example, there is a geographical component to the decision beyond ideology. This is typically the case for decisions where to build a nuclear plant or where to hold an international event such as the Olympic Games. Another possibility is to consider a “metacommittee” composed of all regional parliaments of a country.

In the following, we outline two specific applications of our results to *binary* collective decisions taken through voting.

7.2.1 Application to lobbying

In both elections and committee decisions made through voting, votes are cast in favor of some alternative. This means that a vote (or a ballot) is not a standard consumption good. Voters may obtain the same utility no matter who is eventually casting their ballot, as long as it is done in favor of their preferred alternative. In accordance with this view, we assume that the adversary herself has some preference regarding the two alternatives at hand—and hence she does not simply act as a reseller of ballots—and this preference is common knowledge. The adversary can be a lobbying group or a special interest group seeking to influence the voting decision, say by buying some citizens’ votes or bribing some committee members. In the case of an electorate, it is clear that the adversary has, in general, little means to know *ex ante* what the preferences of particular individuals are. For committees, even if voting records are public, there might still be (large) uncertainty about whether particular committee members favor one alternative over the other (see Buzard and Saiegh, 2016; Le Breton and Salanié, 2003, for the case of ideological uncertainty in a committee).

As already mentioned, we consider that voting takes place to decide which of two alternatives should be implemented.²⁹ Absent the adversary, any voter’s reservation price is V regardless of his preferences. Parameter V reflects both how much he cares about his preferred alternative being implemented compared to the case where the other alternative is implemented, and the degree to which he cares that his vote is cast in favor of the alternative he prefers. In the presence of the adversary, however, the voters who prefer the same alternative as the adversary would anticipate that if they sold their vote to her, she would then use it as they intended themselves, with the same consequences for the outcome. Hence, their willingness to sell would drop from V to 0.³⁰ This means that as far as modeling goes, we can assume that the voters

²⁹Voting does not necessarily have to be public but the adversary must have means to contract on individual voting decisions. This guarantees that the adversary can actually “buy a vote”. A possibility is to use a smart contract. A similar comment applies to Section 7.2.2.

³⁰The inefficiencies associated with spending resources trying to persuade voters who are already persuaded has been analyzed e.g. in Casas (2018).

who prefer the same alternative as the adversary hold low-quality items, while the voters who prefer the other alternative to be implemented hold high-quality items.³¹

Our results show that the adversary (or principal) can buy as many votes as she wants at the second-best price, and, in particular, she can capture a committee's decision even if her pocket is not very deep. These findings complement Dal Bo (2007) in that they show additional ways how a committee or polity may become vulnerable to external influence. It is crucial for our and his results that the adversary can make payment promises that are contingent not only on voting behavior at the individual level but also at the aggregate level. To be effective, these promises must reward *pivotality* off-equilibrium, and at the same time yield incentives for all sellers of a same type to act homogeneously so that no one ends up being pivotal on the equilibrium path. Unlike Dal Bo (2007), we assume that agents have heterogeneous preferences. Hence the main inefficiency source for the principal in our model is that she needs to tell agents apart with regard to preference to avoid excessive payments. Importantly, as we have seen in Section 7, our results are very robust. In particular, they hold even if the principal has low commitment power.³² For their part, agents cannot commit to selling: they simply hear offers which they might eventually accept if they find it in their best interest to do so, but do not sign contracts ahead of voting. This aims at reflecting real-world situations in which swaying a collective decision through some sort of bribing is possible.

Finally, our finding that voter collusion either does not harm the adversary at all (when collusion occurs only within districts) or only does if it is sufficiently significant (when collusion occurs across districts) is also relevant for the design of political institutions.³³ It means that political systems where parties are sufficiently large and cut across regional lines might be most effective against vote buying and manipulation of voting outcomes by special interest groups since collusion across districts is difficult or impossible to organize. This rationalizes the view that a (large) political party is an effective way for voters and committee members sharing the same interests not only to coordinate themselves, but also to defend themselves against external influence.

7.2.2 Application to decoy ballots

In the past few years, electronic voting has become popular in many countries.³⁴ The possibility that voting can be carried out electronically opens up many new options for representative and

³¹Alternatively, one can assume that all citizens have the same preferences but differ in terms of the extent to which they care about the alternative being implemented or, equivalently, in the extent of their corruptibility. From this perspective, voters whose reservation price is V care more than voters whose reservation price is 0. The latter can sell their vote at any price, but are willing to pose as if they were less corruptible to obtain higher rents.

³²See Footnote 11 in Dal Bo (2007).

³³Recall Footnote 19.

³⁴See https://en.wikipedia.org/wiki/Electronic_voting_by_country, retrieved on 18 September, 2017.

direct democracies alike, as the marginal cost of voting is typically lower than for traditional voting. One possibility that has been argued is to randomly select a subgroup of citizens from the entire population and have each of these citizens vote on a single issue. This is called *random sample voting* (Chaum, 2016).³⁵ Assuming that the chosen subpopulation (or *sample voting group*) represents the entire citizenry, random sample voting may improve decision-making by yielding the same decision as standard voting, albeit at a less costly voter-participation level. For instance, several issues could be considered at once by having different samples vote separately on each of them. This way, citizens may acquire more information only about the issue on which they have a say, and decisions may be better informed.

Electronic voting procedures of this type entail risks, particularly when an adversary is interested in buying some citizens' right to vote in a particular instance of random sample voting. To prevent this, Chaum (2016) has proposed the following mechanism: not only the members of the sample voting group receive a ballot, but so do all other citizens. The difference is that only the ballots of the members of the sample voting group are real, i.e., they alone are counted. The remaining ballots act as a decoy—and hence they are called *decoy ballots*. Crucially, whether a ballot is real or a decoy is a citizen's private information, so the adversary cannot distinguish real ballots from decoy ballots.³⁶ Voting systems based on decoy ballots revolve around the idea that selling them to an adversary is valuable for society—and can constitute a social norm—, since doing so may prevent adversaries from buying a large number of real ballots and then manipulating the voting outcome.³⁷ Parkes et al. (2017) have recently shown that under some assumptions, decoy ballots *do* indeed discourage attempts by an adversary with a limited budget to try and buy (real) ballots.³⁸

For an application of our results to a setting with decoy ballots, it suffices to identify real-ballot holders and decoy-ballot holders with high-quality and low-quality sellers, respectively. Our findings then suggest that decoy ballots lose their discouragement power once the adversary has applied a Devil's Menu, since she can buy them at negligible costs.³⁹

³⁵See also Gersbach et al. (2019).

³⁶One could conceive of a ballot as a password in the electronic voting system, where a decoy ballot is simply an invalid password.

³⁷Parallels can be drawn between decoy ballots preventing vote-buying and the idea of producing fake, harmless drugs and selling them in the drug market to destroy this market.

³⁸For other types of attack on electronic voting, see Basin et al. (2017).

³⁹In real voting settings for large populations, it is reasonable to expect the number of real voters in any given district to be fixed—and easy to guess for the adversary—, since otherwise the aggregate preferences of the society may not be well represented by the citizens receiving the real ballots. This would contravene the minimal requirement that no district should obtain an advantage in the number of real ballots allocated as a result of some randomness. One possibility is that while such numbers are fixed, real ballots are allocated randomly *within* each district. This would ensure that the form of seller partition needed for the functioning of a Devil's Menu holds.

7.3 Implementation

Implementing a Devil’s Menu could be particularly appealing on the Internet, in which case item purchase could be set up as in online auctions. For instance, the implementation of any of the above mechanisms can be effected using so-called *smart contracts*.⁴⁰ These contracts are computer protocols intended to facilitate, verify, and enforce the exchange between individuals. A Devil’s Menu can be coded in any programming language and run on the blockchain. Payments may be carried out in any of the crypto-currencies implementing smart contracts. Since the main property of smart contracts is that they are self-executing and self-enforcing, the adversary would be committed to the protocol and the payments. This is in accord with the assumptions we made for Theorem 1 about agent behavior. Agents can join the contract at any time before timeout by sending their messages.

8 Conclusion

We have presented a mechanism that can help to solve the adverse selection problem faced by one buyer who is interested in buying goods from multiple, heterogeneous sellers. The novelty of our approach is that the pool of sellers can be partitioned in different groups according to some characteristics different from quality in a way that the buyer has information about the (*relative*) number of high-quality sellers in each group. In such cases, the buyer may significantly reduce her expenditures on low-quality sellers and thus mainly use her budget to buy high-quality items. This can be achieved through a multiple-stage buying mechanism that builds on seller partition to create downward pressure on the competing groups, thereby allowing the buyer to extract all surplus from the market.

Our findings are robust with respect to weaker commitment power of the buyer, lack of information about the distribution of types, and the presence of market saboteurs. All these features could hurt the buyer. Our results also highlight the importance on and off the equilibrium path that the buyer is committed to buying some low-quality items. This is necessary for her to be able to screen low-quality items from high-quality items.

As a first application of our analysis and results, we have suggested that a Devil’s Menu might be a useful tool for market makers in markets that exhibit damaging levels of adverse selection. As a second application, we have shown that lobbying or special interest groups could capture a political body’s decision if they used a Devil’s Menu, even if they have low commitment power and they lack deep pockets. A third insight of our paper is that (electronic) voting systems based on decoy ballots may be vulnerable to sophisticated attacks. This would render decoy ballots an ineffective tool against vote-buying by trying to blow up the adversary’s budget.

⁴⁰We refer to Wood (2014). See also https://en.wikipedia.org/wiki/Smart_contract, retrieved on 16 November, 2017.

While the latter two applications exhibit some elements that are specific to voting, the market for votes and political influence shares some important features with many other markets. Hence the insights we develop for these applications extend beyond voting markets.

While we have limited ourselves to a simple scenario, numerous further scenarios are also possible. First, sellers may have more than one item to sell or, similarly, they could adopt a unified strategy. Second, there could be two (big) buyers in the market trying to buy items. In this scenario, it would be interesting to know if both would collude to use a Devil’s Menu to screen low-quality sellers, or if, instead, both buyers would use such a mechanism to throw low-quality sellers against each other, thereby incurring (potentially big) losses. Third, as a further application, a Devil’s Menu could be tailored to situations where some buyer wants to buy a valuable secret from a group of individuals, say, a firm or a bureaucratic unit, knowing that only one of the individuals has the secret. Examples could be a password or the identity of the person who has important knowledge. These issues are left to future research.

References

- Akerlof, G. A. (1970). The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84(3):488–500.
- Aoyagi, J. and Adachi, D. (2018). Economic implications of blockchain platforms. *arXiv preprint arXiv:1802.10117*.
- Basin, D., Gersbach, H., Mamageishvili, A., Schmid, L., and Tejada, O. (2017). Election Security and Economics: It’s All about Eve. In *Proceedings of the Second International Joint Conference on E-Vote-ID, Bregenz, Austria*, pages 1–20.
- Biglaiser, G. and Friedman, J. W. (1994). Middlemen as guarantors of quality. *International Journal of Industrial Organization*, 12(4):509–531.
- Bó, I. and Hakimov, R. (2020). Pick-an-object mechanisms, mimeo. *Available at SSRN*.
- Börger, T. (2015). *An Introduction to the Theory of Mechanism Design*. Oxford University Press, New York, NY.
- Buzard, K. and Saiegh, S. (2016). Lobbying and legislative uncertainty. *Unpublished manuscript, Syracuse University, NY, USA*.
- Casas, A. (2018). Distributive politics with vote and turnout buying. *American Political Science Review*, 112(4):1111–1119.
- Chaum, D. (2016). Random Sample Voting. *White Paper*.
- Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 1(1):17–33.

- Cong, L. W. and He, Z. (2019). Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5):1754–1797.
- Cr mer, J. and McLean, R. P. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257.
- Dal Bo, E. (2007). Bribing voters. *American Journal of Political Science*, 51(4):789–803.
- Daley, B. and Green, B. (2012). Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504.
- Dekel, E., Jackson, M. O., and Wolinsky, A. (2008). Vote Buying: General Elections. *Journal of Political Economy*, 116(2):351–380.
- Elfenbein, D. W., Fisman, R., and McManus, B. (2015). Market structure, reputation, and the value of quality certification. *American Economic Journal: Microeconomics*, 7(4):83–108.
- Fang, H. (2001). Social Culture and Economic Performance. *American Economic Review*, 91(4):924–937.
- Felgenhauer, M. and Gr ner, H. P. (2008). Committees and special interests. *Journal of Public Economic Theory*, 10(2):219–243.
- Finan, F. and Schechter, L. (2012). Vote-Buying and Reciprocity. *Econometrica*, 80(2):863–881.
- Freixas, X., Guesnerie, R., and Tirole, J. (1985). Planning under incomplete information and the ratchet effect. *Review of Economic Studies*, 52(2):173–191.
- Fuchs, W. and Skrzypacz, A. (2015). Government interventions in a dynamic market with adverse selection. *Journal of Economic Theory*, 158:371–406.
- Gersbach, H., Mamageishvili, A., and Tejada, O. (2019). The effect of handicaps on turnout for large electorates: An application to assessment voting. *CEPR Working Paper*.
- Groseclose, T. and Snyder, J. M. (1996). Buying supermajorities. *American Political Science Review*, 90(2):303–315.
- Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4):617–631.
- Hart, O. D. and Tirole, J. (1988). Contract renegotiation and coasian dynamics. *Review of Economic Studies*, 55(4):509–540.
- Kawai, K. (2014). Dynamic market for lemons with endogenous quality choice by the seller. *Games and Economic Behavior*, 84:152–162.
- Kim, K. (2012). Endogenous Market Segmentation for Lemons. *RAND Journal of Economics*, 43(3):562–576.

- Kim, K. (2017). Information about sellers' past behavior in the market for lemons. *Journal of Economic Theory*, 169:365–399.
- Laine, C. R. (1977). Strategy in Point Voting: A note. *Quarterly Journal of Economics*, 91(3):505–507.
- Lalley, S. P. and Weyl, E. G. (2018). Quadratic Voting: How Mechanism Design Can Radicalize Democracy. In *AEA Papers and Proceedings*, volume 108, pages 33–37.
- Le Breton, M. and Salanié, F. (2003). Lobbying under political uncertainty. *Journal of Public Economics*, 87(12):2589–2610.
- Leaver, C. and Makris, M. (2006). Passive industry interests in a large polity. *Journal of Public Economic Theory*, 8(4):571–602.
- Lin, M., Prabhala, N. R., and Viswanathan, S. (2013). Judging borrowers by the company they keep: Friendship networks and information asymmetry in online peer-to-peer lending. *Management Science*, 59(1):17–35.
- Louis-Sidois, C. and Musolff, L. (2020). Collusion-proof vote buying, mimeo. *Available at SSRN*.
- Mailath, G. J., Samuelson, L., and Shaked, A. (2000). Endogenous Inequality in Integrated Labor Markets with Two-sided Search. *American Economic Review*, 90(1):46–72.
- Moreno, D. and Wooders, J. (2016). Dynamic markets for lemons: Performance, liquidity, and policy intervention. *Theoretical Economics*, 11(2):601–639.
- Moscatti, I. and Tubaro, P. (2011). Becker random behavior and the as-if defense of rational choice theory in demand analysis. *Journal of Economic Methodology*, 18(2):107–128.
- Mueller, D. C. (1973). Constitutional Democracy and Social Welfare. *Quarterly Journal of Economics*, 87(1):60–80.
- Muthoo, A. and Mutuswami, S. (2011). Imperfect competition and efficiency in lemons markets. *Economic Journal*, 121(552):580–594.
- O'Dair, M. and Owen, R. (2019). Financing new creative enterprise through blockchain technology: opportunities and policy implications. *Strategic Change*, 28(1):9–17.
- Parkes, D. C., Tylkin, P., and Xia, L. (2017). Thwarting Vote Buying Through Decoy Ballots. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 3784–3790.
- Peterson, J. R. and Schneider, H. S. (2017). Beautiful lemons: Adverse selection in durable-goods markets with sorting. *Management Science*, 63(9):3111–3127.

- Piketty, T. (1993). Implementation of first-best allocations via generalized tax schedules. *Journal of Economic Theory*, 61:23–23.
- Riley, J. G. (2001). Silver Signals: Twenty-five Years of Screening and Signaling. *Journal of Economic Literature*, 39(2):432–478.
- Schneider, M. T. (2014). Interest-group size and legislative lobbying. *Journal of Economic Behavior & Organization*, 106:29–41.
- Skreta, V. (2006). Sequentially Optimal Mechanisms. *Review of Economic Studies*, 73(4):1085–1111.
- Skreta, V. (2015). Optimal Auction Design under Non-commitment. *Journal of Economic Theory*, 159:854–890.
- Spence, M. (1973). Job Market Signalling. *Quarterly Journal of Economics*, 87(3):355–374.
- Stiglitz, J. E. (1975). The Theory of “Screening,” Education, and the Distribution of Income. *American Economic Review*, 65(3):283–300.
- Vickrey, W. (1961). Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16(1):8–37.
- Wood, G. (2014). Ethereum: A Secure Decentralized Generalised Transaction Ledger. *Mimeo*.
- Zavolokina, L., Miscione, G., and Schwabe, G. (2019). Buyers of lemons: How can a blockchain platform address buyers needs in the market for lemons? *Electronic Markets*, pages 1–13.

Appendix

Proof of Theorem 1. Let $\sigma = (\sigma_i)_{i \in N}$ be a strategy profile that constitutes an equilibrium of \mathcal{G}_4 . Because all sellers with a high-quality item send message h , we can assume w.l.o.g. that

$$1 \leq \rho_1 \leq \rho_2 \leq \dots \leq \rho_{\bar{k}}.$$

We next observe that in any equilibrium, set O is empty and thus o is equal to 0. Indeed, assume that $o > 0$. Then, in any of the groups k that belong to O such that $\rho_k > 1$, any low-quality seller i for which $\sigma_i = h$ have an incentive to deviate and send message l . This deviation strictly improves expected utility. The reason is as follows: As a consequence of the deviation, either group N_k becomes a selected group, and the expected utility is strictly larger than ε (since $\theta > \varepsilon$), or group N_k still belongs to O , in which case the deviation increases the payment to seller i by ε .

Since $o = 0$ in any equilibrium, it follows that $t = \bar{k} - c$. We prove next that σ cannot be an equilibrium if $\rho > 1$. This property can be derived as follows: Suppose that $\rho > 1$. Then, in each group belonging to T , there is at least one low-quality seller who has sent message h in σ . Consider one of these groups, say N_k . The probability that N_k is selected by fair randomization is equal to $\frac{q-c}{t}$. Now let i be any low-quality seller for which $\sigma_i = h$ and who belongs to N_k . Then, i 's expected payoff is equal to $\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon$. By contrast, if i deviates from his strategy in σ and sends message l , group k is selected with probability 1, and the payoff to low-quality seller i is equal to θ . Consequently, set T consists of $t - 1$ groups, and set C consists of $c + 1$ groups. The payoff in the latter case, namely θ , is strictly higher than $\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon$, since

$$\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon = \frac{q-c}{\bar{k}-c} \cdot V + \varepsilon \leq \frac{q}{\bar{k}} \cdot V + \varepsilon < \frac{q}{\bar{k}} \cdot V + 2\varepsilon \leq \theta.$$

Accordingly, it must be that $\rho = 1$. To sum up, we have shown that in any equilibrium of game \mathcal{G}_4 , we have $\rho = 1$ and $o = 0$. Finally, it can easily be seen that σ^* is indeed an equilibrium and that it is the only one satisfying $\rho = 1$ and $o = 0$. This completes the proof. \square

Proof of Proposition 1. Proposition 1 can be proved by backward induction. At the last date, there are $\bar{k} - q + 1$ groups, of which only one is selected by the buyer. By the same arguments as in the proof of Theorem 1, one derives the critical value for θ ensuring that sellers must act as prescribed by σ^* at the last date. It suffices to note that

$$\frac{1}{\bar{k}-q+1} \cdot (V + \varepsilon) + \frac{\bar{k}-q}{\bar{k}-q+1} \cdot \varepsilon = \frac{1}{\bar{k}-q+1} \cdot V + \varepsilon < \frac{1}{\bar{k}-q+1} \cdot V + 2\varepsilon.$$

At all preceding dates, we have a smaller critical value for θ . This ensures that σ^* prescribes at

any such stage only those actions that are compatible with equilibrium behavior, anticipating that σ^* will also determine the actions chosen at any future date. The reason is that there are more groups from which one is selected. This completes the proof. \square

Proof of Proposition 2. Under this modified Devil's Menu with four prices, sending message h weakly dominates sending message l in the case of high-quality sellers, because the payoff matrix for high-quality sellers remains the same as before, since they only sell their item if they have sent message h and their group has been selected. If sellers behave in accordance with σ^* , the payoff of a low-quality seller is 2ε , with every group having the same chance of being selected. A deviation by a low-quality seller i switching from $\sigma_i = \sigma^* = l$ to $\sigma_i = h$ would then cause the deselection of his group. In such a case, seller i would end up with a payoff of ε . This means that such a deviation is not profitable, which completes the proof. \square

Proof of Theorem 2. As in the proof of Theorem 1, abstaining is weakly dominated for all sellers, while sending message l is weakly dominated for all high-quality sellers. We also observe that $o = 0$ must also hold in any equilibrium $\sigma = (\sigma_i)_{i \in N}$. Indeed, if $o > 0$, there must exist low-quality sellers of groups in O who sent message h . However, these sellers can strictly improve their payoff by sending message l . The reason is as follows: Let $i \in N_k$ be a low-quality seller such that $N_k \in O$ and $\sigma_i = h$. If seller i deviated and sent message l instead, there are three possible cases.

First, N_k could belong to set C after i 's deviation, which would result in a payoff strictly larger than ε (since $V - \varepsilon > \varepsilon$). Second, N_k could belong to set T after i 's deviation, which would result in a strictly larger expected payoff. This follows from $\theta > \varepsilon$ and the fact that $N_k \in T$ implies $t \geq 1$, which leads to

$$\frac{q-c}{t} \cdot \theta + \frac{t-q+c}{t} \cdot 2\varepsilon > \varepsilon.$$

Third, N_k could still belong to set O after i 's deviation, in which case the deviation would increase the payment to seller i (since $2\varepsilon > \varepsilon$). All in all, we obtain that $o = 0$, and hence $t = \bar{k} - c$.

Next we show that $\rho > 1$ cannot occur in equilibrium either. Suppose that $\rho > 1$ for σ . Then, in each group that belongs to T , there must be at least one low-quality seller i such that $\sigma_i = h$. Consider one of these groups, say N_k , and a low-quality seller i in N_k who has sent message h . The probability that N_k is selected by fair randomization is equal to $\frac{q-c}{t}$. Then i 's expected payoff is equal to $\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon$. By contrast, if i deviates from his strategy in σ and sends message l , set T consists of $t - 1$ groups and set C consists of $c + 1$ groups. In this latter case, group N_k is chosen with probability 1, and the payoff to low-quality seller i is equal to $V - \varepsilon$. We claim that the latter payoff is strictly larger than $\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon$ if ε is sufficiently small, so i is better off if he sends message l . Indeed, using $t = \bar{k} - c$ and $q < \bar{k}$,

and taking $\varepsilon > 0$ arbitrarily small, we obtain

$$\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon = \frac{q-c}{k-c} \cdot V + \varepsilon \leq \frac{q}{k} \cdot V + \varepsilon < V - \varepsilon.$$

Finally, $\rho = 1$ is an equilibrium if no low-quality seller wants to send message h . In the unique equilibrium featuring $\rho = 1$ (which is exactly σ^*), the expected payoff for a low-quality seller is

$$\frac{q}{k} \cdot \theta + \frac{\bar{k}-q}{\bar{k}} \cdot 2\varepsilon,$$

since all groups are in set T . If the low-quality seller sent instead message h , his group would belong to O , which would result in a strictly lower payoff of ε (since $\theta > \varepsilon$). Hence, the deviation is not profitable, which completes the proof. \square

Proof of Proposition 3. Let $\sigma = (\sigma_i)_{i \in N}$ be a strategy profile that constitutes an ex post equilibrium in pure strategies of game \mathcal{G}_4 . As in the rest of the paper, we assume that no seller plays a weakly dominated strategy. Then consider i to be any low-quality seller with $\sigma_i = h$ (otherwise we are done). We use N_k to denote the partition group to which i belongs. For a given draw of types t_{-i} , let $\mathbb{1}_C, \mathbb{1}_T, \mathbb{1}_O$ be the indicator functions (conditional on σ) that group N_k belongs to C, T, O , respectively. This means that for given t_{-i} and σ_{-i} , ex-post utility when $\sigma_i = h$ is

$$u_i(\sigma_i = h) = \mathbb{1}_C \cdot (V + \varepsilon) + \mathbb{1}_T \cdot \left(\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon \right) + \mathbb{1}_O \cdot \varepsilon,$$

while ex-post utility when $\sigma_i = l$ is

$$u_i(\sigma_i = l) = \mathbb{1}_{C'} \cdot \theta + \mathbb{1}_{T'} \cdot \left(\frac{q-c'}{t'} \cdot \theta + \frac{t'-q+c'}{t'} \cdot 2\varepsilon \right) + \mathbb{1}_{O'} \cdot 2\varepsilon,$$

where C', T', O' stand respectively for the sets of chosen, tied, and out when $\sigma_i = l$.

Now set t_{-i} and distinguish three cases.

Case I: $\mathbb{1}_O = 1$

In this case, either $\mathbb{1}_{T'} = 1$ or $\mathbb{1}_{O'} = 1$, so

$$u_i(\sigma_i = l) - u_i(\sigma_i = h) = \left[\mathbb{1}_{T'} \cdot \left(\frac{q-c'}{t'} \cdot \theta + \frac{t'-q+c'}{t'} \cdot 2\varepsilon \right) + \mathbb{1}_{O'} \cdot 2\varepsilon \right] - \varepsilon > 0,$$

where the inequality holds (see the proof of Theorem 1) since

$$\theta \geq \frac{q}{k} \cdot V + 2\varepsilon.$$

This means that seller i has an ex-post incentive to deviate from his strategy $\sigma_i = h$, regardless

of his beliefs about t_{-i} . Hence, there is no ex post equilibrium in which a low-type seller i who chooses $\sigma_i = h$ is from a group N_k that belongs to set O . Suppose now that $o > 0$ and let $N_k \in O$. Let also $N_j \notin O$ be a group with $j \succeq k$ (such a group must exist by construction of the Devil's Menu). From (8) and (10), and using the fact that high-type sellers always report their type truthfully because doing otherwise is weakly dominated, we obtain

$$m_k - m_j = n_k^H - m_j \leq n_k^H - n_j^H = \mathbb{E}[n_k^H] - \mathbb{E}[n_j^H], \quad (12)$$

where the first equality holds because we have shown that in a group belonging to O all sellers report their types truthfully. Accordingly, $k \succeq j$. But this is in contradiction with $N_k \in O$ and $N_j \notin O$. Hence, in any ex post equilibrium, it must be that $o = 0$. We stress that this holds independently of the sellers' beliefs about the type distribution for the other sellers.

Case II: $\mathbb{1}_T = 1$

In this case, $\mathbb{1}_{C'} = 1$ and

$$u_i(\sigma_i = l) - u_i(\sigma_i = h) = \theta - \left(\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon \right) > 0,$$

since

$$\frac{q-c}{t} \cdot (V + \varepsilon) + \frac{t-q+c}{t} \cdot \varepsilon < \frac{q}{k} \cdot V + 2\varepsilon \leq \theta,$$

where the strict inequality holds since $o = 0$. Hence, there is no ex post equilibrium in which a low-type seller i that chooses $\sigma_i = h$ is from a group N_k that belongs to set T . As in the above case, this holds independently of the beliefs held by the sellers about the type distribution for the other sellers.

Case III: $\mathbb{1}_C = 1$

Let N_j be a group belonging to T with $k \succ j$. Such a group must exist due to the construction of a Devil's Menu, as we have already seen that set O must be empty in all ex post equilibria. From Case II, we know that $m_j = n_j^H$. Moreover, from (8) and (10), and using the fact that high-type sellers always report their type truthfully,

$$m_k - m_j = n_k^H - m_j > n_k^H - n_j^H = \mathbb{E}[n_k^H] - \mathbb{E}[n_j^H], \quad (13)$$

so $j \succ k$, a contradiction. Hence, there is no ex post equilibrium in which a low-type seller i who chooses $\sigma_i = h$ is from a group N_k that belongs to set C . Once more, this holds independently of the beliefs held by the sellers about the type distribution for the other sellers.

Finally, we claim that if all sellers report their type truthfully, this is an ex post equilibrium. Recall that high-quality sellers always send message h , as doing otherwise is weakly dominated. To prove the claim, let σ^* denote the truth-telling strategy profile (conditional on type) and suppose that a low type-quality citizen decides to deviate from σ^* and send message h . By

doing so, the difference in ex-post utility—assuming (8) is known—is equal to

$$\varepsilon - \left(\frac{q}{\bar{k}} \cdot \theta + \frac{\bar{k} - q}{\bar{k}} \cdot 2\varepsilon \right) < 0.$$

Hence, seller i has no incentive to deviate, so σ^* is an ex post equilibrium. □

Proof of Proposition 5. Consider an equilibrium of game \mathcal{DSG} . The first observation is that for any given period, it is weakly dominated for high-quality sellers to not reveal their type. This means that these sellers will send message h in both period $t = 1$ and period $t = 2$. Hence, for our analysis, it suffices to focus on sellers owning items of low quality. Let i denote a low-quality seller different from seller j . Assume now that no matter what seller i 's message in period $t = 1$ is, the game moves to period $t = 2$. In this case, the payoff to seller i does not depend on his action in period $t = 1$. Because we assume that sellers do not play stage-dominated strategies, we can further focus on the case where the probability that the game will proceed to period $t = 2$ depends on seller i 's message in period $t = 1$.

Note that if period $t = 2$ is attained along the equilibrium path, all sellers will report their type truthfully in such a period. This follows from Theorem 1. This means that in this case, seller i (belonging to some group N_k) expects utility

$$\bar{u}_i := \delta \cdot \left[\frac{q}{\bar{k}} \cdot \theta + \frac{\bar{k} - q}{\bar{k}} \cdot 2\varepsilon \right] > 0,$$

while seller j expects

$$\bar{u}_j := \frac{q}{\bar{k}} \cdot \theta + \frac{\bar{k} - q}{\bar{k}} \cdot 2\varepsilon > 0.$$

Since we have ruled out the case where the probability that the game proceeds to period $t = 2$ is the same no matter whether seller i sends message h or message l , we distinguish two cases. In either case, the message sent by seller i must have an impact on the outcome of the Devil's Menu in period $t = 1$ with positive probability.

The first case is that the probability that the game proceeds to period $t = 2$ is higher if seller i sends message h than if he sends message l . There are two possibilities: first, $N_k \in O$ if seller i sends message h and $N_k \in T$ if seller i sends message l ; second, $N_k \in T$ if seller i sends message h and $N_k \in C$ if seller i sends message l . In either case, let p_h (p_l) denote the probability that the game moves to period $t = 2$ if seller i sends message h (l), with $0 \leq p_l < p_h \leq 1$. The expected gain for seller i from sending message l instead of message h in the first case is equal to

$$(1 - p_l) \cdot \left(\frac{q - c}{t} \cdot \theta + \frac{t - q + c}{t} \cdot 2\varepsilon \right) + p_l \cdot \bar{u}_i - (1 - p_h) \cdot \varepsilon - p_h \cdot \bar{u}_i, \quad (14)$$

while in the second case, the expected gain is equal to

$$(1 - p_l) \cdot \theta + p_l \cdot \bar{u}_i - (1 - p_h) \cdot \left(\frac{q - c}{t} \cdot (V + \epsilon) + \frac{t - q + c}{t} \cdot \epsilon \right) - p_h \cdot \bar{u}_i. \quad (15)$$

Because ϵ can be chosen to be arbitrarily small and that our arguments below are not tight, we let $\epsilon = 0$. On the one hand, since $p_h \leq 1$, we can easily see that

$$(1 - p_l) \cdot \frac{q - c}{t} \cdot \theta > (p_h - p_l) \cdot \delta \cdot \frac{q}{\bar{k}} \cdot \theta$$

is implied by

$$\frac{q - c}{t} > \delta \cdot \frac{q}{\bar{k}} = \frac{\bar{k} \cdot (q - c)}{q \cdot (\bar{k} - o - c)}.$$

The latter inequality holds provided that

$$\delta < \frac{\bar{k}}{q \cdot (\bar{k} - 1 + q)}.$$

Indeed, it suffices to note that

$$\frac{\bar{k}}{q \cdot (\bar{k} - 1 + q)} \leq \frac{\bar{k} \cdot (q - c)}{q \cdot (\bar{k} - c)} \leq \frac{\bar{k} \cdot (q - c)}{q \cdot (\bar{k} - o - c)},$$

where the first inequality holds since it must be the case that $c \leq q - 1$ and $(q - c)/(\bar{k} - c)$ is a decreasing function of c (also recall that $q < \bar{k}$). In turn, the second inequality holds since $o \geq 0$. Thus expression (14) is positive if $\epsilon > 0$ is sufficiently small. On the other hand,

$$\begin{aligned} & (1 - p_l) \cdot \theta + p_l \cdot \bar{u}_i - (1 - p_h) \cdot \theta - p_h \cdot \bar{u}_i \\ &= \theta \cdot (p_h - p_l) - \bar{u}_i \cdot (p_h - p_l) = (\theta - \bar{u}_i)(p_h - p_l) > 0. \end{aligned}$$

Thus expression (15) is also positive if $\epsilon > 0$ is sufficiently small.

The second case is that the probability that the game proceeds to period $t = 2$ is higher if seller i sends message l than if he sends message h . However, by construction of a Devil's Menu, this cannot occur, as all else being equal, the probability that the buyer will attain her goal in a given execution of a Devil's menu is (weakly) larger if the number of sellers sending message l in a given group increases.

That is, we have proved that truth-telling is a necessary condition for an equilibrium for all sellers $i \neq j$. If seller j also sends message l in period $t = 1$, it must be the case that all groups have the same chance to be chosen in such a period. In particular, seller j expects utility

$$\frac{q}{\bar{k}} \cdot \theta + \frac{\bar{k} - q}{\bar{k}} \cdot 2\epsilon.$$

Now, if seller j deviates and sends message h in period $t = 1$, his group will not be selected in period $t = 1$. Yet, the buyer will be able to attain her goal in this period (and she will be certain of that with probability one), and thus the game will not proceed to period $t = 2$. Accordingly, upon deviation, seller j expects utility ε , which is lower than the utility he expects by reporting his type truthfully in period $t = 1$, since

$$\frac{q}{\bar{k}} \cdot \theta + \frac{\bar{k} - q}{\bar{k}} \cdot 2\varepsilon \geq \varepsilon.$$

Hence, seller j will report his type in period $t = 1$. This completes the proof of the proposition.

□