Markets and Transaction Costs*

Simon Jantschgi[†], Heinrich H. Nax[‡], Bary S. R. Pradelski[§] and Marek Pycia[¶]

First draft: April 2022; Latest version: June 2023

Abstract

Transaction costs are ubiquitous in markets. We show that their presence can fundamentally alter incentives and welfare in markets in which the price equates supply and demand. We categorize transaction costs into two types. Asymptotically uninfluenceable transaction costs—such as fixed and price fees—preserve the key asymptotic properties of markets without transaction costs, namely strategyproofness, efficiency, and robustness to misspecified beliefs and to aggregate uncertainty. In contrast, influenceable transaction costs—such as spread fees—lead to complex strategic behavior (which we call price guessing) and may result in severe market failure. In our analysis of optimal design we focus on transaction costs that are fees collected by a platform as revenue. We show how optimal design depends on the traders' beliefs. In particular, with common prior beliefs, any asymptotically uninfluenceable fee schedule can be scaled to be optimal, while purely influenceable fee schedules lead to zero revenue. Our insights extend beyond markets equalizing supply and demand.

Keywords: Transaction Costs, Markets, Incentives, Efficiency, Robustness, Market Design.

1 Introduction

Transaction costs are ubiquitous in markets. Markets are organized in a variety of ways. In economists' thinking about them, central has been the idea that buyers offer prices (bids) and sellers ask for prices (asks) and that the trade occurs at a price in between these buyers' and sellers' prices, a price that equalizes supply and demand. This idea was formalized as a Double Auction (DA) by

*We are grateful for comments and suggestions by Paul Klemperer, Jacob Leshno, Meg Meyer, Jonathan Newton, Sven Seuken, Eran Shmaya, Philipp Strack, Alex Teytelboym and participants of the 7th Workshop on Stochastic Methods in Game Theory (Erice, Italy), the Student Research Seminar in Microeconomic Theory (U. of Oxford), the 8th Swiss Theory Day (U. of Bern), the 12th Conference on Economic Design (University of Padova), the 33rd International Conference on Game Theory (Stony Brook U.), the EEA Annual Conference 2022 (Bocconi U.), the TOM Seminar (PSE), the PhD Seminar on Combinatorics, Games and Optimization (LSE), the Parisian Seminar of Game Theory, and the EC'22 Program Committee.

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement No 866376 (MP), from the Swiss National Science Foundation (SNSF) under an Eccellenza Grant (SJ and HHN), and by the French National Research Agency (ANR) through grant ANR-19-CE48-0018-01 (SJ and BSRP).

[†]University of Oxford, United Kingdom; simon.jantschgi@economics.ox.ak.uk

[‡]University of Zurich, Switzerland; heinrich.nax@uzh.ch

[§]CNRS, Université Grenoble Alpes, France; bary.pradelski@cnrs.fr

[¶]University of Zurich, Switzerland; marek.pycia@econ.uzh.ch

Chatterjee and Samuelson (1983); Wilson (1985b); Rustichini et al. (1994) but the mechanism has been used for centuries as modeled as early as Walras (1874).[1]

The importance of transaction costs for the functioning of markets has been recognized at least since Demsetz (1968). Transaction costs may depend on the underlying trades in various ways. A *fixed fee* charged to a trader depends only on whether the trader participates in trade; examples include handling costs such as packaging and shipping and a transaction cost covering the market makers internal operations costs. A *price fee* is a percentage of the price; examples include stamp duties set by governments, Tobin taxes as levied in Sweden and Latin America, the 'buyer's premium' charged by art auction houses, and 'service fees' or 'final value fees' as charged by Airbnb, eBay, Uber and Lyft, etc. A *spread fee* is a percentage of the difference between a trader's bid or ask and the market clearing price (that is often unknown to the trader); examples include commissions charged by intermediaries such as car dealers, limit orders on stock markets, and markets where trader's pay their bid (e.g., Priceline.com). Transaction costs can sometimes be small (e.g., a stock market transaction fee), and sometimes substantive (e.g., a service charge from a matching platform). Conceptually, one may think of any difference between what a buyer pays and what a seller receives as a transaction cost.

Early models of Double Auctions abstracted away from informational issues and transaction costs. The subsequent rich literature integrated incomplete information (starting with Chatterjee and Samuelson (1983); Myerson and Satterthwaite (1983); Wilson (1985b); Rustichini et al. (1994)), while the impact of transaction costs was less explored. An impression one might obtain from the rare analysis taking transaction costs into account is that they have no substantive impact on strategic behavior.[2] This conclusion hinges on the focus of this literature on fixed fees and price fees. In this paper, we generalize the analysis of transaction costs: *How do general transaction costs affect traders' incentives? What are the resulting welfare properties of markets? How would auctioneers design the costs optimally?* We find that transaction costs can be classified into two categories. The conclusions gleamed from studying fixed and price fees hold for transaction costs belonging to one of these categories but not for the other. In particular, we show that the presence of such common types of transaction costs as spread fees might have substantive impact on strategic behavior and the resulting market performance.

Allowing general continuous and monotonic transactions costs, we consider a market in which the price equates supply and demand: a Double Auction.[3] In the absence of transaction costs, in large DAs the gains from misreporting have been shown to vanish and the resulting outcome to be efficient (cf. Wilson 1985b; Rustichini et al. 1994; Cripps and Swinkels 2006; Azevedo and Budish 2019). We characterize optimal strategic behavior and categorize transaction costs into two types, *asymptotically uninfluenceable transaction costs* that preserve the latter desirable properties—

[1] See Friedman and Rust (1993) for a survey of the DA in history, theory and practice.

[2] See Tatur (2005) and other papers we discuss below.

[3] Notably, stock exchanges, including the New York Stock Exchange, run opening auctions at the start of each trading day to equate supply and demand. Their mechanism closely resembles a Double Auction with transaction costs. During a trading day stock exchanges run quasi-continuous markets, which can be thought of as continuous open-bid Double Auctions.

asymptotic truthfulness and efficiency—and *influenceable transaction costs* that do not preserve them. We also analyze the robustness of our findings to market participants having misspecified beliefs and aggregate uncertainty.

A transaction cost is *asymptotically uninfluenceable* if, conditional on a market participant trading in the market, the participant's impact on the transaction cost they pay vanishes as the market grows large; the transaction cost is *influenceable* if the cost depends on the trading participant's actions independently of the market size. Price fees are examples of asymptotically uninfluenceable transaction costs as, in the limit, the market participants impact on the fee vanishes (and, relatedly, all participants who trade pay the same fee). Spread fees are examples of influenceable transaction costs as, for any market size, the spread and hence the fee paid depends on the trading participant's action. Not surprisingly, under asymptotically uninfluenceable transaction costs, the traders behave similarly to traders in markets with no transaction costs and they are approximately *truthful* in large markets. In contrast, influenceable transaction costs distort incentives fundamentally, and, asymptotically, lead to what we call *price-guessing* behavior whereby traders bid close to estimated market prices in order to try to minimize their transaction cost.

Asymptotically uninfluenceable transaction costs lead to some unavoidable welfare losses in finite markets that are due to strategic behavior and possible direct loss due to unprofitability of trades whose surplus is insufficient to cover the cost. Because truthfulness emerges in the limit, in large markets the outcomes are not much affected when the transaction costs are small; and the same obtains even when agents have misspecified beliefs. In contrast, in large market equilibria, influenceable transaction costs lead to no loss due to strategic behavior, but again may lead to a direct loss as described above. However, even slight belief misspecification often leads to substantive market failure. The risk of market failure occurs for all influenceable transaction costs, and the degree of inefficiency does not vanish for small transaction cost.

We apply our insights on strategic behavior in the presence of transaction costs to analyze the optimal design of fees collected by a platform; such fees are one of the canonical instances of transaction costs. Considering an objective function that incorporates traders' welfare and platform's revenue, we show how optimal design depends on the traders' beliefs. With common prior beliefs, any asymptotically uninfluenceable fee schedule can be scaled to be optimal, while purely influenceable fee schedules lead to zero revenue. For some heterogeneous prior beliefs, purely influenceable fee schedules can strictly outperform any asymptotically uninfluenceable fee schedule.

Finally, we discuss how our insights remain valid in any market organization in which the participants believe that they have no influence on market prices.

Related literature

The idea that trade occurs at the price that equates revealed supply and demand goes back many centuries and is at the core of economics until today (cf. Smith 1776; Hosseini 1995). In finite markets, the DA is the standard mechanism to compute the allocation and the market price. Strategic behavior has consequently been widely studied. Prominently, Myerson and Satterthwaite (1983) showed that for finite markets with incomplete information, there generally exists no budget-balanced, incentive-compatible, and individually rational mechanism that is Pareto efficient.[4]

Without transaction costs, it has been shown that the incentive to misrepresent becomes arbitrarily small in large markets (Roberts and Postlewaite 1976). For the DA-mechanism in finite markets, Wilson (1985b), Rustichini et al. (1994) and Cripps and Swinkels (2006) show that market participants have incentives to be increasingly truthful, which results in asymptotic efficiency; any given participant's influence on demand or supply—and therefore the market clearing price—vanishes. Rustichini et al. (1994) establish this key insight for the DA mechanism with independent private values (cf. Satterthwaite and Williams 1989b). Reny and Perry (2006) extend those findings to continuum markets (à la Aumann (1964)) in which supply and demand satisfy certain continuity and monotonicity assumptions.[5]

In the presence of transaction costs we know much less about strategic behavior in double auctions. One notable exception is the treatment of constant transaction costs in Tatur (2005). Chen and Zhang (2020) study revenues in linear equilibria of DAs with transaction costs; they allow transaction costs to depend on the size of individual trade but not on price, bid-ask spread, nor other parameters of the market. Marra (2021) studies market entry in DAs with fixed transaction costs. Noussair et al. (1998) provides experimental evidence that fixed transaction costs lead to efficiency loss. Fixed transaction costs have also been the focus in the finance literature on limit order books (Colliard and Foucault 2012, Foucault et al. 2013, Malinova and Park 2015).[6] Where this literature focuses on specific (fixed) transaction costs, we look at transaction costs more generally and our classification has no counterpart in the literature. Our general incentive, efficiency, and robustness results are also new.

As we shall show a key difficulty in analyzing DAs with transaction costs is that the continuity and monotonicity assumptions as in Reny and Perry (2006) fail for many cost structures including the widely used spread fees. Even the definition of DAs in continuum markets that fail the continuity and monotonicity assumptions was missing till recently. We provide such a definition in our companion paper Jantschgi et al. (2022). This definition does not require any regularity assumptions, is consistent with earlier definitions where they exist, and it allows us to use the same definition for finite and infinite markets.

As we allow for general belief structures, our analysis contributes to the literature on market behavior under belief misspecifications, a topic of interest since Ledyard (1978) and Wilson (1985a).[7] The main thrust of this literature (see review by Spiegler 2020) is that robustness to misspecification

[4]The impossibility hinges on the quasilinearity of the preferences, which we also assume; see Garratt and Pycia (2016).

[5]See also Fudenberg et al. (2007) who generalize the convergence results of Rustichini et al. (1994). Earlier work on equilibrium existence in DAs includes Chatterjee and Samuelson (1983), Leininger et al. (1989), Satterthwaite and Williams (1989a), Williams (1991), and Cripps and Swinkels (2006). Of interest is also Jackson and Swinkels (2005) who study equilibrium existence in a broad class of private value auctions that includes DAs, and Azevedo and Budish (2019) who show that DAs are strategy-proof in the large, that is, truthfulness is approximately optimal against regular action distributions in large finite markets.

[6]See also Shi et al. (2013) for a numerical model of marketplace competition with transaction costs.

[7]See, for instance, Bergemann and Morris (2005); Chu and Shen (2006); Chassang (2013); Carroll (2015); Wolitzky (2016); Madarász and Prat (2017).

requires the mechanism to be simple (Li, 2017; Börgers and Li, 2019; Pycia and Troyan, 2023); other than our design analysis (which shows that simple fee structures are as good as the more complex ones), we focus on the complementary problem of how traders with misspecified beliefs behave. In the context of Walrasian markets, the impact of heterogeneous, misspecified, beliefs has been analyzed, e.g., by Harrison and Kreps (1979) and Eyster and Piccione (2013).[8] While they studied Walrasian equilibria without transaction costs, we look at DAs and allow for transaction costs.

As discussed above transaction costs are ubiquitous. They might take the form of transportation costs, platform fees, or taxes. Various transaction taxes are perhaps the most studied form of transaction costs. Common taxes include per-unit and ad-valorem commodity taxes, which in our formulation correspond to constant fees and price fees. Such taxes have been extensively studied in the context of perfectly competitive markets, cf., e.g., Ramsey (1927); Diamond and Mirrlees (1971a,b); Mirrlees (1986). In our optimal design analysis we use the welfare measure pioneered by Harberger (1962, 1964): the deadweight loss due to taxation (or in our case more general transaction costs) known as the 'Harberger triangle'.[9]

The issue taken up by our analysis of strategic behavior and optimal fee design have also been independently explored by Kang and Muir (2022) in the context of optimal contract design by a hybrid platform; they show that when facing competition in the upstream market, the platform wants to respond by exclusive dealing and "killer" acquisitions (as in Cunningham et al. 2021). As we show that the welfare critically depends on traders' beliefs, our study might be seen as highlighting the importance of the recent literature on information design by platforms, cf. Jullien and Pavan (2019). Early platform studies simplified the behavior of platform users by focusing on entry decisions (Rochet and Tirole, 2003; Armstrong, 2006); a notable recent analysis in this spirit is, e.g., Galeotti and Moraga-González (2009), who study the impact of outside options on platform's revenues.

Finally, the literature on market microstructure also pays attention to transaction costs which are set to ensure that auctioneers make nonnegative profits, see e.g., Glosten and Milgrom (1985) who explain how differences between sale and buy prices may arise as a function of informational asymmetries even if the auctioneer makes zero profits.

2 Example

We consider a special case of our general model. Assume there is a continuum of traders on each side of the market. One of the main results of our paper is that there are two qualitatively different categories of transaction costs. In the example, we focus on two common transaction costs that are representative of these categories: price fees and spread fees. The exposition is parallel to the structure of the general results so that the reader can easily read it as both a preview and an illustration.

[8]See also, e.g., Heidhues et al. (2018) who study overconfidence in markets, and de Clippel and Rozen (2018) who study the misperception of tastes.

[9]The Harberger triangle is also the standard tool to assess the welfare loss of monopoly; see Harberger (1971); Hines Jr (1999) for surveys.

Model (cf. Section 3)

The market (cf. Section 3.1)**.** We consider a two-sided infinite market with a unit mass of buyers and sellers interested in buying or selling an indivisible good. *Types*, giving the *gross value* of the item to a trader i, are uniformly distributed with $t_i \in T = [1, 2]$. The *utility* of each trader is the sum of the gross value of the object (if they have it) and their money holdings, normalized such that a trader who does not trade has utility 0.

The mechanism (cf. Section 3.2)**.** Every trader i submits an *action* $a_i \in \mathbb{R}^{\geq 0}$ representing a buyer's *bid* and a seller's *ask*. Given all actions, the *double auction* selects subsets of buyers and sellers involved in trade at a unique *market price* P^*. The market price is set to balance supply and demand, which are the total mass of sellers and buyers, who, given their actions, weakly prefer trading over not trading at that price. Additionally, every trader involved in trade has to pay a *transaction cost*. In the example, we consider representative transaction costs, price and spread fees. A *price fee* is given by a fixed percentage $\phi \in [0, 1]$ of the market price and a *spread fee* is given by a fixed percentage $\phi \in [0, 1]$ of the spread between the action of a trader and the market price.

Beliefs and aggregate uncertainty (cf. Section 3.3)**.** We assume that traders know the market mechanism, but have incomplete information about the market environment, that is the distribution of gross values and the behavior of other traders. Both market sides may have incorrect and heterogeneous *beliefs*, and *aggregate uncertainty*. We work with traders' beliefs over actions. In an infinite market—as considered in the example—this simplifies to considering beliefs directly over the market price. Suppose that all buyers believe the market price to be $\beta \in [1, 2]$ and all sellers believe it to be $\sigma \in [1, 2]$. We say that beliefs have a *common prior*, if $\beta = \sigma$. Otherwise, we call them *heterogeneous prior beliefs*. Traders might be uncertain about the market price and believe that it is distributed according to a *Beta*-distribution over $[1, 2]$, with mean equal to β respectively σ.[10].

Key Concepts (cf. Section 4)

Truthfulness (cf. Section 4.1)**.** In a double auction without transaction costs bidding one's gross value is the only action that (1) never results in a loss, (2) dominates all less aggressive actions (that is, actions that are higher for the buyer and lower for the seller), and (3) is not dominated by any more aggressive action. With transaction costs, bidding one's gross value may no longer satisfy these properties. We define the *net value*, t_b^Φ of a buyer with gross value t_b as the largest action satisfying (1)-(3). In analogy, for a seller with gross value t_s, the net value t_s^Φ is the smallest action satisfying (1)-(3). With no transaction costs, the net value is the gross value, and motivated by this we say that bidding is *truthful* if the trader bids their net value. Consider price and spread fees. With price fees, for a buyer with gross value t_b, the net value is $t_b^\Phi = {t_b}/{(1+\phi)}$ and for a seller with gross value t_s, the net value is $t_b^\Phi = {t_s}/{(1-\phi)}$. With positive price fees, trading at the market price equal to gross value results in negative utility while trading at the price equal to net value results in the utility of 0. With spread fees, the net values are equal to the gross values, that is, $t_b^\Phi = t_b$ and $t_s^\Phi = t_s$. A trader is indifferent between trading and not trading if the market price is equal to their gross value.

[10]As other simplifying assumptions, the *Beta*-distribution is just assumed for concreteness in the example

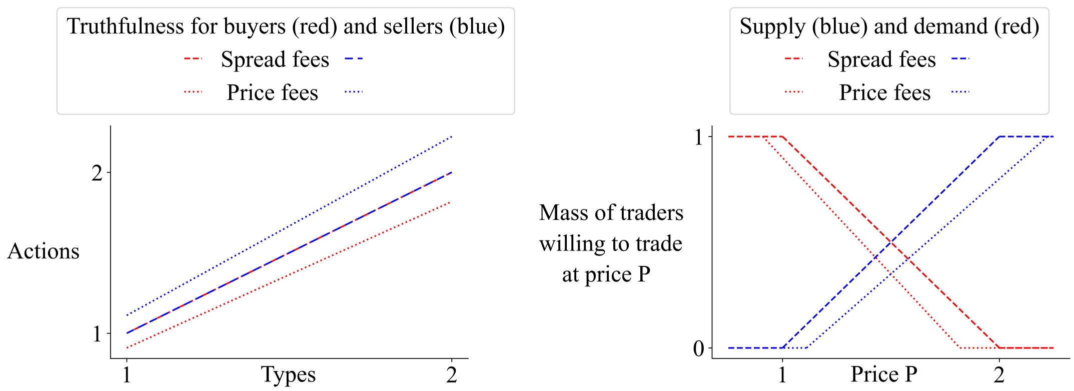

Figure 1: *Left.* Truthful strategy profiles for a 10% price and any spread fee. *Right.* Demand and supply functions, if traders act truthfully, again with a 10% price and any spread fee.

Predictability of trade (cf. Section 4.2)**.** Without uncertainty, a buyer believes to trade, if their bid is above the market price. Similarly, a seller believes to trade, if their ask is below the market price. If their action is equal to the market price they believe to be involved in *tie-breaking* and trade with some probability. In the presence of uncertainty, the probability to be involved in trade is a continuous function of a trader's action. Decreasing the *aggressiveness* of one's action, that is the distance to truthfulness, increases the probability of being involved in trade.

Profitability of trade (cf. Section 4.3)**.** In an infinite market, as a trader cannot influence the market price, a price fee is independent of a trader's action. In contrast, a spread fee is directly influenced by the action of a trader and decreases, if a trader reports an action that is closer to the market price. As a general analysis shows, a trader's influence on their transaction cost or its lack plays a crucial role in determining their optimal strategy.

Trader's behavior (cf. Section 5)

Optimal behavior maximizes the expected utility of a trader given their beliefs by finding the right amount of aggressiveness to balance probability of trade with profitability of trade. In the absence of tie-breaking, optimal strategies exist. With tie-breaking, existence of optimal strategies depends on the nature of the transaction cost.

Truthfulness is optimal for price fees (cf. Section 5.1)**.** As a trader cannot influence their payment, in order to maximize expected utility, it is optimal to maximize trading probability as long as the involvement in trade is individually rational. This is achieved by a trader truthfully bidding their net value. Note that truthfulness is independent of beliefs and uncertainty.

Price-guessing is optimal for spread fees (cf. Section 5.2)**.** In the absence of uncertainty and tie-breaking, it is optimal to bid the market price, if this is individually rational given a trader's gross value. We call this behavior *price-guessing*. If there is uncertainty or tie-breaking, the trade-off between decreasing the spread fee and increasing the probability of trade is non-trivial. If the uncertainty is sufficiently small, the incentive on the former outweighs the latter and it is optimal to bid close to the market price. Note that price-guessing crucially depends on beliefs.

Market performance and design (cf. Section 6)

Suppose that the fees are collected by a market platform (as opposed to, for example, transportation costs). Then a social planner evaluates market outcomes using standard performance metrics. If the social planner can design the fee structure, what is the optimal choice?

Market Performance (cf. Section 6.1)**.** The *trading volume* Q^* is the mass of active traders. The *trading excess* Ex measures for the two market sides the difference in mass of traders, who are willing to trade at the market price. The *trader's welfare* W is the utility of all active traders. The *platform's revenue* R is the total amount of collected fees. Their sum is called the *gains of trade*. We distinguish between *realized*, *net*, and *gross gains of trade*, write G^{real}, G^{net}, and G^{gross}, depending on whether trader's use some action profile, or report their net or gross values. The *total loss* is the difference $L = G^{gross} - G^{real}$, which measures how much gains of trade are lost due to transaction cost considerations and strategic behavior. We split it up into $L = L_d + L_s$, where $L_d = G^{gross} - G^{net}$ is the *direct loss* due to transaction cost constraints and $L_s = G^{net} - G^{real}$ is the *strategy-induced loss*. G^{gross} can then be decomposed into welfare, revenue, and loss: $G^{gross} = W + R + L_d + L_s$.

Optimal design (cf. Section 6.2)**.** Suppose that the social planner setting the fee schedule is *revenue-maximizing* and traders' beliefs are independent of the fee schedule. In Section 6, we will consider general objective functions that are induced by social planner's that care about the trader's welfare as well and traders' beliefs that depend on the fee schedule. We will show that price fees can be optimally scaled independent of traders' beliefs. For spread fees, optimal design depends on the trader's beliefs. For common prior beliefs, any spread fee leads to zero revenue due to price-guessing. For some heterogeneous prior beliefs, spread fees can strictly outperform price fees, while for others, they lead to complete market failure.

Optimal price fees are independent of beliefs (cf. Section 6.2.1)**.** Independent of beliefs and uncertainty, truthfulness is optimal. The market price does not depend on the symmetric fee parameter ϕ and is equal to $P^* = {}^3\!/_2$. The trading volume $Q^* = {}^{(1-3\phi)}\!/_2$ decreases linearly in ϕ with maximal trading volume without price fees equal to ${}^1\!/_2$ and full market failure occurring at $\phi = {}^1\!/_3$. Trading excess is equal to 0, so no tie-breaking is needed. The gross gains of trade are $G^{gross} = {}^1\!/_4$ and the realized gains of trade are equal to the net gains of trade $G^{net} = {}^{(1-9\phi^2)}\!/_4$. There is no strategy-induced loss, as traders report truthfully. The direct loss is equal to ${}^{9\phi^2}\!/_4$, which is strictly increasing in the fee parameter. Welfare is equal to $W = {}^{(1-6\phi+9\phi^2)}\!/_4$ and revenue is equal to $R = {}^{(3\phi-9\phi^2)}\!/_2$. Revenue is maximized at $\phi = {}^1\!/_6$, where individuals' fee payments and market volume are balanced. At this point, 25% of the gross gains of trade are lost, 50% are revenue and 25% remain as welfare to the traders. The second column of Figure 2 shows the decomposition of the gross gains of trade as a function of the fee parameter ϕ.

Optimal spread fees depend on beliefs (cf. Section 6.2.2)**.** Optimal behavior in the presence of spread fees depends on beliefs and uncertainty. Without uncertainty, price-guessing is optimal. With uncertainty, traders might deviate from price-guessing: Traders with profitable gross values are less aggressive, while traders with gross value close to the true market price might submit actions that are more aggressive. We show that depending on the beliefs β and σ about the market price, market outcomes range from full efficiency (with different decomposition of the gross gains of trade into welfare and revenue) to complete market failure. Note that inefficiency is only due to strategic behavior, as spread fees do not lead to a direct loss. Furthermore, depending on the beliefs, uncertainty can either improve or worsen the market outcome, both from traders and the market maker's perspective. To illustrate the range of possibilities, we analyze five different belief scenarios:

1. **Calibrated beliefs ($\beta = \sigma = 3/2$).** The market is fully efficient. There is no revenue, as there is no bid-ask spread for traders involved in trade. Uncertainty leads to a strategy-induced loss and some revenue.
2. **Homogeneous bias ($\beta = \sigma \neq 1.5$).** The market is not fully efficient. The strategy-induced loss is increasing in the distance between $\beta = \sigma$ and $3/2$. Similar to calibrated beliefs, there is no revenue. Uncertainty diminishes the strategy-induced loss and leads to positive revenue.
3. **Conservative bias ($\beta \geq 1.5 \geq \sigma$).** The market is fully efficient. The revenue decrease, if traders act more aggressive, and β and σ approach $3/2$. Uncertainty decreases the revenue and adds a strategy-induced loss.
4. **Aggressive bias ($\sigma \geq 1.5 \geq \beta$).** Complete market failure occurs. There is no trade, leading to zero revenue and surplus. Uncertainty lessens this effect, as traders are less aggressive, leading to trade, and hence some revenue and surplus.
5. **Mixed bias ($1.5 \geq \beta \geq \sigma$).** The market is not fully efficient. The loss is increasing in σ, more aggressive price-guessing by sellers leads to an efficiency loss. The revenue depends on the spread $\beta - \sigma$ and is generated entirely by buyers. Uncertainty leads to greater revenue and less strategy-induced loss.

The third and fourth column of Figure 2 show the decomposition of the true gains of trade as a function of the fee parameter ϕ for examples of the five belief scenarios with or without aggregate uncertainty. The optimal design for a revenue-maximizing social planner crucially depends on beliefs and uncertainty. First, consider the absence of aggregate uncertainty. If traders have homogeneous prior beliefs, so either calibrated beliefs or a homogeneous bias, revenue is zero regardless of the fee percentage. In that case, it is optimal to not charge any spread fee and avoid price-guessing, which would lead to the fully efficient market. If the beliefs are such that there is a spread, e.g., a conservative or a mixed bias, it is optimal for revenue maximization to charge a 100% spread fee, as price-guessing does not depend on the fee parameter. In the presence of aggregate uncertainty, the optimal fee percentage is given via a non-trivial optimization problem that can be solved analytically.

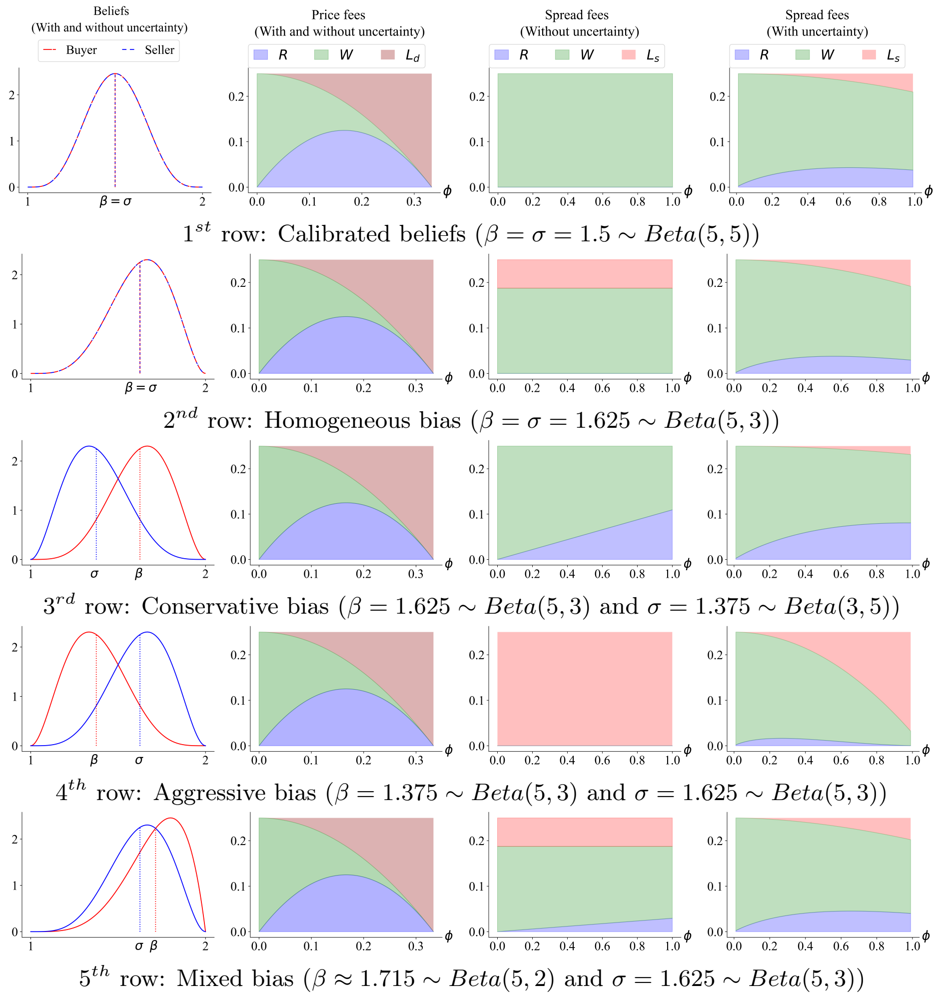

Figure 2: Decomposition of the gross gains of trade $G^{gross} = 0.25$ of an infinite uniform market into revenue R (blue), welfare W (green), direct loss L_d (dark-red) and strategy-induced loss L_s (light-red) as a function of price (2^{nd} column, independent of uncertainty) or spread fees ϕ (3^{rd} column without uncertainty and 4^{th} column with uncertainty), if traders best respond to their beliefs. The first column in each row shows the beliefs.

3 The model

With the developed intuition, we now turn to formally define our model to then be able to state our general results.

3.1 The market

We study a market in which traders play one of two roles: sellers sell and buyers buy a commodity. $\mathcal{B}$ denotes the set of buyers and $\mathcal{S}$ denotes the set of sellers. Each seller s has one unit to sell and each buyer b has single-unit demand. We allow both the finite case, with m buyers $\mathcal{B} = \{1, 2, ..., m\}$ and n sellers $\mathcal{S} = \{1, 2, ..., n\}$, and the infinite case, with $\mathcal{B} \subset \mathbb{R}$ and $\mathcal{S} \subset \mathbb{R}$ being two compact intervals.

Denote by μ_B and μ_S the counting measure (in the finite case) or the Lebesgue measure (in the infinite case) on the sets $\mathcal{B}$ and $\mathcal{S}$. Let $R = \frac{\mu_S(\mathcal{S})}{\mu_B(\mathcal{B})}$.

We focus on large finite and on infinite markets. We say that a property $\mathcal{P}$ holds *in sufficiently large finite markets*, if there exist $m, n \geq 1$ such that $\mathcal{P}$ holds in any finite market with at least m buyers and n sellers. If the property also holds in infinite markets, we say that it holds *in sufficiently large markets*.

Every trader $i \in \mathcal{B} \cup \mathcal{S}$ has a *type* $t_i \in T = [\underline{t}, \overline{t}] \subset \mathbb{R}^{\geq 0}$ that gives a trader's *gross value*, that is, the trader's valuation or reservation price for the item. We assume that the distribution from which types are drawn are absolutely continuous with probability densities f_B^t and f_S^t that are continuous and strictly positive on their support T, which we call the *type space*. Let (F_B^t, F_S^t) be the corresponding pairs of cumulative distribution functions of types. In finite markets, we assume that traders' types are independent random variables that are identically distributed according to (f_B^t, f_S^t) for each of the two market sides.[11] Given the random variables $t_b^1, ..., t_b^m$ and $t_s^1, ..., t_s^n$, we consider the *random empirical measures* on the sets of types $\mu_B^t = \sum_{j=0}^m \delta_{t_b^j}$ and $\mu_S^t = \sum_{k=0}^n \delta_{t_s^k}$. Letting n and m tend to infinity, normalized versions of μ_B^t and μ_S^t converge uniformly to measures with densities f_B^t and f_S^t; for details see Vapnik and Chervonenkis (1971). In an infinite market, we scale these measures by $\mu_B(\mathcal{B})$ and $\mu_S(\mathcal{S})$ to achieve the market ratio $R = \mu_S(\mathcal{S})/\mu_B(\mathcal{B})$ and we denote these measures again by μ_B^t and μ_S^t. Given realizations of types in finite markets and distributions of types in infinite markets, let $t_B : \mathcal{B} \to T$ and $t_S : \mathcal{S} \to T$ denote the functions assigning each trader their type.[12] The type distributions μ_B^t and μ_S^t are then the push-forward measures of μ_B and μ_S via the functions t_B and t_S, i.e., $\mu_B^t(\cdot) = \mu_B(t_B^{-1}(\cdot))$ and $\mu_S^t(\cdot) = \mu_S(t_S^{-1}(\cdot))$. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the probability space describing the randomness of sampling type distributions. Denote by $\mathbb{E}[\cdot]$ the *expectation* with respect to the probability measure $\mathbb{P}$. We write $t = (t_i, t_{\text{-}i})$, where t_i is trader i's type and $t_{\text{-}i}$ is the type distribution of all traders excluding trader i. In finite markets, t is obtained by adding a point mass at t_i to $t_{\text{-}i}$. In infinite markets, single traders do not change the type profile.

Every trader i submits an *action* $a_i \in \mathbb{R}^{\geq 0}$ representing a buyer's *bid* and a seller's *ask*. Denote by $a_B : \mathcal{B} \to A_B$ with $a_B(b) = a_b$ and by $a_S : \mathcal{S} \to A_S$ with $a_S(s) = a_s$ Borel-functions that assign an action to each trader. Let the *action distributions* μ_B^a and μ_S^a be two induced σ-additive and finite measures on $\mathbb{R}^{\geq 0}$ with support in the *action spaces* $A_B = [\underline{a}_B, \overline{a}_B]$ and $A_S = [\underline{a}_S, \overline{a}_S]$. That is, $\mu_B^a(\cdot) = \mu_B(a_B^{-1}(\cdot))$ and $\mu_S^a(\cdot) = \mu_S(a_S^{-1}(\cdot))$. Write $a = (a_i, a_{\text{-}i})$, where a_i is trader i's action and $a_{\text{-}i}$ is the action distribution of all traders excluding trader i. In finite markets a is obtained by adding a point mass to $a_{\text{-}i}$. In infinite markets, single traders do not influence the action profile. We will sometimes consider *strategies* $a_i : T \to A_i$, where $a_i(t_i)$ specifies the action given i's type. Given type distributions t, strategies of traders induce action distributions a as the push-forward measure of the type distributions.

We compare actions with respect to their *aggressiveness*, which refers to the amount of misrepresentation of their type: A buyer's bid a_b^1 is *(strictly) less aggressive* than a_b^2, write $\underset{(\succ)}{\succcurlyeq}$, if $a_b^1 \underset{(>)}{\geq} a_b^2$

[11] This is a common assumption in the literature, cf. Rustichini et al. (1994); Azevedo and Budish (2019).

[12] Given the assumptions on f_B^t and f_S^t, t_B and t_S are Borel functions.

and a seller's offer a_s^1 is *(strictly) less aggressive* than a_s^2, write $\underset{(\succ)}{\succsim}$, if $a_s^1 \underset{(<)}{\leq} a_s^2$.

The *utility* of each trader is the sum of the gross value of the object (if they have it) and their money holdings, normalized such that a trader who does not trade has utility 0. A buyer b involved in trade makes a payment, $P_b(a_b, a_{\text{-}b})$, in order to obtain an item and their resulting utility is $u_b(t_b, a_b, a_{\text{-}b}) = t_b - P_b(a_b, a_{\text{-}b})$. Similarly, a seller s involved in trade receives a payment, $P_s(a_s, a_{\text{-}s})$, for their item and their utility is $u_s(t_s, a_s, a_{\text{-}s}) = P_s(a_s, a_{\text{-}s}) - t_s$.

3.2 Double auction with transaction costs

Demand $D(P)$ and *supply* $S(P)$ at a price $P \geq 0$ are defined as $D(P) = \mu_B(\mathcal{B}_{\geq}(P))$ and $S(P) = \mu_S(\mathcal{S}_{\leq}(P))$, where $\mathcal{B}_{\geq}(P) = \{b \in \mathcal{B} : a_b \geq P\}$ and $\mathcal{S}_{\leq}(P) = \{s \in \mathcal{S} : a_s \leq P\}$. $\mathcal{B}_{>}(P)$, $\mathcal{B}_{=}(P)$, $\mathcal{S}_{<}(P)$, and $\mathcal{S}_{=}(P)$ are defined analogously.

A *double auction with transaction costs* maps an action profile a to a *market outcome*:

- A *market price* $P^*(a)$. Given a pricing parameter $k \in (0, 1)$, the market price is

  $$P^*(a) = k \cdot \min \mathcal{P}^{MC}(a) + (1 - k) \cdot \max \mathcal{P}^{MC}(a),$$

  where $\mathcal{P}^{MC}(a)$ is the set of *market clearing prices* equating revealed demand and supply. This pricing rule corresponds to the k-double auction introduced by Wilson (1985b); Rustichini et al. (1994).

- An *allocation* $A^*(a) = \mathcal{B}^*(a) \cup \mathcal{S}^*(a)$ identifying subsets of traders $\mathcal{B}^*(a) \subset \mathcal{B}$ and $\mathcal{S}^*(a) \subset \mathcal{S}$ involved in trade. Given $P^*(a)$, the allocation is:

  $$\begin{array}{lll} \mathcal{B}^*(a) = \mathcal{B}_{\geq}(P^*(a)) & \text{and } \mathcal{S}^*(a) = \mathcal{S}_{\leq}(P^*(a)) & \text{if there is no excess at } P^*(a) \\ \mathcal{B}^*(a) = \mathcal{B}_{>}(P^*(a)) \cup \tilde{\mathcal{B}} & \text{and } \mathcal{S}^*(a) = \mathcal{S}_{\leq}(P^*(a)) & \text{if there is excess demand at } P^*(a) \\ \mathcal{B}^*(a) = \mathcal{B}_{\geq}(P^*(a)) & \text{and } \mathcal{S}^*(a) = \mathcal{S}_{<}(P^*(a)) \cup \tilde{\mathcal{S}} & \text{if there is excess supply at } P^*(a) \end{array}$$

  $\tilde{\mathcal{B}} \subset \mathcal{B}_{=}(P^*(a))$ and $\tilde{\mathcal{S}} \subset \mathcal{S}_{=}(P^*(a))$ are sets of bid and ask orders selected uniformly at random to ensure that the trade is balanced, that is $\mu_B(\mathcal{B}^*(a)) = \mu_S(\mathcal{S}^*(a))$.

- *Transaction costs* $\Phi(a) = \Phi_i(a)_{i \in \mathcal{B}^* \cup \mathcal{S}^*} = (\Phi_i(a), \Phi_{-i}(a))$ for all active traders.

For a more detailed account of this unified double auction mechanism without transaction costs, see our companion paper Jantschgi et al. (2022).

Whenever the dependence on the action profile is clear, we write P^*, $\mathcal{B}^*$ and $\mathcal{S}^*$. When focusing on a single trader with action a_i, we write, e.g., $P^*(a_i, a_{\text{-}i})$.

The *payments of traders* $i \in \mathcal{B}^* \cup \mathcal{S}^*$ are determined by the market price $P^*(a_i, a_{\text{-}i})$ and transaction cost $\Phi_i(a_i, a_{\text{-}i})$. The payment a buyer $b \in \mathcal{B}^*$ makes is $P_b = P^*(a_b, a_{\text{-}b}) + \Phi_b(a_b, a_{\text{-}b})$ and the payment a seller $s \in \mathcal{S}^*$ receives is $P_s = P^*(a_s, a_{\text{-}s}) - \Phi_s(a_s, a_{\text{-}s})$. We assume that the payments $P_i(a_i, a_{\text{-}i})$ are continuous and increasing in a_i. Hence, for a buyer bidding more aggressively leads to a lower payment and for a seller bidding more aggressively leads to a higher payment. In Appendix B.1, we prove that the function $a_i \mapsto P^*(a_i, a_{\text{-}i})$ is continuous and increasing in a_i. Therefore, a sufficient condition for the monotonicity and continuity of the payment is that the transaction cost is continuous and increasing. The payments of traders $i \notin \mathcal{B}^* \cup \mathcal{S}^*$ are normalized to 0; these traders do not participate in trade.

Commonly observed transaction cost structures result in payments that are continuous and increasing. Examples include *constant* fees, *price* fees, and *spread* fees. We say that a transaction cost Φ_i is a *constant fee* if $\Phi_i(a) = c_i$ for some constant $c_i \geq 0$, is a *price fee* if $\Phi_i(a) = \phi_i \cdot P^*(a)$ for some constant $\phi_i \in [0, 1]$, and is a *spread fee* if $\Phi_i(a) = \phi_i \cdot |P^*(a) - a_i|$ for some constant $\phi_i \in [0, 1]$.[13,14]

3.3 Beliefs

We assume that traders commonly know the market mechanism, but have incomplete information regarding the market environment. In general, traders may have heterogeneous priors and incorrect beliefs.

Trader i has beliefs about the number of traders, the distribution of their gross values, and their market behavior. Denote by $R_i = {}^{\mu_S(\mathcal{S}_i)}/_{\mu_B(\mathcal{B}_i)}$ the ratio of the number of sellers to buyers. It is customary in the literature to assume correct beliefs about the number of traders and their gross value distribution. In this common prior belief setting it is then standard to study symmetric equilibrium strategies, see Rustichini et al. (1994). In an equilibrium, the traders' beliefs over fundamentals then induce their beliefs over other traders' actions. In a more recent strand of work, e.g., by Azevedo and Budish (2019) on Strategy-proofness in the Large, best response behavior to arbitrary action distributions is studied, not only those induced by common prior beliefs and perfectly rational play. With some analytical assumptions, beliefs over actions incorporate the classical model of traders having beliefs about type distributions and strategies of other traders.[15] Results on best response behavior can therefore be translated to symmetric Bayesian Nash equilibria. We adopt this line of thought and work directly with beliefs over actions, as we will also study the influence of misspecified beliefs on market performance in Section 6.

The distribution of actions of other traders is assumed to be absolutely continuous with probability densities $f^a_{B,i}$ and $f^a_{S,i}$ that are continuous and strictly positive on their supports. Let $\underline{a}_{B,i} = \min\{a_b : f^a_{B,i}(a_b) > 0\}$, $\overline{a}_{B,i} = \max\{a_b : f^a_{B,i}(a_b) > 0\}$, $\underline{a}_{S,i} = \min\{a_s : f^a_{S,i}(a_s) > 0\}$, $\overline{a}_{S,i} = \max\{a_s : f^a_{S,i}(a_s) > 0\}$. We assume that $\overline{a}_{S,i} \geq \overline{a}_{B,i} > \underline{a}_{S,i} \geq \underline{a}_{B,i}$; that is, the action spaces intersect. We also assume that trader i believes that when being truthful (up to transaction cost considerations), traders on both market sides will submit both less and more aggressive actions with positive probability. We shall define in Section 4.1 trader i's net value t^Φ_i, giving the truthful action (up to transaction cost considerations), that then satisfies $t^\Phi_i \in (\underline{a}_{S,i}, \overline{a}_{B,i})$.

[13] If $\phi_i = 0$ or $c_i = 0$, the setting simplifies to the classical DA without transaction costs. Further, for spread fees, if $\phi_i = 1$ a trader's payment is equal to their bid/ask.

[14] An interesting case is the price fee for the seller, the expected payment received by the seller is increasing when the seller bids more aggressively despite the fee paid by the seller being also increasing.

[15] If trader i believes that types are distributed according to (F^t_B, F^t_S) and all traders use a symmetric strategy profile (a_B, a_S), where both strategies are strictly increasing C^1-functions, then actions are distributed according to $F^t_B(a_B^{-1}(\cdot))$ on $A_{B,i}$ and $F^t_S(a_S^{-1}(\cdot))$ on $A_{S,i}$.

In finite markets, we impose two additional assumptions. First, we assume that other traders' actions are independent random variables, identically distributed for each of the two market sides. Second, we assume that the supports of distribution of actions of other traders are convex, that is, $A_{B,i} = [\underline{a}_{B,i}, \overline{a}_{B,i}]$ and $A_{S,i} = [\underline{a}_{S,i}, \overline{a}_{S,i}]$. Let $(F^a_{B,i}, F^a_{S,i})$ be the pair of corresponding C^1 distribution functions. Realizations of these random variables induce random empirical action distributions $\mu^a_{B,i}$ and $\mu^a_{S,i}$.

In infinite markets, we allow trader i to believe any action distribution $\mu^a_{B,i}$ and $\mu^a_{S,i}$. One class of absolutely continuous action distributions is obtained by viewing infinite markets as the limit of finite markets. Letting n and m tend to infinity, the random empirical probability measures converge uniformly to measures with densities $f^a_{B,i}$ and $f^a_{S,i}$.[16] Scaling these measures by $\mu_B(\mathcal{B}_i)$ and $\mu_S(\mathcal{S}_i)$ results in deterministic beliefs about absolutely continuous action distributions in infinite markets.

Given the beliefs of trader i, let $(\Omega_{\text{-}i}, \mathcal{F}_{\text{-}i}, \mathbb{P}_{\text{-}i})$ be the probability space describing the randomness of the action distribution $a_{\text{-}i}$ and tie-breaking. Denote by $\mathbb{E}_{\text{-}i}[\cdot]$ the *expectation* with respect to the probability measure $\mathbb{P}_{\text{-}i}$.

Let the *belief system* $\mathfrak{B}$ be the collection of all traders beliefs. $\mathfrak{B}$ is thus a mapping from the set of traders $\mathcal{B} \cup \mathcal{S}$ into the space of beliefs. We say that a belief system $\mathfrak{B}$ has a *common prior*, if all traders' beliefs lead to the same critical value. The critical value of trader i, which will be formally introduced in Section 4.2, is the market price that results when all traders in an infinite market behave according to trader i's beliefs; thus also approximating the market price for large finite markets. A key example of a common prior belief system is that all traders have exactly the same beliefs (common prior). Moreover, we say that a common prior belief system is *calibrated*, if the traders' belief of the critical value coincides with the critical value resulting from the type distributions. If a belief system does not have a common prior, we say that it has a *heterogeneous prior*. For tractability, we assume that in infinite markets, traders with the same type and on the same market side have the same critical value (cf. Section 6).

To evaluate the robustness of our findings we allow that traders are uncertain about their beliefs in infinite markets. We give a detailed definition of *aggregate uncertainty* in Appendix B.4. In the main text, after each result, we will state qualitatively how they extend to markets with aggregate uncertainty. The formal results are again relegated to Appendix B.4.

4 Key Concepts

In this section we introduce three key concepts which will allow to analyse optimal behavior. First, we give a definition of how a trader can ensure to not be loss-making ex-post in the presence of transaction costs. The second concerns a trader's ability to estimate their trade probability. Third, we introduce the key distinction between influenceable and asymptotically uninfluenceable transaction costs and their relationship to the profitability of trade.

[16] See Vapnik and Chervonenkis (1971) for these convergence results.

4.1 Truthfulness and Net Values

Without transaction costs, if trader i bids their gross value ($a_i(t_i) = t_i$), they maximize the probability to be involved in trade, conditional on guaranteeing ex-post individual rationality. An action a_i is ex-post individually rational, if for all $a_{\text{-}i}$ it holds that $u_i(t_i, a_i, a_{\text{-}i}) \geq 0$. Such behavior is often called *truthful* because a trader reveals their type. Buyers prefer not to trade at market prices above their gross value, and sellers prefer not to trade at market prices below their gross value. Indeed, bidding gross values represents the maximal bids that constitute undominated actions for buyers, and similarly the minimal asks that constitute undominated actions for sellers. We say that an action a_i^1 dominates an action a_i^2, if for all $a_{\text{-}i}$ it holds that $u_i(t_i, a_i^1, a_{\text{-}i}) \geq u_i(t_i, a_i^2, a_{\text{-}i})$.

In the presence of transaction costs, actions may have to be more aggressive than gross values in order to guarantee ex-post individual rationality, and bidding gross values may be dominated. For some transaction costs, e.g., for constant and price fees, bidding one's gross value would result in negative utility when the market price is equal to the gross value. Taking transaction costs into account, we define a buyer's *net value* (or *net value bid*), t_b^Φ, as the supremum of the set of undominated and ex-post individually rational bids. Similarly, we define a seller's *net value* (or *net value ask*), t_s^Φ, as the infimum of the set of undominated and ex-post individually rational asks. The net values always exist, though they might take an infinite value if the corresponding set is empty. The net values are actions (bids or asks) and thus the above-defined properties of actions apply to them.

The net values are finite, ex-post individually rational, undominated, strictly increasing, and continuous in the gross value for many transaction costs, including constant, price, and spread fees. The net values satisfy these properties for a large class of transaction costs. We say that transaction costs are *regular* if they only depend on the action of the trader and the market price (that is $\Phi_i(a_i, a_{\text{-}i}) = \Phi_i(a_i, P^*(a_i, a_{\text{-}i}))$, $P^* \mapsto P_i(a_i, P^*)$ is increasing,$a_i \mapsto P_i(a_i, a_i)$ is strictly increasing, and both are continuous. When transaction costs are regular, the sets of gross values that allow for profitable trade take the form $T_b^+ = \{t_b : \exists a_b : t_b - a_b - \Phi_b(a_b, a_b) > 0\}$ and $T_s^+ = \{t_s : \exists a_s : a_s - t_s - \Phi_s(a_s, a_s) > 0\}$. In Appendix A.1 we prove the following:

Proposition 1 (Regularity of net values)**.** *For regular transaction costs and $t_i \in T_i^+$, the net value actions are finite, undominated, ex-post individually rational, continuous and strictly increasing in the gross values. Furthermore, the net values are the unique solution of the equation $t_b - x - \Phi_b(x, x) = 0$ for a buyer and $x - t_s - \Phi_s(x, x) = 0$ for a seller.*

Motivated by this proposition we impose in the sequel the following global assumption:

Assumption—Regularity of net values. The net value actions are finite, ex-post individually rational, undominated, strictly increasing, and continuous in the gross value.

While the regularity of transaction costs is sufficient for the above assumption, our analysis does not rely on the regularity of transaction costs but only on the weaker assumption of the regularity of net values. Proposition 1 also tells us that the net value is the unique action, at which a trader is indifferent between trading and not trading, when the market price is equal to their action. For constant, price and spread fees, this characterization allows us to express the net values as follows:

Corollary 2 (Net values for constant, price, and spread fees)**.** *For constant fees, the net value shifts the gross value, that is,* $t_b^\Phi = \max(0, t_b - c_b)$ *and* $t_s^\Phi = t_s + c_s$*. Similarly, for price fees the net value scales the gross value, that is,* $t_b^\Phi = {}^{t_b}/_{1+\phi_b}$ *and* $t_s^\Phi = {}^{t_s}/_{1-\phi_s}$*. By contrast, for spread fees the gross value equals the net value, that is,* $t_b^\Phi = t_b$ *and* $t_s^\Phi = t_s$*.*

For proof see Appendix A.2. Finally, in the presence of transaction costs, we say that a trader is *truthful* if they bid their net value. Without transaction costs the net value is the gross value. We say that an action a_i is *(strictly) individually rational*, if it is (strictly) smaller than the net value for buyers and (strictly) greater than the net value for sellers.

4.2 Predictability of trade

Consider trader i's *probability of trading*, $\mathbb{P}_{\text{-}i}\left[i \in A^*(a_i, a_{\text{-}i})\right]$. In finite markets and infinite markets with aggregate uncertainty, the function $a_i \mapsto \mathbb{P}_{\text{-}i}\left[i \in A^*(a_i, a_{\text{-}i})\right]$ is continuous and can be expressed in terms of $F^a_{S,i}$ and $F^a_{B,i}$.[17] In infinite markets without aggregate uncertainty, trader i believes that the market price is deterministic and equal to the unique solution of the equation $\mu_S(\mathcal{S})F^a_{S,i}(\cdot) = \mu_B(\mathcal{B})(1 - F^a_{B,i}(\cdot))$. Call this solution the *critical value* P_i^∞.[18] The probability of trading is equal to 1, if trader i's action is less aggressive than P_i^∞. If their action is equal to P_i^∞ they believe to be involved in tie-breaking and trade with some probability between 0 and 1. If their action is more aggressive, trader i believes that they are not involved in trade.

The critical value is also of central importance for the study of trading probabilities in large finite markets. Given trader i's beliefs about others' behaviors, they can compute the market price with increasing accuracy as the market grows. With increasing numbers of traders on both market sides the variance of the realized market price decreases and it converges to the critical value. The probability of trading then converges to a step function at the critical value P_i^∞.

Proposition 3 (Predictability of trade)**.** *Consider trader i with action a_i. For every $\epsilon > 0$, in sufficiently large markets, the probability of trade for i is (1) bounded from below by $1 - \epsilon$ if a_i is strictly less aggressive than the critical value P_i^∞ and (2) bounded from above by ϵ if a_i is strictly more aggressive than the critical value P_i^∞.*

In the omitted case, when $a_i = P_i^\infty$, the trading probability in finite markets is determined by the action distributions and lies strictly between 0 and 1.[19] This results remains true, if trader's have sufficiently small uncertainty about the market, see Appendix B.4.

Proof Outline. Growing market size in finite markets is formalized with respect to a single parameter. Consider a sequence of strictly increasing market sizes $(m(l), n(l))_{l\in\mathbb{N}}$ with $m(l), n(l) = \Theta(l)$ and

[17]This is proven in Appendix B.1.2 (see Equations (42) and (43)).
[18]Existence and uniqueness are proven in Appendix A.3.
[19]E.g., for uniform action distributions and equally many buyers and sellers, the trading probability is independent of the market size and equal to $\frac{1}{2}$; we provide more details in the proof of Theorem 10 (2).

$|R - \frac{n(l)}{m(l)}| = \mathcal{O}(l^{-1})$ for $R \in (0, \infty)$.[20] A buyer b is involved in trade, if their action a_b is greater (or equal, if they win tie-breaking) than at least $m(l)$ actions of other traders, that is $\mathbb{P}_{\text{-}b}\left[b \in A^*(a_b, a_{\text{-}b})\right] = \mathbb{P}_{\text{-}b}\left[a_b \geq a_{\text{-}b}^{m(l)}\right]$. The probability that the action of any other buyer and seller is below a_b is $p_{a_b} = F_{B,b}(a_b)$ and $q_{a_b} = F_{S,b}(a_b)$. If $X_i^{p_{a_b}}$ and $X_j^{q_{a_b}}$ are Bernoulli random variables with parameters p_{a_b} and q_{a_b}, then the total number of traders with actions below a_b has the same distribution as the sum $S_l^{a_b} = \sum_{i=1}^{m(l)-1} X_i^{p_{a_b}} + \sum_{j=1}^{n(l)} X_j^{q_{a_b}}$. It follows that $\mathbb{P}_{\text{-}b}\left[b \in A^*(a_b, a_{\text{-}b})\right] = \mathbb{P}[S_l^{a_b} \geq m(l)] = 1 - \mathbb{P}[S_l^{a_b} \leq m(l) - 1]$. By the Berry-Esseen Theorem (Tyurin, 2012) an appropriately normalized version of $S_l^{a_b}$ converges in distribution to a standard normal random variable with CDF Φ. We show that there exists a sequence $(A_{a_b}(l))_{l \in \mathbb{N}} = \Theta(\sqrt{l})$ with $|\mathbb{P}[S_l^{a_b} \leq m(l) - 1] - \Phi(A_{a_b}(l))| \in \mathcal{O}(l^{-\frac{1}{2}})$. For $a_b \prec P_b^{\infty}$ we show for sufficiently large l that $A_{a_b}(l) < 0$, which yields that $A_{a_b}(l) \in \Theta(-\sqrt{l})$. Using a concentration inequality for a standard Gaussian random variable gives $\Phi(A_{a_b}(l)) \in \mathcal{O}(e^{-l})$. It therefore holds that $\mathbb{P}[S_l^{a_b} \leq m(l) - 1] = \mathcal{O}(l^{-\frac{1}{2}})$. The statement for $a_b \succ P_b^{\infty}$ and for sellers can be derived analogously. In infinite markets, the statement follows directly from the model. Proof details are relegated to Appendix A.4. □

We sometimes focus on *in-the-market* gross values that are gross values t_i such that $t_i^{\Phi} \prec P_i^{\infty}$. Traders with such gross values are able to submit individually rational actions that make them likely to be involved in trade when the market is sufficiently large. By contrast, for an *out-of-the-market* trader, that is, one with gross value $t_i^{\Phi} \succ P_i^{\infty}$, the probability of trade, when acting individually rationally, vanishes in large markets.

4.3 Profitability of trade

We now turn to the *expected utility conditional on trading*. Write $\mathbb{E}_{\text{-}i}\left[\cdot | i \in A^*\left(a_i, a_{\text{-}i}\right)\right]$ for the conditional expectation of trader i given their beliefs. Recall, that we assume that payments are monotone in the aggressiveness of one's action. Further, payments are composed of the market price and a transaction cost. For the former, it is known from Rustichini et al. (1994), that in large markets traders have vanishing influence on the market price. On the other hand, this is not necessarily the case for transaction costs. To this end, a classification of transaction costs into two broad classes turns out to be useful.

Definition (Asymptotically uninfluenceable vs. influenceable transaction costs)**.** Two actions a_i^1 and a_i^2, such that a_i^1 is less aggressive than a_i^2 and both are less aggressive than the critical value, that is $a_i^1 \prec a_i^2 \prec P_i^{\infty}$, lead to *asymptotically different transaction costs*, if there exists $\epsilon > 0$ such that in sufficiently large markets

$$\mathbb{E}_{\text{-}i}\left[\Phi_i(a_i^1, a_{\text{-}i}) | i \in A^*\left(a_i^1, a_{\text{-}i}\right)\right] - \mathbb{E}_{\text{-}i}\left[\Phi_i(a_i^2, a_{\text{-}i}) | i \in A^*\left(a_i^2, a_{\text{-}i}\right)\right] \geq \epsilon. \tag{1}$$

Otherwise, the two actions lead to *asymptotically equal transaction costs*. Transaction costs Φ_i are *influenceable* if every two such actions $a_i^1 \prec a_i^2 \prec P_i^{\infty}$ lead to asymptotically different transaction

[20] If there exists a parameter l, such that for every $l' \geq l$ Proposition 3 holds in markets with $m(l')$ buyers and $n(l')$ sellers, then the statement also holds in sufficiently large finite markets.

costs. Transaction costs Φ_i are *asymptotically uninfluenceable* if for every $\epsilon > 0$ in sufficiently large markets

$$\sup_{a_i^1 \prec a_i^2 \prec P_i^\infty} \mathbb{E}_{\text{-}i}\left[\Phi_i\left(a_i^1, a_{\text{-}i}\right) | i \in A^*\left(a_i^1, a_{\text{-}i}\right)\right] - \mathbb{E}_{\text{-}i}\left[\Phi_i\left(a_i^2, a_{\text{-}i}\right) | i \in A^*\left(a_i^2, a_{\text{-}i}\right)\right] \leq \epsilon. \tag{2}$$

In infinite markets, the definitions simplify, as there is no randomness due to sampling. Influenceability is then equivalent to the map $a_i \mapsto \Phi_i(a_i, a_{\text{-}i})$ being strictly increasing for buyers and strictly decreasing for sellers. Uninfluenceability is equivalent to the map $a_i \mapsto \Phi_i(a_i, a_{\text{-}i})$ being constant. For regular transaction costs that only depend on the trader's action and the market price, this implies that in infinite markets, asymptotically uninfluenceable transaction costs are a function of the market price, i.e. $\Phi_i(P^*)$, while for influenceable transaction costs, the function $a_i \mapsto \Phi_i(a_i, P^*)$ is again strictly monotone.[21]

An influenceable transaction cost might still include an asymptotically uninfluenceable part, e.g., the sum of a price and spread fee. We say that a regular influenceable transaction cost $\Phi_i(a_i, P^*)$ is *purely influenceable*, if it holds that $\Phi_i(P^*, P^*) = 0$. Note that for purely influenceable transaction costs the net value equals the gross value. Spread fees are an example of purely influenceable transaction costs. For regular transaction costs, it is possible to decompose any transaction cost into an asymptotically uninfluenceable and purely influenceable part.

Lemma 4 (Decomposition of regular transaction costs.)**.** *A regular influenceable transaction cost can be written as the sum of an asymptotically uninfluenceable transaction cost and a purely influenceable transaction cost.*

Proof details are relegated to Appendix A.5. Moreover, the two types are not mutually exclusive, as one can construct transaction costs that are asymptotically uninfluenceable in some price regions and influenceable at others. However, focusing on these two cases (rather than on hybrids) allows us to study the key strategic differences that in fact yield completely opposing behavior. In particular, the two canonical examples of transaction costs, price and spread fees, fall under the two definitions: Price fees are asymptotically uninfluenceable, and spread fees are purely influenceable.

5 Trader's behavior

Best responses maximize individual expected utility given beliefs. The maximization identifies the optimal amount of aggressiveness, balancing the opposing forces of increasing the probability of trade versus increasing the utility when trading.[22] Given trader i's beliefs and gross value t_i, an action a_i is an ϵ-*best response* if $\mathbb{E}_{\text{-}i}\left[u_i\left(t_i, a_i, a_{\text{-}i}\right)\right] \geq \sup_{a_i' \in \mathbb{R}} \mathbb{E}_{\text{-}i}\left[u_i\left(t_i, a_i', a_{\text{-}i}\right)\right] - \epsilon$. For $\epsilon = 0$ a_i is a *best response*.

[21] These monotonicity conditions can be equivalently stated for the payment function $a_i \mapsto P_i(a_i, P^*)$.

[22] A detailed analysis of this trade-off for price and spread fees in finite markets via first order conditions can be found in Appendix B.3.

The analysis of best responses includes the special case of *symmetric Bayesian Nash equilibria.* If all buyers use the same strictly increasing and continuous strategy a_B and all sellers use the same strictly increasing and continuous strategy a_S, call (a_B, a_S) a *symmetric strategy profile.* Given type distributions, the corresponding action distributions are given by $\mu_B^a(\cdot) = \mu_B\left(t_B^{-1}(a_B^{-1}(\cdot))\right)$ and $\mu_S^a(\cdot) = \mu_S\left(t_S^{-1}(a_S^{-1}(\cdot))\right)$. Assume that beliefs over action distributions originate from beliefs over gross value distributions and over the symmetric strategy profiles of the traders (a_B, a_S). If, for every trader and every gross value, the action specified by these strategies are ϵ-best responses, then the strategy profile constitutes a symmetric ϵ-Bayesian Nash equilibrium.[23]

Proposition 5 (Existence of best responses)**.** *Suppose the market is finite or infinite with tie-breaking occurring with probability zero. Then a best response exists for trader i.*

In infinite markets the no-tie-breaking assumption is necessary. In its absence, a best response might not exist for a trader i with $t_i \prec P_i^\infty$. This is the case, for example, when spread fees are charged. Under spread fees, it is not optimal to bid P_i^∞ (or more aggressive) due to the risk of loosing out on trading. But for any less aggressive bid, bidding slightly more aggressively would lead to a higher payoff. The results also extends to aggregate uncertainty, see Appendix B.4.

Proof Outline. We show that a best response is necessarily located in a compact action space. Given the continuity assumption of the payment, it follows that the expected utility is continuous in the action a_i and therefore attains a maximum by the Extreme Value Theorem. Proof details are relegated to Appendix A.6. □

The following theorem is a first indication that transaction costs have significant strategic consequences.

Theorem 6 (Asymptotically equal transaction costs)**.** *Let T^* be the set of gross values at which bidding the critical value is strictly individually rational. If trader i is best responding, then the expected transaction costs of any two types $t, t' \in T^*$ are asymptotically equal.*

For asymptotically uninfluenceable transaction costs this result holds by definition. For influenceable transaction costs, the result is non-trivial and will be useful in later analyses (see Section 5.2).

Proof Outline. Assume that two actions $a_i^1 \prec a_i^2 \prec P_i^\infty$ lead to asymptotically different transaction costs. We show that in sufficiently large markets, a trader can increase their expected utility, when switching from action a_i^1 to a_i^2, proving that a_i^1 is not a best response. Formally, as $a_i^1 \prec a_i^2 \prec P_i^\infty$, Proposition 3 yields that for every $\epsilon_1 > 0$, $\mathbb{P}_{\text{-}i}\left[i \in A^*(a_i^1, a_{\text{-}i})\right], \mathbb{P}_{\text{-}i}\left[i \in A^*(a_i^2, a_{\text{-}i})\right] \geq 1 - \epsilon_1$ in sufficiently large markets. The difference in trading probability between a_i^1 and a_i^2 is then upper bounded by ϵ_1. If ϵ_1 is sufficiently small, the loss in trading probability and possible influence on the market price is compensated by a decrease in expected transaction cost by at least some

[23]Therefore all of the results that we shall present in this paper about best responses directly apply to the study of symmetric Bayesian Nash equilibria.

$\epsilon_2 > 0$ because transactions are assumed to be asymptotically different. For sufficiently small ϵ_1, the difference in expected utility between actions a_i^1 and a_i^2 is negative, if the market is sufficiently large, proving that a_i^1 is indeed not a best response. Proof details are relegated to Appendix A.7. □

5.1 For asymptotically uninfluenceable transaction costs truthfulness is approximately optimal

Strategic misrepresentation is driven by the incentive to influence market price and transaction cost. Reporting truthfully maximizes one's trading probability, while remaining individually rational. In large markets, the influence on the market price is vanishing 'faster' than the influence on one's trading probability, which is what drives the asymptotic truthfulness result in the literature, see Rustichini et al. (1994). Therefore, if the influence on one's own transaction cost is also vanishing 'fast' enough, then it is close to optimal to maximize one's trading probability by reporting truthfully. This is the case for asymptotically uninfluenceable transaction costs, such as constant or price fees.

Theorem 7 (In large markets with asymptotically uninfluenceable transaction costs truthfulness is an approximate best response)**.** *If the transaction cost is asymptotically uninfluenceable and trader i's best response is uniformly bounded away from the critical value P_i^∞, then for every $\epsilon > 0$, in sufficiently large markets, truthfulness is an ϵ-best response.*

In infinite markets, the presence of aggregate uncertainty strenghtens this result, as truthfulness is then the unique best response, see Appendix B.4.

Proof Outline. Consider a best response a_i of trader i. If $a_i \prec t_i^\Phi$, then t_i^Φ is a best response by weak domination. Suppose now that $a_i \succ t_i^\Phi$. By assumption, there exists $\delta > 0$, such that in sufficiently large markets, (i) $a_i \prec P_i^\infty - \delta$ or (ii) $a_i \succ P_i^\infty + \delta$ holds. If (i) holds, then Proposition 3 implies that $\mathbb{P}_{\text{-}i}\left[i \in A^*(a_i, a_{\text{-}i})\right]$ converges to zero as the market gets large. Therefore for all $\epsilon > 0$ the expected utility of a_i is then upper bounded by ϵ, which also proves that that the net value is an ϵ-best response, because it leads to a non-negative expected utility. If (ii) holds, consider $\mathbb{E}_{\text{-}i}[u_i(t_i, a_i, a_{\text{-}i})] - \mathbb{E}_{\text{-}i}[u_i(t_i, t_i^\Phi, a_{\text{-}i})]$. We split the difference into two components and show that for every $\forall \epsilon > 0$ both components are less or equal than $\frac{\epsilon}{2}$ if the market is sufficiently large: (a) Difference in expected transaction costs and (b) Terms corresponding to a classical DA without transaction costs. To bound (a), we can use Proposition 3 and uninfluenceability. For (b), we will use that for a DA without transaction costs truthfulness is an ϵ-best response in sufficiently large markets, see Theorem 8.2 with price fees equal to zero. Proof details are relegated to Appendix A.8. □

Price fees. Fixing a specific transaction cost allows sharper results than Theorem 7. In particular, for a price fee, any best response can be explicitly shown to be close to truthful in large finite markets.

Theorem 8 (In large markets with price fees best responses are approximately truthful and truthfulness is an approximate best response)**.** *If the transaction cost is a price fee, then for every*

$\epsilon > 0$ *it holds that (1) in sufficiently large markets truthfulness is an ϵ-best response and (2) in sufficiently large finite markets all best responses are ϵ-truthful.*

In infinite markets, truthfulness is not unique as a best response. Every action $a_i \succ P_i^\infty$ that is individually rational is also a best response. Theorem 8 is robust to aggregate uncertainty in infinite markets, in which case truthfulness is also the unique best response, see Appendix B.4.

Proof Outline. Consider a buyer b. For (2), a best response satisfies the first order condition $\frac{d\mathbb{E}_{\text{-}b}[u_b(t_b,a_b,a_{\text{-}b})]}{da_b} = 0$, see Appendix B.3. Explicit calculations yield that there exists a constant $\kappa > 0$, such that $t_b - (1+\phi_b)\, a_b \leq \kappa q(n,m)$, with $q(m,n) = \max\left\{\frac{1}{n}\left(1+\frac{m}{n}\right), \frac{1}{m}\left(1+\frac{n}{m}\right)\right\} = O(\max(m,n)^{-1})$, from which the statement follows.[24] For (1), we estimate $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right]$, where a_b denotes the best response. This difference is shown to be upper bounded by $-2k(1+\phi_b)|t_b^\Phi - a_b|$. It follows from (2) that $\forall \delta > 0$ it holds in sufficiently large finite markets that $t_b^\Phi - a_b \leq \delta$. If for a given $\epsilon > 0$, $\delta > 0$ is chosen such that $\delta \leq \frac{\epsilon}{2k(1+\phi_b)}$, it holds that t_b^Φ is ϵ-close to a best response a_b in sufficiently large finite markets. In infinite markets, the expected utility is deterministic and truthfulness is a best response, as the only strategic incentive is to be involved in trade. Proof details are relegated to Appendix A.9. □

5.2 For influenceable transaction costs price-guessing is approximately optimal

If a trader can influence their transaction cost, then there remains a (non-vanishing) incentive to act strategically in large markets. Moreover, given a trader will almost certainly trade as long as their action meets the required threshold of the critical value, the incentive to influence their transaction cost asymptotically outweighs the concern of loosing out on the deal. Therefore, it is optimal to bid close to the critical value that corresponds to the predicted price, which is why we call such behavior *Price-Guessing*. While our analysis only covers the case of a trader for whom bidding the critical value is individually rational, the case of traders for whom it is not is covered in Proposition 19.

Theorem 9 (In large markets with influenceable transaction costs best responses are close to price guessing)**.** *If the transaction cost is influenceable and bidding the critical value P_i^∞ is strictly individually rational for trader i, then for every $\epsilon > 0$, in sufficiently large finite markets, all best responses of i are in an ϵ-neighbourhood of the critical value P_i^∞.*

This result extend to infinite markets with sufficiently small aggregate uncertainty, see Appendix B.4.

Proof Outline. Consider a buyer with action $a_b > P_b^\infty$. We show that if $a_b - P_b^\infty \geq \epsilon$, then the difference in expected utility from playing a_b versus $P_b^\infty + \frac{\epsilon}{2}$ is strictly negative in sufficiently large markets, proving that a_b is then not a best response. Similar to the proof of Theorem 6, we show that in such markets, the buyer will be involved in trade with high probability with both actions.

[24] A similar proof technique has been used to show that Bayesian Nash equilibria are approximately truthful in DAs without fees, see Rustichini et al. (1994, Theorem 3.1).

Using that the transaction cost is influenceable, the decrease of the transaction cost when switching to the more aggressive action $P_b^\infty + \frac{\epsilon}{2}$ outweighs the decrease in trading probability. Proof details are relegated to Appendix A.10. □

Spread fees. As a spread fee depends linearly on a trader's action, it is an example of an influenceable transaction cost. A best response exists given the spread fee is continuous and must be close to the critical value. However, an analogous statement to Theorem 8.2, i.e., the utility at the critical value is close to optimal, is not true in general. We show that there exist markets, such that bidding the critical value is in general not ϵ-optimal in large markets.

Theorem 10 (In large markets with spread fees best responses are close, but not necessarily equal, to the critical value)**.** *If the transaction cost is a strictly positive spread fee, then a best response exists for a trader i in finite and infinite markets without tie-breaking. Further, if bidding the critical value is strictly individually rational, then (1) for every $\epsilon > 0$, in sufficiently large markets, all best responses of i are in an ϵ-neighbourhood of the critical value P_i^∞ and (2) for sufficiently small $\epsilon > 0$, there exist beliefs, such that in sufficiently large finite markets the critical value P_i^∞ is not an ϵ-best response.*

Theorem 10 is robust to small aggregate uncertainty in infinite markets, see Appendix B.4.

Proof Outline. We show that in finite markets and infinite markets without tie-breaking the expected transaction cost and therefore the expected utility is continuous in a_i. The existence of a best response again follows as in Theorem 8. Consider a buyer b with $t_b^\Phi > P_i^\infty$. (1) is proven in complete analogy to Theorem 9. For (2), consider beliefs such that the number of traders is equal to l for both market sides, where beliefs are uniformly distributed over $A_B = A_S = [0,1]$. It follows that $P_b^\infty = \frac{1}{2}$. We prove that for every $l > 1$ it holds that $\mathbb{P}_{\text{-}b}[b \in \mathcal{B}^*(P_b^\infty, a_{\text{-}b})] = \frac{1}{2}$. Therefore, for every bid $a_b > P_b^\infty$ and for every $\epsilon > 0$, it follows from Proposition 3 that the buyer can increase their trading probability by $\frac{1}{2} - \epsilon$ when switching from P_b^∞ to a_b. If a_b is chosen close to P_b^∞, then this outweighs the increase in spread fee payment. Proof details are relegated to to Appendix A.11. □

5.3 Best responses and Bayesian Nash equilibria for price versus spread fees

Consider finite markets with sizes (i) 2×2 (that is, two buyers and two sellers), (ii) 5×5 and (iii) the infinite market with a unit mass of both buyers and sellers in the presence of either a price fee $\phi_i = 0.1$ or a spread fee $\phi_i = 1$, and $k = 0.5$. Figure 3 shows best response strategies (*top*) for uniform beliefs over others' actions in $[1,2]$ and a symmetric Bayesian Nash Equilibrium (*bottom*) for uniform beliefs over gross values in $[1,2]$ for price fees (*left*) and spread fees (*right*).

Best responses and Bayesian Nash equilibria are calculated using the first order conditions described in Appendix B.3. For best responses, the calculation is straight-forward. For Bayesian Nash equilibria, we adapt the method developed in Rustichini et al. (1994) for double auctions without transaction costs. Note that Figure 3 displays qualitatively representative Bayesian Nash equilibria as there is generally a multiplicity of equilibria.

In line with Theorem 8, optimal strategic behavior converges to truthfulness with growing market size, if price fees are charged. In a small market with two buyers and two sellers traders have an incentive to be more aggressive and misrepresent their net value, as can be measured by the distance between their respective best response (dashed red/blue lines) and the net value (solid black lines). In contrast, and in line with Theorem 8.1, the best responses (dotted red/blue line) in the larger market (5×5) are approaching truth-telling.

Note that in line with Theorem 10, best responses converge towards price-guessing with growing market size if spread fees are charged. In a small market with two buyers and two sellers traders have an incentive to be aggressive and misrepresent their true net value in order to influence the price and reduce their fee payment. Best responses in a larger market with five buyers and sellers (dotted line) do not approach truth-telling, if $t_i \prec P_i^\infty$. Instead traders remain aggressive as they aim to reduce their fee payment.

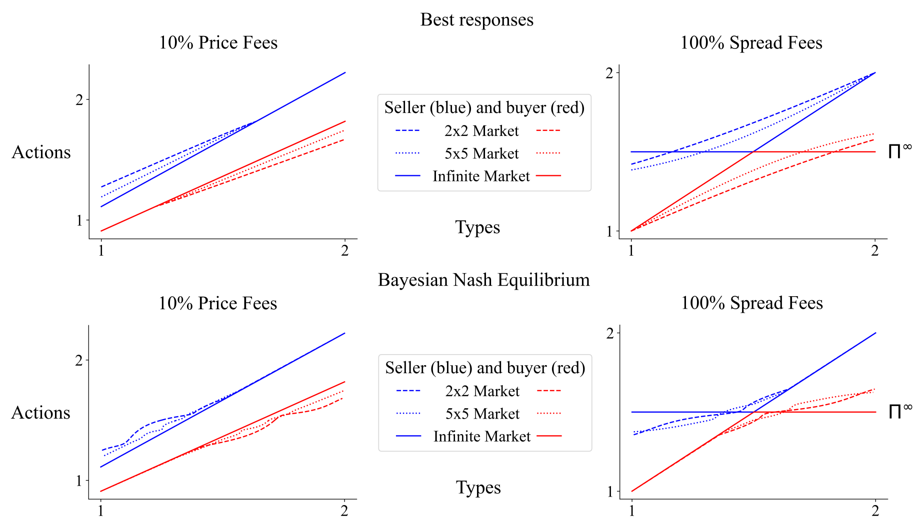

Figure 3: Best responses for uniform beliefs over actions (top) and a symmetric Bayesian Nash equilibrium for uniform beliefs over types (bottom) for buyers (red) and sellers (blue) as functions of their gross value for 2×2 (dashed lines) and 5×5 (dotted lines) markets with price fee $\phi_i = 0.1$ (left) or spread fee $\phi_i = 1$ (right).

6 Market Performance and Design

So far we did not commit to a specific nature of transaction costs. For example, transaction costs could have included shipping costs, taxes, and revenue of a market platform. As we study welfare metrics in this section, we now assume that all transaction costs are collected by some market platform. Then a social planner evaluates the market outcome by considering the social welfare and platform revenue according to some objective function. If the social planner can design the transaction cost, what is the optimal choice?

Suppose the social planner chooses the transaction cost (e.g., constant fee, price fee, spread fee) and its scale. For the latter, define the *scaling* of transaction costs Φ; for a two-dimensional parameter $\gamma = (\gamma_B, \gamma_s)$ with $\gamma_B, \gamma_S \geq 0$ the linear γ-scaling of transaction costs Φ are $\Phi_B^\gamma = \gamma_B \cdot \Phi$ and $\Phi_S^\gamma = \gamma_S \cdot \Phi$. For instance, for price and spread fees, γ-scaling linearly scales the fee percentage ϕ_i.

For analytical tractability, we restrict our analysis to infinite markets with type distributions μ_B^t and μ_S^t and regular transaction costs (Φ_B, Φ_S) that are charged to all buyers and all sellers involved in trade. Recall that those are transaction costs that only depend on a trader's action and the market price. As we work in an infinite market, we will use the term uninfluenceable transaction costs instead of asymptotically uninfluenceable transaction costs.

As in Section 5, traders best respond to their beliefs $\mathfrak{B}$ about the market environment. We allow that the belief system of traders may depend on the transaction costs, that is $\mathfrak{B}(\Phi)$.[25] This implies that a change in transaction costs, including their scaling, might influence the belief system of the traders. Recall from Section 3.3 the definitions of common and heterogeneous prior belief systems. We say that a belief system $\mathfrak{B}$ has a common prior, if all traders' beliefs lead to the same critical value, otherwise $\mathfrak{B}$ has a heterogeneous prior. We extend this definition to belief systems that depend on the transaction costs, and say that $\mathfrak{B}(\cdot)$ has a *common prior*, if for every transaction cost Φ, $\mathfrak{B}(\Phi)$ has a common prior. Otherwise, we say that $\mathfrak{B}(\cdot)$ has a *heterogeneous prior*. Moreover, we say that a belief system $\mathfrak{B}(\cdot)$ is *scaling-independent* for transaction costs Φ, if for any scaling γ it holds that $\mathfrak{B}(\Phi) = \mathfrak{B}(\Phi^\gamma)$.

For uninfluenceable transaction costs we focus on traders who truthfully report their net value. Truthfulness is a best response in infinite markets but other behaviors are also possible. However, we focus on truthfulness, as we show in Theorems 7 and 8, limit best-response behavior in large finite markets approaches truthfulness, and it is the unique best response in infinite markets with aggregate uncertainty, c.f. Appendix B.4. For influenceable transaction costs suppose that traders price-guess. This behavior is the unique best response in infinite markets and the approximate best response in large finite markets, see Theorems 9 and 10. This also holds in markets with sufficiently small aggregate uncertainty, c.f. Appendix B.4. Therefore the results in this section are qualitatively robust to small aggregate uncertainty, but nevertheless may differ, c.f. Section 2 for an analysis of market performance for spread fees in the presence of aggregate uncertainty.

6.1 Market Performance

The social planner evaluates the market outcome using the following performance metrics.

The *traders' welfare* $W = \int_{\mathcal{B}^*} u_b(t_b, a_b, a_{\text{-}b}) d\mu_B(b) + \int_{\mathcal{S}^*} u_s(t_s, a_s, a_{\text{-}s}) d\mu_S(s)$ is the overall utility of all traders involved in trade.[26] The *platform revenue* $R = \int_{\mathcal{B}^*} \Phi_b(a_b, a_{\text{-}b}) d\mu_B(b) + \int_{\mathcal{S}^*} \Phi_s(a_s, a_{\text{-}s}) d\mu_S(s)$ is the total amount of transaction costs that is collected by the market maker. The sum $G^{real} = W + R$ are the *realized gains of trade*, and note that $G^{real} = \int_{\mathcal{B}^*} (t_b - P^*)\, d\mu_B(b) + \int_{\mathcal{S}^*} (P^* - t_s)\, d\mu_S(s)$. If

[25] Note that this assumption has no bearing on our previous analysis of incentives in Section 5, where traders' behavior was studied for fixed transaction costs and beliefs.

[26] Because best responses are individually rational, W is non-negative.

agents report truthfully in the presence of transaction costs Φ, we denote by G^{net} the *net gains of trade*. If no transaction costs are charged reporting truthfully yields the *gross gains of trade* G^{gross}, where $G^{gross} \geq G^{net} \geq G^{real}$ will be shown to hold. The *loss* $L = G^{gross} - G^{real}$ measures how much gains of trade are lost. It can be split into the *direct loss* $L_\phi = G^{gross} - G^{net}$, that is due to transaction costs, and the *strategy-induced loss* $L_F = G^{net} - G^{real}$.

The gross gains of trade are equal to the sum of platform revenue, social welfare and loss, that is $G^{gross} = W + R + L$. We identify *market performance* with the triple (W, R, L). We normalize $G^{gross} = 1$ and hence the set of all performance triples lie on a triangle Δ in a 2-dimensional hyperplane in $\mathbb{R}^3$.

We say that a performance triple (W, R, L) is *achievable* for transaction costs Φ and belief system $\mathfrak{B}(\cdot)$, if there exists a γ-scaling, such that optimal behavior of all traders according to the belief system $\mathfrak{B}(\Phi^\gamma)$ leads to that market performance.

6.2 Optimal transaction cost design

Suppose that the social planner aims to maximize a continuous objective function $U : \Delta \to \mathbb{R}$ on performance triples $(W, R, L) \in \Delta$. We will consider objective functions, such that a Pareto improvement of welfare and revenue leads to an increase in utility, that is, for any performance triplet (W, R, L) and for $\alpha, \beta \geq 0$ with $0 < \alpha + \beta \leq L$ it holds that $U(W, R, L) < U(W + \alpha, R + \beta, L - (\alpha + \beta))$. In particular, it holds that $U(1, 0, 0) > U(0, 0, 1)$, that is the social planner prefers a fully efficient market with zero revenue for the market platform over a fully inefficient market. Moreover, we assume that the social planner values revenue, that is, there exists $R \in (0, 1]$, such that $U(1-R, R, 0) > U(1, 0, 0)$. Given a belief system $\mathfrak{B}(\cdot)$, we say that transaction costs Φ *(strictly) dominate* transaction costs Φ' if the resulting market performance $U(W, R, L)$ is (strictly) greater for Φ than for Φ'.

Optimal design of transaction costs depends on the nature of traders' beliefs. The following theorem shows that for common prior beliefs optimal design is possible for any uninfluenceable transaction cost and crucially, independent of the specific belief system. By contrast, for heterogeneous prior belief systems there exists no transaction cost that is always optimal.

Theorem 11 (Optimal Design)**.** *For all common prior belief systems, any uninfluenceable transaction cost can be scaled to dominate all transaction costs. Furthermore, the optimal scaling does not depend on the belief system. For some heterogeneous prior belief systems, there exists a purely influenceable transaction cost that dominates all transaction costs and strictly dominates all uninfluenceable transaction costs.*

Notably, for common prior belief systems, the optimal design problem is reduced from the space of all transaction costs to a one-dimensional optimization problem of finding the optimal scaling for any uninfluenceable transaction cost. For some heterogeneous prior belief systems, influenceable transaction costs, even without an uninfluenceable part, can strictly outperform any uninfluenceable transaction cost. However, there also exist heterogeneous prior belief systems, such that any uninfluenceable transaction cost U-dominates all purely influenceable transaction costs, as the latter class would lead to market failure.

The proof is relegated to Appendix A.12. We will omit a proof outline for Theorem 11 as it combines results of the more detailed analysis that follows. Concretely, we analyse uninfluenceable and purely influenceable transaction costs separately, to detail what market performances are achievable and how they depend on the belief system and scaling. Finally, we return to mixed transaction costs to discuss under what circumstances it may be optimal to use those.

6.2.1 Uninfluenceable Transaction Costs

Suppose that the market platform charges uninfluenceable transaction costs Φ. The following proposition characterizes the set of all achievable market performances: First, it is fully specified by the type distributions μ_B^t and μ_S^t, but does not depend on the choice of the uninfluenceable transaction costs, or the traders' belief system. This implies that any market performance achievable with one uninfluenceable transaction cost can be achieved with another provided it is properly scaled. Second, this set is one-dimensional, as scaling is the only way to influence the market performance, which in turn implies that most performances are not achievable with uninfluenceable transaction costs.

Proposition 12 (Equivalence of uninfluenceable transaction costs)**.** *The set of achievable performance triples (W, R, L) is the same for all uninfluenceable transaction costs and belief systems. The set is a curve $c_P : [0,1] \to \Delta$ in the simplex Δ of all performance triples.*

Proof Outline. We prove that for any uninfluenceable transaction cost Φ and belief system $\mathfrak{B}$, the market performance can be represented as a continuous function of the trading volume Q^*_{net}, that is $Q^*_{net} \mapsto (W(Q^*_{net}), R(Q^*_{net}), L(Q^*_{net}))$. The gross gains of trade are the area under the gross demand and supply curve. The net gains of trade are then given by truncating this area at the height of the net trading volume. The loss is therefore the area under the gross demand and supply curve, with height between net and gross trading volume. The revenue is equal to the maximum rectangle with height equal to the net trading volume that fits under the gross demand and supply curve. The welfare is equal to the area under gross demand and supply that is left (for sellers) and right (for buyers) of the revenue rectangle. See Figure 4. We prove that for any $V \in [0, Q^*_{gross}]$, there exists a scaling, such that the trading volume Q^*_{net} is equal to V. The trading volume is the intersection of net demand and supply. Proof details are relegated to Appendix A.13. $\square$

The performance curve c_P has several interesting properties: First, it connects the fully efficient market outcome with zero revenue $(1, 0, 0)$ and the fully inefficient market outcome $(0, 0, 1)$ that corresponds to complete market failure. Second, c_P is strictly increasing in the loss and strictly decreasing in the welfare. Therefore, for any level of welfare in $[0, 1]$, there exists a scaling to achieve it. The revenue, as well as the loss, are then uniquely determined by the curve c_P. This implies that positive platform revenue with uninfluenceable transaction costs is directly tied to a positive loss of efficiency. See Figure 4.

This has immediate consequences for the optimal design of uninfluenceable transaction costs. The market maker is restricted to a one-dimensional set of achievable performance triples that is

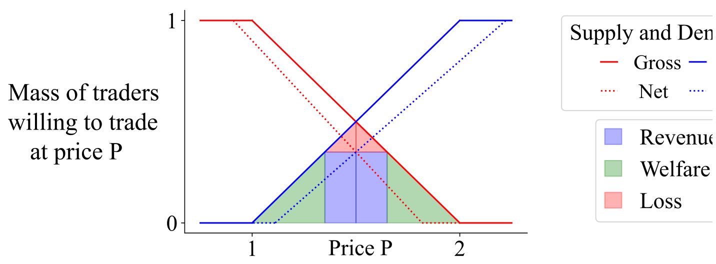

(a) Market performance

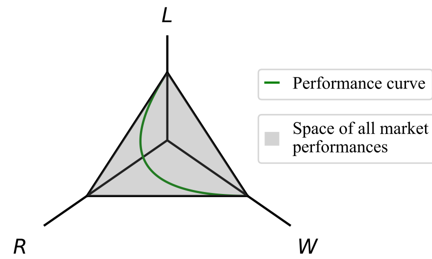

(b) Performance curve

Figure 4: Market performance for price fees and uniformly distributed types. *Left.* Decomposition of the gross gains of trade into revenue (blue), welfare (green), and loss (red) for symmetric price fees $\phi_B = \phi_S = 0.1$. *Right.* Performance curve for price fees.

fully specified by the type distributions. Given their objective function U, the U-dominant market performance is then achieved by scaling any uninfluenceable transaction cost properly. The belief system of traders does not influence the optimal design.

Corollary 13 (Optimal design of uninfluenceable transaction costs)**.** *For any uninfluenceable transaction costs, there exists a scaling, such that for any belief system the scaled transaction cost dominates all uninfluenceable transaction costs.*

Proof. It follows from Proposition 12 that the set of achievable market performances is the same curve in Δ for all uninfluenceable transaction costs. Because the objective function U is continuous, the Extreme Value Theorem implies that there exists a maximum (W, R, L) on this compact subset of Δ. In the proof of Proposition 12, we have shown that any achievable market performance is fully specified by the net trading volume and that for any trading volume $V \in [0, Q^*_{gross}]$, there exists a scaling, such that the trading volume Q^*_{net} is equal to V. Therefore, for any uninfluenceable transaction cost, there exists a scaling that achieves the U-maximum (W, R, L). $\square$

If the social planner wants to maximize efficiency (that is, minimize the loss), zero transaction costs are optimal. This leads to maximum traders' welfare and zero revenue. For a revenue-maximizing market-maker, there is a non-trivial trade-off between higher transaction costs per trader and trading volume. It follows from the proof of Proposition 12 that the total platform revenue is equal to the area of the rectangle with height equal to the net trading volume that fits under the true demand and supply curve. Maximizing platform revenue is therefore an optimization problem with respect to the net trading volume, for which an optimal solution exists. Once the optimal trading volume is determined, any uninfluenceable transaction cost can be scaled to lead to that trading volume, that is net demand and supply intersect at that height. The horizontal component of the crossing point, that is the market price, determines, how much of the revenue is paid by buyers and sellers. The area to the rectangle left to the market price is paid by sellers, and the area to the right by buyers. With the right scaling – different to buyers and sellers – any market price on the horizontal rectangle can be achieved. Hence, if the total revenue is equal to R, for any $\alpha \in [0, 1]$, there exists a scaling, such that the revenue generated by buyers is $\alpha \cdot R$ and the revenue generated by sellers is $(1 - \alpha) \cdot R$.

It is natural to ask how the maximum revenue depends on the distributions of buyer and seller types. In the example of Section 2, with uniform type distributions and price fees, we showed that the maximum revenue is equal to 0.5 and thus half the gains of trade can be extracted as revenue. It turns out, that the extractable revenue can be anywhere between zero and G^{gross}, dependent on type distributions. For any $\epsilon > 0$, there exist type distributions, such that the generated revenue is less or equal to ϵ or greater or equal to $G^{gross} - \epsilon$. [27]

6.2.2 Purely Influenceable Transaction Costs

Suppose that the market platform charges purely influenceable transaction costs Φ. The following proposition characterizes the set of all achievable market performances: For common prior belief systems, platform revenue is always zero. If the belief system is scaling-independent, the market maker cannot influence the distribution of welfare and loss, that is, there is a unique achievable market performance. Second, if traders have heterogeneous prior beliefs, revenue might be positive and the market platform has some influence on the welfare-revenue distribution via scaling.

Proposition 14 (Non-equivalence of scaled influenceable transaction costs)**.** *Suppose the transaction costs Φ are purely influenceable.*

- *For any common prior belief system, achievable market performances have zero revenue, $R = 0$. Moreover, if the belief system is scaling-independent, then there exists a unique achievable market outcome. If this belief system is calibrated, then the market is fully efficient, $W = 1$.*
- *For any heterogeneous prior belief system that is scaling-independent, the set of achievable market performances is a singleton or line-segment with constant loss L. Furthermore, for any $L \in [0, 1]$, there exists such a belief system that leads to loss L.*

Proof Outline. We show that the loss L is fully characterized by the belief system as players price-guess and therefore the loss is independent of the scaling of a purely influenceable transaction cost. For common prior beliefs, price guessing leads to market outcomes, where all traders involved in trade submitted an action equal to the realized market price P^*. Therefore, purely influenceable transaction costs lead to zero revenue. If the belief system is calibrated around the true critical value, then the trading volume is maximized and the market is fully efficient. For heterogeneous prior beliefs, scaling of the transaction costs leads to a continuous increase or decrease in revenue. As the loss is fixed, this yields that the set of achievable market performances is a line-segment or singleton. To show that any loss can be realized, we construct belief systems such that the traders with the most profitable gross values are involved in trade with price-guessing. Then, the loss is a continuous function of the trading volume. We prove that any trading volume can be realized with some heterogeneous prior beliefs. Proof details are relegated to Appendix A.14. □

[27] The construction is based on demand and supply functions, that are either extremely convex or concave.

Note that if the belief system has a common prior and is scaling-independent, then the market maker has no influence on the market performance (W, R, L) via the choice of the purely influenceable transaction cost. For heterogeneous prior beliefs, the market maker can influence the welfare-revenue distribution. For the special case of spread fees, any market performance (W, R, L) is achievable for some belief system and scaling. This is in stark contrast to uninfluenceable transaction cost, where only a one-dimensional subset of the space of all performance triples is achievable regardless of the belief system.

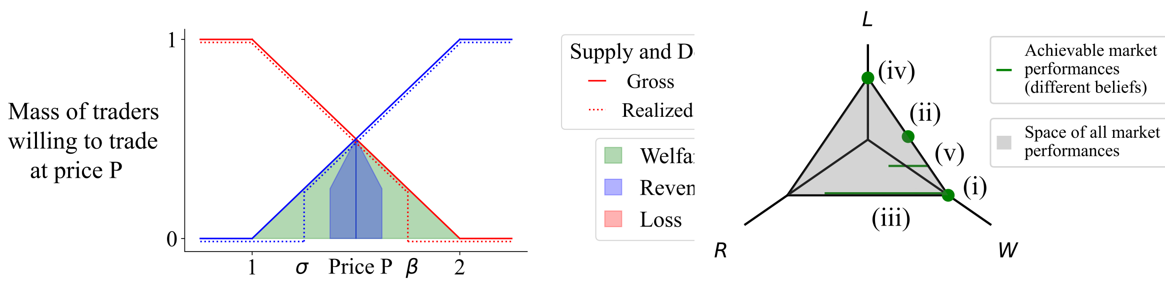

(a) Market performance, case (iii)

(b) Performance curve

Figure 5: Market performance for spread fees, uniformly distributed types, and different scaling-independent belief systems. *Left.* Decomposition of the gross gains of trade into revenue (blue), welfare (green), and loss (red; zero) for symmetric spread fees $\phi_B = \phi_S = 0.5$ and beliefs $\beta = 1.75$ and $\sigma = 1.25$. *Right.* Achievable market performances for (i) calibrated beliefs $\beta = \sigma = 1.5$, (ii) a homogeneous bias $\beta = \sigma = 1.75$, (iii) a conservative bias $\beta = 1.75$, $\sigma = 1.25$, (iv) an aggressive bias $\beta < \sigma$, and (v) a mixed bias $\beta = 1.8$, $\sigma = 1.6$

.

Note that for some belief systems (e.g., a scaling-independent aggressive bias as illustrated in Section 2, complete market failure, that is, $(0, 0, 1)$, is the only achievable market performance. However, the market maker has the possibility to scale the transaction cost to zero, which is not purely influenceable and thus leads to the fully efficient market with zero revenue, that is $(1, 0, 0)$. See Figure 5.

Proposition 14 implies that the optimal design of purely influenceable transaction costs crucially depends on the traders' belief system. For some belief systems, including the ones that have a common prior, it turns out to be optimal to not charge any purely influenceable transaction costs at all.

Corollary 15 (Optimal design of purely influenceable transaction costs)**.** *Suppose the transaction costs Φ are purely influenceable.*

- *For any common prior belief system, zero transaction cost leads to a fully efficient market, dominates all purely influenceable transaction costs.*

- *For some heterogeneous prior belief systems, zero transaction cost dominates all purely influenceable transaction costs.*

Proof. It follows from Proposition 14 that for common prior beliefs, there exists a unique achievable market performance with zero revenue, that is $(1 - L, 0, L)$ for some loss $L \in [0, 1]$. By assumption, the social planner values welfare over loss and it is thus optimal to not charge any purely influenceable transaction costs. Then, traders act truthfully, which yields the fully efficient market performance without revenue, that is $(1, 0, 0)$. For some heterogeneous prior beliefs, it follows from Proposition 14 that any purely uninfluenceable transaction cost will lead to complete market failure, that is $(0, 0, 1)$, due to price-guessing behavior. In that case not charging any transaction cost with the fully efficienct market outcome $(1, 0, 0)$ is again U-optimal. □

For some belief systems and influenceable transaction costs, in contrast to uninfluenceable transaction costs, it is possible to achieve market outcomes with strictly positive revenue and zero loss, that is $(x, 1 - x, 0)$. Note that for some belief systems and objective functions, a high scaling might be optimal. For example, if spread fees are charged and beliefs are scaling-independent, have a heterogeneous prior, and lead to strictly positive gains of trade, it is optimal for a revenue-maximizing market maker to scale the fees to 100%. Moreover, for certain belief systems, it is possible to achieve the optimal performance triplet (W, R, L) in the space of all market performances. More formally, for any objective function U, there exist spread fees and a belief system, such that the corresponding market performance maximizes the objective function U over the space of all market performances Δ. For example, for a revenue-maximizing market maker, there exist belief systems such that the market performance $(0, 1, 0)$ is achievable.

7 Conclusion

We have studied a market environment in which the price is set to equate revealed supply and demand and we have shown that the presence of transaction costs may fundamentally alter incentives and welfare in markets. We have shown that the traders' strategic behavior and market performance hinge on whether transaction costs are asymptotically uninfluenceable (as, for instance, fixed fees and price fees) or influenceable (as, for instance, spread fees). Uninfluenceable transaction costs don't fundamentally alter strategic incentives and, in large markets, inefficiency only arises from the direct loss that resembles the dead-weight loss of taxation or monopoly power. By contrast influenceable transaction costs starkly alter strategic consideration and market performance; total market failure may even occur.

While we focus on the popular Double Auction mechanism, our results remain valid for any mechanism in which:

- A trader's expected utility $\mathbb{E}[u(a)]$, given their action a, can be expressed as the product of the probability of trade $\mathbb{P}[\text{trading given } a]$ and the expected utility conditional on trading $\mathbb{E}[u(a)|\text{trading given } a]$; we assume here that the utility when not trading is zero.
- A buyer's $\mathbb{P}[\text{trading given } a]$ is increasing in a and $\mathbb{E}[u(a)|\text{trading given } a]$ is decreasing in a. A seller's $\mathbb{P}[\text{trading given } a]$ is decreasing in a and $\mathbb{E}[u(a)|\text{trading given } a]$ is increasing in a.

- Trade is *predictable* in the sense that the dependence of the probability of trade on the trader's action approaches a 0-1 step-function.

Note that trade is predictable in the above sense, for instance, in large markets without aggregate uncertainty (c.f. Proposition 3) and in posted-price mechanisms. For this class of mechanisms, an analog of our categorization into asymptotically uninfluenceable and influenceable payments remains crucial and our analysis of the two categories carries over to this more general setting. Vickrey mechanisms are examples of the asymptotically uninfluenceable category, and first-price auctions in environments with one seller and many buyers are examples of the influenceable category.

The key role of the type of transaction costs that we have uncovered suggests further empirical and theoretical questions. Given that both asymptotically uninfluenceable and influenceable transaction costs are charged in practice, what explains the choices? May the choice depend on differences in sophistication of traders; for example influenceable transaction costs might be charged in situations where traders have incorrect beliefs or face aggregate uncertainty. How to optimally design information that market participants have? How do our results extend to more complex market interactions, where traders are interested in bundles or where platforms compete?

References

Armstrong, Mark (2006), "Competition in two-sided markets." *RAND Journal of Economics*, 37, 668–691.

Aumann, Robert (1964), "Markets with a continuum of traders." *Econometrica*, 32, 39–50.

Azevedo, Eduardo and Eric Budish (2019), "Strategy-proofness in the Large." *Review of Economic Studies*, 86, 81–116.

Bergemann, Dirk and Stephen Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1771–1813.

Börgers, Tilman and Jiangtao Li (2019), "Strategically simple mechanisms." *Econometrica*, 87, 2003–2035.

Carroll, Gabriel (2015), "Robustness and linear contracts." *American Economic Review*, 105, 536–563.

Chassang, Sylvain (2013), "Calibrated incentive contracts." *Econometrica*, 81, 1935–1971.

Chatterjee, Kalyan and William Samuelson (1983), "Bargaining under incomplete information." *Operations Research*, 31, 835–851.

Chen, Daniel and Anthony Lee Zhang (2020), "Subsidy schemes in double auctions." *Available at SSRN 3668394*.

Chu, Leon Yang and Zuo-Jun Max Shen (2006), "Agent competition double-auction mechanism." *Management Science*, 52, 1215–1222.

Colliard, Jean-Edouard and Thierry Foucault (2012), "Trading fees and efficiency in limit order markets." *Review of Financial Studies*, 25, 3389–3421.

Cripps, Martin and Jeroen Swinkels (2006), "Efficiency of large double auctions." *Econometrica*, 74, 47–92.

Cunningham, Colleen, Florian Ederer, and Song Ma (2021), "Killer acquisitions." *Journal of Political Economy*, 129, 649–702.

de Clippel, Geoffroy and Kareen Rozen (2018), "Consumer theory with misperceived tastes." *Available at SSRN 3223654*.

Demsetz, Harold (1968), "The cost of transacting." *Quarterly Journal of Economics*, 82, 33–53.

Diamond, Peter and James Mirrlees (1971a), "Optimal taxation and public production i: Production efficiency." *American Economic Review*, 61, 8–27.

Diamond, Peter and James Mirrlees (1971b), "Optimal taxation and public production ii: Tax rules." *American Economic Review*, 61, 261–278.

Eyster, Erik and Michele Piccione (2013), "An approach to asset pricing under incomplete and diverse perceptions." *Econometrica*, 81, 1483–1506.

Foucault, Thierry, Ohad Kadan, and Eugene Kandel (2013), "Liquidity cycles and make/take fees in electronic markets." *Journal of Finance*, 68, 299–341.

Friedman, Dan and John Rust (1993), *The double auction market: institutions, theories, and evidence.* Westview Press.

Fudenberg, Drew, Markus Mobius, and Adam Szeidl (2007), "Existence of equilibrium in large double auctions." *Journal of Economic Theory*, 133, 550–567.

Galeotti, Andrea and José Luis Moraga-González (2009), "Platform intermediation in a market for differentiated products." *European Economic Review*, 53, 417–428.

Garratt, Rodney and Marek Pycia (2016), "Efficient bilateral trade." *Available at SSRN 2444279*.

Glosten, Lawrence and Paul Milgrom (1985), "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders." *Journal of Financial Economics*, 14, 71–100.

Harberger, Arnold (1962), "The incidence of the corporation income tax." *Journal of Political Economy*, 70, 215–240.

Harberger, Arnold (1964), "The measurement of waste." *American Economic Review*, 54, 58–76.

Harberger, Arnold (1971), "Three basic postulates for applied welfare economics: an interpretive essay." *Journal of Economic Literature*, 9, 785–797.

Harrison, John Michael and David Kreps (1979), "Martingales and arbitrage in multiperiod securities markets." *Journal of Economic Theory*, 20, 381–408.

Heidhues, Paul, Botond Kőszegi, and Philipp Strack (2018), "Unrealistic expectations and misguided learning." *Econometrica*, 86, 1159–1214.

Hines Jr, James (1999), "Three sides of harberger triangles." *Journal of Economic Perspectives*, 13, 167–188.

Hosseini, Hamid (1995), "Understanding the market mechanism before adam smith: economic thought in medieval islam." *History of Political Economy*, 27, 539–561.

Jackson, Matthew and Jeroen Swinkels (2005), "Existence of equilibrium in single and double private value auctions." *Econometrica*, 73, 93–139.

Jantschgi, Simon, Heinrich Nax, Bary Pradelski, and Marek Pycia (2022), "Double auctions and walrasian equilibrium." *Available at SSRN 4043532*.

Jullien, Bruno and Alessandro Pavan (2019), "Information management and pricing in platform markets." *Review of Economic Studies*, 86, 1666–1703.

Kang, Zi Yang and Ellen V Muir (2022), "Contracting and vertical control by a dominant platform." *Unpublished manuscript, Stanford University.*

Ledyard, John (1978), "Incentive compatibility and incomplete information." *Journal of Economic Theory*, 18, 171–189.

Leininger, Wolfgang, Peter Linhart, and Roy Radner (1989), "Equilibria of the sealed-bid mechanism for bargaining with incomplete information." *Journal of Economic Theory*, 48, 63–106.

Li, Shengwu (2017), "Obviously strategy-proof mechanisms." *American Economic Review*, 107, 3257–3287.

Madarász, Kristóf and Andrea Prat (2017), "Sellers with misspecified models." *Review of Economic Studies*, 84, 790–815.

Malinova, Katya and Andreas Park (2015), "Subsidizing liquidity: The impact of make/take fees on market quality." *Journal of Finance*, 70, 509–536.

Marra, Marleen (2021), "Estimating an auction platform game with two-sided entry." Technical report, Working Paper, Sciences Po.[507].

Mirrlees, James (1986), "The theory of optimal taxation." *Handbook of Mathematical Economics*, 3, 1197–1249.

Myerson, Roger and Mark Satterthwaite (1983), "Efficient mechanisms for bilateral trading." *Journal of Economic Theory*, 29, 265–281.

Noussair, Charles, Stephane Robin, and Bernard Ruffieux (1998), "The effect of transaction costs on double auction markets." *Journal of Economic Behavior & Organization*, 36, 221–233.

Pycia, Marek and Peter Troyan (2023), "A theory of simplicity in games and mechanism design." *Econometrica*, Forthcoming.

Ramsey, Frank (1927), "A contribution to the theory of taxation." *Economic Journal*, 37, 47–61.

Reny, Philip and Motty Perry (2006), "Toward a strategic foundation for rational expectations equilibrium." *Econometrica*, 74, 1231–1269.

Roberts, Donald John and Andrew Postlewaite (1976), "The incentives for price-taking behavior in large exchange economies." *Econometrica*, 44, 115–127.

Rochet, Jean-Charles and Jean Tirole (2003), "Platform competition in two-sided markets." *Journal of the European Economic Association*, 1, 990–1029.

Rustichini, Aldo, Mark Satterthwaite, and Steven Williams (1994), "Convergence to efficiency in a simple market with incomplete information." *Econometrica*, 62, 1041–1063.

Satterthwaite, Mark and Steven Williams (1989a), "Bilateral trade with the sealed bid k-double auction: Existence and efficiency." *Journal of Economic Theory*, 48, 107 – 133.

Satterthwaite, Mark and Steven Williams (1989b), "The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large." *Review of Economic Studies*, 56, 477–498.

Shi, Bing, Enrico Gerding, Perukrishnen Vytelingum, and Nicholas Jennings (2013), "An equilibrium analysis of market selection strategies and fee strategies in competing double auction marketplaces." *Autonomous Agents and Multi-Agent Systems*, 26, 245–287.

Smith, Adam (1776), *An Inquiry into the Nature and Causes of the Wealth of Nations*. Strahan and Cadell.

Spiegler, Ran (2020), "Behavioral implications of causal misperceptions." *Annual Review of Economics*, 12, 81–106.

Tatur, Tymon (2005), "On the trade off between deficit and inefficiency and the double auction with a fixed transaction fee." *Econometrica*, 73, 517–570.

Tyurin, Ilya (2012), "A refinement of the remainder in the lyapunov theorem." *Theory of Probability & Its Applications*, 56, 693–696.

Vapnik, Vladimir and Alexey Chervonenkis (1971), "On the uniform convergence of relative frequencies of events to their probabilities." *Theory of Probability and its Applications*, 16, 264.

Walras, Leon (1874), *Principe d'une theorie mathematique de l'echange, memoire lu a l'Academie des sciences morales et politiques*. Guillaumin.

Williams, Steven (1991), "Existence and convergence of equilibria in the buyer's bid double auction." *Review of Economic Studies*, 58, 351–374.

Wilson, Robert (1985a), *Game-theoretic Analyses of Trading Processes*. Institute for Mathematical Studies in the Social Sciences, Stanford University.

Wilson, Robert (1985b), "Incentive efficiency of double auctions." *Econometrica*, 53, 1101–1115.

Wolitzky, Alexander (2016), "Mechanism design with maxmin agents: Theory and an application to bilateral trade." *Theoretical Economics*, 11, 971–1004.

A Proofs

A.1 Proof of Proposition 1

Proof. Consider a buyer b with gross value $t_b \in T_b^+$. First, we prove that there exists a unique solution to the equation $t_b - x - \Phi_b(x,x) = 0$. Because $t_b \in T_b^+$, there exists an action a_b such that $t_b - a_b - F_b(a_b, a_b) > 0$. Furthermore, for $a_b > t_b$, it holds that $t_b - a_b - F_b(a_b, a_b) < 0$. Because the function $x \mapsto t_b - x - F_b(x,x)$ is continuous and strictly decreasing, there exists a unique zero point by the Intermediate Value Theorem.

Existence. Next, we show that this solution x is equal to the net value t_b^Φ, by proving that x is undominated, it dominates every larger action a_b, it is ex-post individually rational, and no larger action a_b is ex-post individually rational. Consider $a_b > x$. If $a_{\text{-}b}$ is such that buyer b is not involved in trade with x and a_b, then the utility is equal to 0 for both actions. If $a_{\text{-}b}$ is such that b is involved in trade with both actions, then it follows that $u_b(t_b, x, a_{\text{-}b}) \geq u_b(t_b, a_b, a_{\text{-}b})$, because the fee is monotone. If $a_{\text{-}b}$ is such that b is only involved in trade with a_b, then then the market price $P^*(a_b, a_{\text{-}b})$ is greater or equal than x. It holds that $u_b(t_b, a_b, a_{\text{-}b}) \leq u_b(t_b, P^*(a_b, a_{\text{-}b}), a_{\text{-}b}) = t_b - P^*(a_b, a_{\text{-}b}) - \Phi_b(P^*(a_b, a_{\text{-}b}), P^*(a_b, a_{\text{-}b})) \leq t_b - x - \Phi_b(x,x) = 0$. The first inequality follows from the monotonicity of the fee, the second inequality follows, because the map $a_i \mapsto P_i(a_i, a_i)$ is strictly increasing, and the final equality follows from the definition of x. Therefore a_b is dominated by x. Consider $a_b < x$. We show that there exists $a_{\text{-}b}$ such that $u_b(t_b, x, a_{\text{-}b}) > u_b(t_b, a_b, a_{\text{-}b})$. Take $a_{\text{-}b}$, such that buyer b is involved in trade only with x and the market price is strictly less than x. It holds that $u_b(t_b, x, a_{\text{-}b}) = t_b - P^*(x, a_{\text{-}b}) - \Phi_b(x, P^*(x, a_{\text{-}b}) > t_b - x - \Phi_b(x,x) = 0$. The inequality follows from regularity of the fee. Therefore x is not dominated by a_b. To show that x is ex-post individually rational, take any distribution of actions $a_{\text{-}b}$. If buyer b is involved in trade with x, it holds that $P^*(x, a_{\text{-}b}) \leq x$ and therefore $u_b(t_b, x, a_{\text{-}b}) = t_b - P^*(x, a_{\text{-}b}) - \Phi_b(x, P^*(x, a_{\text{-}b})) \geq t_b - x - \Phi_b(x,x) = 0$, where the inequality follows from regularity. Finally, we show that $a_b > x$ is not ex-post individually rational. Take $a_{\text{-}b}$, such that buyer b is involved in trade with a_b and $P^*(a_b, a_{\text{-}b}) > x$. It holds that $u_b(t_b, a_b, a_{\text{-}b}) \leq u_b(t_b, P^*(a_b, a_{\text{-}b}), a_{\text{-}b}) = t_b - P^*(a_b, a_{\text{-}b}) - \Phi_b(P^*(a_b, a_{\text{-}b}), P^*(a_b, a_{\text{-}b})) < t_b - x - \Phi_b(x,x) = 0$, where the first inequality follows from monotonicity, and the second one follows, because the map $a_i \mapsto P_i(a_i, a_i)$ is strictly increasing. This finally proves that $x = t_b^\Phi$. Therefore, the net value exists and the supremum is attained as a maximum.

Continuity. It was proven above that the net value exists on T_b^+ and is equal to the unique zero point of the function $x \mapsto t_b - x - \Phi_b(x,x)$. Because this function is strictly increasing and continuous, the zero point continuously depends on the gross value t_b.

Monotonicity. The map $t_b \mapsto t_b - x - \Phi_b(x,x)$ is strictly increasing. Therefore, the zero point of the map $x \mapsto t_b - x - \Phi_b(x,x)$ is strictly increasing in t_b.

The statements for sellers can be proven analogously. $\square$

A.2 Proof of Corollary 2

Proof. **Spread fees.** It holds that $\Phi_b(a_b, a_{\text{-}b}) = \phi_b(a_b - P^*(a_b, a_{\text{-}b})) = F_b(a_b, P^*(a_b, a_{\text{-}b}))$ with the function $F_b(x, y) = \phi_b(x - y)$. It holds that the map $y \mapsto y + F_b(x, y) = \phi_b x + (1 - \phi_b)y$ is increasing, the map $x \mapsto x + F_b(x, x) = x$ is strictly increasing in y and both are continuous. Therefore spread fees satisfy the conditions of Proposition 1. For any t_b, there exists a unique solution of $t_b - t_b^\Phi - F_b(t_b^\Phi, t_b^\Phi) = 0$. It is given by $t_b^\Phi = t_b$, proving that the net value equals the gross value.

Price fees. It holds that $\Phi_b(a_b, a_{\text{-}b}) = \phi_b P^*(a_b, a_{\text{-}b}) = F_b(a_b, P^*(a_b, a_{\text{-}b}))$ with the function $F_b(x, y) = \phi_b y$. It holds that the maps $y \mapsto y + F_b(x, y) = (1 + \phi_b)y$ and $x \mapsto x + F_b(x, x) = (1 + \phi_b)x$ are strictly increasing and continuous. Therefore price fees satisfy the conditions of Proposition 1. The unique solution of $t_b - t_b^\Phi - F_b(t_b^\Phi, t_b^\Phi) = 0$ is given by $t_b^\Phi = \frac{t_b}{1+\phi_b}$, proving that the net value scales the gross value.

Constant fees. It holds that $\Phi_b(a_b, a_{\text{-}b}) = c_b = F_b(a_b, P^*(a_b, a_{\text{-}b}))$ with the function $F_b(x, y) = c_b$. It holds that the maps $y \mapsto y + F_b(x, y) = y + c_b$ and $x \mapsto x + F_b(x, x) = x + c_b$ are continuous and strictly increasing in y. Therefore constant fees satisfy the conditions of Proposition 1. There exists a solution to $t_b - t_b^\Phi - F_b(t_b^\Phi, t_b^\Phi) = 0$, if $t_b \geq c_b$. It is given by $t_b^\Phi = t_b - c_b$, proving that the net value shifts the gross value.

The statement for sellers can be proven analogously. □

A.3 Proof that the critical value P_i^∞ exists and is unique

Proof. At the point $\underline{a}_{S,i}$, it holds that $F_{B,i}^a(\underline{a}_{S,i}) < 1$. That is because $F_{B,i}^a$ has a strictly positive density $f_{B,i}^a$ on $[\underline{a}_{B,i}, \overline{a}_{B,i}]$ and $\underline{a}_{S,i} < \overline{a}_{B,i}$ by assumption. Second, it holds that $F_{S,i}^a(\underline{a}_{S,i}) = 0$, because the corresponding density $f_{S,i}^a$ has support $[\underline{a}_{S,i}, \overline{a}_{B,i}]$. Therefore, at $\underline{a}_{S,i}$, it holds that $F_{B,i}^a(\underline{a}_{S,i}) + R_i F_{S,i}^a(\underline{a}_{S,i}) < 1$. A similar argument yields that at the point $\overline{a}_{B,i}$, it holds that $F_{B,i}^a(\overline{a}_{B,i}) = 1$ and $F_{S,i}^a(\overline{a}_{S,i}) > 0$. This implies that $F_{B,i}^a(\overline{a}_{B,i}) + R_i F_{S,i}^a(\overline{a}_{B,i}) > 1$. Because $F_{B,i}^a$ and $F_{S,i}^a$ are both continuous, it follows from the Intermediate Value Theorem, that there exists $P_i^\infty \in (\underline{a}_{S,i}, \overline{a}_{B,i})$ with $F_{B,i}^a(P_i^\infty) + R_i F_{S,i}^a(P_i^\infty) = 1$. Because both $F_{B,i}^a$ and $F_{S,i}^a$ are strictly monotone on $(\underline{a}_{S,i}, \overline{a}_{B,i})$, the uniqueness of P_i^∞ follows. □

A.4 Proof of Proposition 3

Proof. For trader i, consider a sequence of strictly increasing market sizes $(m(l), n(l))_{l\in\mathbb{N}}$ with $m(l), n(l) = \Theta(l)$ and $|R - \frac{n(l)}{m(l)}| = \mathcal{O}(l^{-1})$ for $R \in (0, \infty)$.[28]

Consider a buyer b. It holds that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] = \mathbb{P}_{\text{-}b}\left[a_b \geq a_{\text{-}b}^{m(l)}\right]$.[29] This is equal to the probability that at least $m(l)$ actions are below a_b in a sample of actions from $m(l) - 1$ buyers and $n(l)$ sellers. Let $p_{a_b} = F_{B,b}(a_b) \in (0, 1)$ be the probability that another buyer's bid is below a_b. In analogy, define $q_{a_b} = F_{S,b}(a_b) \in (0, 1)$ for sellers. For $i > 0$ let $X_i^{p_{a_b}}$ denote an independent Bernoulli

[28]This means that both market sides are assumed to have linear growth with respect to a single parameter l, such that neither side of the market dominates the other asymptotically and the ratio of buyers to sellers converges and fluctuates only slightly in finite markets.

[29]See Rustichini et al. (1994); Jantschgi et al. (2022).

random variable with parameter p_{a_b} and for $j > 0$ let $X_j^{q_{a_b}}$ denote an independent Bernoulli random variable with parameter q_{a_b}. Define $S_l^{a_b} = \sum_{i=1}^{m(l)-1} X_i^{p_{a_b}} + \sum_{j=1}^{n(l)} X_j^{q_{a_b}}$. $S_l^{a_b}$ has the same distribution as the number of traders in a sample of $m(l) - 1$ buyers and $n(l)$ sellers, whose actions are less or equal than a_b. It follows that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] = \mathbb{P}[S_l^{a_b} \geq m(l)] = 1 - \mathbb{P}[S_l^{a_b} \leq m(l) - 1]$. Next, we will show that a properly normalized version of $S_l^{a_b}$ converges in distribution to a standard normal random variable. This follows as an application of the following version of the Berry-Esseen theorem, see Tyurin (2012):

Theorem 16 (Berry-Esseen)**.** *Suppose $X_1, X_2, \ldots$ is a sequence of independent random variables with (i) $\mu_i = \mathbb{E}[X_i] < \infty$, (ii) $\sigma_i^2 = \mathbb{E}[(X_i - \mu_i)^2] < \infty$ and (iii) $\rho_i = \mathbb{E}[|X_i - \mu_i|^3] < \infty$. Set $r_n = \sum_{i=1}^n \rho_i$, $s_n^2 = \sum_{i=1}^n \sigma_i^2$, $F_n(x) = \mathbb{P}\left[\frac{\sum_{i=1}^n (X_i - \mu_i)}{\sqrt{s_n^2}} \leq x\right]$ and let $\Phi(x)$ be the distribution function of a standard random variable. There exists a constant $C = 0.5591$ such that for all $x \in \mathbb{R}$ $|F_n(x) - \Phi(x)| \leq \frac{Cr_n}{s_n^3}$ holds.*

In order to apply Theorem 16, we rewrite $S_l^{a_b}$ as a single sum of random variables and check all requirements. Define $Y_i^{p_{a_b}} = \sum_{j=0}^{m(i)-m(i-1)} X_{i,j}^{p_{a_b}}$ for $i \leq l-1$ and $Y_l^{p_{a_b}} = \sum_{j=1}^{m(l)-m(l-1)-1} X_{i,j}^{p_{a_b}}$ with $X_{i,j}^{p_{a_b}}$ independent Bernoulli random variables with parameter p_{a_b}. In analogy, define $Y_i^{q_{a_b}} = \sum_{j=1}^{n(i)-n(i-1)} X_{i,j}^{q_{a_b}}$ for $i \leq l$ independent Bernoulli random variables with parameter q_{a_b} and $Z_i^{a_b} = Y_i^{p_{a_b}} + Y_i^{q_{a_b}}$. This yields that in distribution $S_l^{a_b} \overset{d}{=} \sum_{i=1}^{l} Z_i^{a_b}$. Recall that a Bernoulli random variable with parameter p has expectation p and variance $p(1-p)$. Using linearity of expectation and, because the random variables are independent, linearity of variance, it holds for $i < l$, that the random variables $Z_i^{a_b}$ satisfy (i) and (ii) in Theorem 16, i.e.,

$$\begin{aligned}
\mu_i &= (m(i) - m(i-1))\, p_{a_b} + (n(i) - n(i-1))\, q_{a_b} < \infty, \\
\sigma_i^2 &= (m(i) - m(i-1))\, p_{a_b}(1 - p_{a_b}) + (n(i) - n(i-1))\, q_{a_b}(1 - q_{a_b}) < \infty.
\end{aligned} \tag{3}$$

For $i = l$ it holds that

$$\begin{aligned}
\mu_l &= (m(l) - m(l-1) - 1)\, p_{a_b} + (n(l) - n(l-1))\, q_{a_b} < \infty, \\
\sigma_l^2 &= (m(l) - m(l-1) - 1)\, p_{a_b}(1 - p_{a_b}) + (n(l) - n(l-1))\, q_{a_b}(1 - q_{a_b}) < \infty.
\end{aligned} \tag{4}$$

Furthermore, for $i < l$ it holds that

$$\begin{aligned}
\rho_i &= \mathbb{E}\left[\left|\sum_{j=0}^{m(i)-m(i-1)} X_{i,j}^{p_{a_b}} + \sum_{j=0}^{n(i)-n(i-1)} X_{i,j}^{q_{a_b}} - (m(i) - m(i-1))\, p_{a_b} - (n(i) - n(i-1))\, q_{a_b}\right|^3\right] \\
&\leq \left((m(i) - m(i-1))(1 - p_{a_b}) + (n(i) - n(i-1))(1 - q_{a_b})\right)^3 \leq K < \infty.
\end{aligned} \tag{5}$$

The first inequality in Equation (5) holds, because $X_{i,j}^{p_{a_b}} \leq 1$ and $X_{i,j}^{q_{a_b}} \leq 1$ almost surely. The second inequality follows for some finite $K > 0$ from the assumption $\sup_{i \geq 1} m(i) - m(i-1) < \infty$ and $\sup_{i \geq 1} n(i) - n(i-1) < \infty$. In analogy, for $i = l$ it holds that $\rho_l \leq K < \infty$, which proves that requirement (iii) is fulfilled. Finally, it holds that $s_l^2 = (m(l) - 1)p_{a_b}(1 - p_{a_b}) + n(l)q_{a_b}(1 - q_{a_b})$.

Next, define the sequence $(A_{a_b}(l))_{l\in\mathbb{N}}$ via

$$A_{a_b}(l) = \frac{m\,(l) - 1 - ((m\,(l) - 1)\,p_{a_b} + n\,(l)\,q_{a_b})}{\sqrt{(m(l) - 1)\,p_{a_b}\,(1 - p_{a_b}) + n(l)q_{a_b}\,(1 - q_{a_b})}}$$
$$= \sqrt{m(l)}\frac{\left(1 - \frac{1}{m(l)}\right) - \left(\left(1 - \frac{1}{m(l)}\right)p_{a_b} + \frac{n(l)}{m(l)}\right)q_{a_b})}{\sqrt{\left(1 - \frac{1}{m(l)}\right)p_{a_b}\,(1 - p_{a_b}) + \frac{n(l)}{m(l)}q_{a_b}\,(1 - q_{a_b})}}. \tag{6}$$

Theorem 16 now implies that

$$|\mathbb{P}[S_l^{a_b} \leq m(l) - 1] - \Phi(A_{a_b}(l))| \leq \frac{Cr_l}{s_l^3} \leq \frac{CKl}{\left(s_l^2\right)^{3/2}} = \mathcal{O}(l^{-\frac{1}{2}}). \tag{7}$$

It follows from Equation (6) that $|A_{a_b}(l)| = \Theta(\sqrt{l})$. We now argue that for $a_b > P_b^\infty$ and sufficiently large l, $A_{a_b}(l) < 0$. This follows, if we show that for sufficiently large l

$$\left(1 - \frac{1}{m(l)}\right) - \left(\left(1 - \frac{1}{m(l)}\right)p_{a_b} + \frac{n(l)}{m(l)}q_{a_b}\right) < 0. \tag{8}$$

Given that a_b is strictly greater than the critical value P_b^∞, there exists $\delta > 0$, such that $p_{a_b} + Rq_{a_b} = 1 + \delta$. By adding and subtracting Rq_{a_b} it follows that Equation (8) is equivalent to $1 - \frac{1}{m(l)}(1 - p_{a_b}) - (1 + \delta) + (R - \frac{n(l)}{m(l)})q_{a_b} < 0.$ and therefore to $R - \frac{n(l)}{m(l)} < \frac{1}{q_{a_b}}(\delta + \frac{(1-p_{a_b})}{m(l)})$. Because it is assumed that $|R - \frac{n(l)}{m(l)}| = \mathcal{O}(\frac{1}{l})$, Equation (8) holds for sufficiently large l. This implies that $A_{a_b}(l) = \Theta(-\sqrt{l})$. A standard concentration inequality for a standard Gaussian random variable Z and $x > 0$ using the Chernoff bound gives $\mathbb{P}\,|Z| \geq x] \leq 2\exp\left(\frac{-x^2}{2}\right)$. It follows that $\Phi(A_{a_b}(l)) = \mathcal{O}(e^{-l})$. Equation (7) therefore implies that $\mathbb{P}[S_l^{a_b} \leq m(l) - 1] = \mathcal{O}(l^{-\frac{1}{2}})$. Recalling $\mathbb{P}_{\text{-}b}\,[b \in \mathcal{B}^*(a_b, a_{\text{-}b})] = \mathbb{P}[S_l^{a_b} \geq m(l)] = 1 - \mathbb{P}[S_l^{a_b} \leq m(l) - 1]$ completes the proof. The statements for $a_b < P_b^\infty$ and for sellers can be proven analogously. $\square$

A.5 Proof of Lemma 4

Proof. For regular transaction costs $\Phi_i(a_i, P^*)$, consider the following two auxiliary transaction costs that decompose $\Phi_i(a_i, P^*)$: $\Phi_i^1(a_i, P^*) = \Phi_i(P^*, P^*)$ and $\Phi_i^2(a_i, P^*) = \Phi_i(a_i, P^*) - \Phi_i^1(a_i, P^*)$. Note that Φ_i^1 depends only on the market price and is therefore uninfluenceable and Φ_i^2 is purely uninfluenceable by construction, that is $\Phi_i^2(P^*, P^*) = 0$. $\square$

A.6 Proof of Proposition 5

Proof. **Finite Markets.** The expected utility is of the form $\mathbb{E}_{\text{-}b}\,[u_b(t_b, a_b, a_{\text{-}b})] = t_b{\cdot}\mathbb{P}_{\text{-}b}\,[b \in \mathcal{B}^*(a_b, a_{\text{-}b})] - \mathbb{E}_{\text{-}b}\,[P^*(a_b, a_{\text{-}b})] - \mathbb{E}_{\text{-}b}\,[\Phi_b(a_b, a_{\text{-}b})]$. First, we will show that the expected utility is continuous in a_b.[30]

[30]The same proof strategy for continuity is used in Williams (1991) for the expected utility in a buyer's bid DA without fees in the context of Bayesian Nash equilibria.

The first term $t_b \cdot \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right]$ is continuous by Equation (42). To show that the expected market price is continuous, consider $\mathbb{E}_{\text{-}b}\left[P^*(a_b'', a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[P^*(a_b', a_{\text{-}b})\right]$ for two bids $a_b'' > a_b'$ as $a_b'' - a_b'$ approaches zero. The buyer increases the expected market price when raising their bid if (1) they are involved in trade at a_b'', but not at a_b' or (2) a_b' influences the market price. For (1), the market price is at most a_b'' and for (2) the change in market price is at most $a_b'' - a_b'$. This implies that $\mathbb{E}_{\text{-}b}\left[P^*(a_b'', a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[P^*(a_b', a_{\text{-}b})\right] \leq a_b'' \left(\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b'', a_{\text{-}b})\right] - \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b', a_{\text{-}b})\right]\right) + (a_b'' - a_b')$. The continuity of $\mathbb{E}_{\text{-}b}\left[P^*(\cdot, a_{\text{-}b})\right]$ follows from the continuity of $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(\cdot, a_{\text{-}b})\right]$. For the expected transaction cost, it holds that $\mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b})\right] = \int_{a_b \geq a_{\text{-}b}^{(m)}} \Phi_b(a_b, a_{\text{-}b}) d\mu(a_{\text{-}b})$. By assumption, the map $a_b \mapsto \Phi_b(a_b, a_{\text{-}b})$ is continuous. Therefore, the map $a_b \mapsto \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b})\right]$ is continuous as well. Thus, the expected utility is indeed continuous in a_b. Every bid $a_b < \underline{a}_{S,b}$ results in zero utility, as the buyer is almost surely not involved in trade. For every bid $a_b > t_b^\Phi$, it follows from weak domination ex post that the expected utility for a_b is smaller or equal than for $t_b^\Phi \leq t_b$. If $t_b^\Phi \leq \underline{a}_{S,b}$, then t_b^Φ is a best response with expected utility equal to zero. Otherwise, in order to compute a best response, it is sufficient to consider the interval $\left[\underline{a}_{S,b}, t_b^\Phi\right]$. Because the expected utility is a continuous function on this compact set, it follows from the Extreme Value Theorem that the expected utility attains a maximum. Therefore, a best response exists.

Infinite Markets. It holds that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] = (t_b - P^*(a_b, a_{\text{-}b}) - \Phi_b(a_b, a_{\text{-}b})) \cdot \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right]$. In an infinite market, the market price $P^*(a_b, a_{\text{-}b})$ and the fee $\Phi_b(a_b, a_{\text{-}b})$ are deterministic. By assumption, $\Phi_b(a_b, a_{\text{-}b})$ is continuous in the action a_b. If there is no tie-breaking, it holds that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] = 1$, if $a_b \geq P^*(a)$, and zero otherwise. If $t_b^\Phi < P^*(a)$, then buyer b has no undominated action with positive probability of trade. Therefore t_b^Φ is a best response with expected utility equal to zero. If $t_b^\Phi = P^*(a)$, then the only undominated action with positive probability of trade is t_b^Φ. If this results in a strictly positive utility, then it is a best response. If not, then any bid below $P^*(a)$ is a best response. Therefore, consider the case $t_b^\Phi > P^*(a)$. If there is no tie-breaking, then the trading probability is constant and equal to 1 on the compact set $[P^*(a), t_b^\Phi]$. Note that any bid above t_b^Φ is not a best response by weak domination. By similar arguments as before, the expected utility on this interval is equal to $t_b - P^*(a_b, a_{\text{-}b}) - \Phi_b(a_b, a_{\text{-}b})$ and therefore a continuous function. The Extreme Value Theorem implies again that the maximum is attained and a best response exists. The statement for sellers can be proven analogously. □

A.7 Proof of Theorem 6

Proof. Consider a buyer b and two actions $a_b^1 > a_b^2 > P_b^\infty$ that lead to asymptotically different transaction costs. We will prove that in sufficiently large markets a buyer can improve their expected utility when switching from action a_b^1 to a_b^2. This in turn implies that best responses for two different gross values must lead to asymptotically equal transaction costs. Otherwise, there is a buyer with a certain gross value, who has an incentive to change their action in sufficiently large markets to increase their expected utility. By assumption, there exists $\epsilon > 0$ such that in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[\Phi_b(a_b^1, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}i}\left[\Phi_b(a_b^2, a_{\text{-}b}) | b \in A^*\left(a_b^2, a_{\text{-}b}\right)\right] \geq \epsilon$. We will show that in sufficiently large markets a_b^1 cannot be a best response. By contradiction, assume that it was a best response for some gross value t_b. The expected utility $\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^1, a_{\text{-}b}\right)\right]$ is greater or equal than 0, otherwise it is trivially not a best response. We will prove that in sufficiently large markets

$\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^2, a_{\text{-}b}\right)\right] < 0$, which proves that a_b^1 is not a best response in such markets, because a_b^2 increases the expected utility. Using the law of total expectation, the expected difference in transaction costs can be lower bounded by

$$\begin{aligned} &\mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^2, a_{\text{-}b}\right)\right] \\ &= \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b^1, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b^1, a_{\text{-}b}\right)\right] \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^1, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b^2, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b^2, a_{\text{-}b}\right)\right] \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right] \\ &\geq \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right]\left(\mathbb{E}_{\text{-}b}\left[\Phi_b(a_b^1, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b^2, a_{\text{-}b}\right)\right]\right) \end{aligned} \tag{9}$$

The inequality from the last line follows from the monotonicity of the trading probability, which implies $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^1, a_{\text{-}b})\right] \geq \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right]$. It follows from Proposition 3 that for every γ it holds in sufficiently large markets that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right] \geq 1 - \gamma$. Combining this with the assumption of asymptotically different transaction costs yields that in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^2, a_{\text{-}b}\right)\right] \geq (1-\gamma)\epsilon$. Using Equation (31) in Lemma 18 it holds in sufficiently large markets that $\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^2, a_{\text{-}b}\right)\right] \leq t_b\gamma - (1-\gamma)\epsilon$. If we now choose $\gamma < \frac{\epsilon}{t_b+\epsilon}$, the difference in expected utility is strictly negative, thus contradicting that a_b^1 is a best response. The statement for sellers can be proven analogously. $\square$

A.8 Proof of Theorem 7

Proof. Consider a buyer b with gross value t_b, such that the best response a_b is uniformly bounded away from the critical value. That is, there exists $\delta > 0$, such that in sufficiently large markets either (i) $a_b \leq P_b^\infty - \delta$ or (ii) $a_b \geq P_b^\infty + \delta$. It suffices to prove that for every $\epsilon > 0$ in sufficiently large markets that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] \geq -\epsilon$, which implies that truthfulness is an ϵ-best response. If it holds that $t_b^\Phi \leq a_b$, then $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] = \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right]$, because t_b^Φ weakly dominates every larger bid and a_b is a best response. Therefore, assume that $t_b^\Phi > a_b$. If (i) holds, then Proposition 3 implies that for all $\gamma > 0$ $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] \leq \gamma$ holds in sufficiently large markets. If $\gamma < \frac{\epsilon}{t_b}$, then $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] \leq t_b\gamma \leq \epsilon$. By assumption, it also holds that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] \geq 0$. Combining the inequalities yields the result. If (ii) holds, then $t_b^\Phi \leq t_b$ implies

$$\begin{aligned} &\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] \geq \\ &t_b^\Phi\left(\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right]\right) - \left(\mathbb{E}_{\text{-}b}\left[P^*(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[P^*(a_b, a_{\text{-}b})\right]\right) \\ &- \left(\mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b})\right]\right). \end{aligned} \tag{10}$$

It follows from Theorem 8 that for a DA without transaction costs for every $\epsilon_1 > 0$ truthfulness is an ϵ_1-best response in sufficiently large markets. Assume that a buyer has gross value equal to t_b^Φ. It therefore holds in sufficiently large markets that for any other bid, i.e., also the best response a_b for gross value t_b

$$t_b^\Phi\left(\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right]\right) - \left(\mathbb{E}_{\text{-}b}\left[P^*(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[P^*(a_b, a_{\text{-}b})\right]\right) \geq -\epsilon_1. \tag{11}$$

Using the law of total expectation, the expected difference in transaction costs in Equation (11) is equal to

$$\begin{aligned} &\mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b})\right] \\ &= \mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b}) | b \in \mathcal{B}^*\left(t_b^\Phi, a_{\text{-}b}\right)\right] \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right] - \\ &\mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b, a_{\text{-}b}\right)\right] \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right]. \end{aligned} \tag{12}$$

Because both actions are by assumption greater or equal than $P_b^\infty + \delta$, for every $\gamma > 0$ it holds in sufficiently large markets that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right], \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] \geq 1 - \gamma$. It therefore holds that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right] \leq \gamma$. This implies that in sufficiently large markets

$$\begin{aligned} &\mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b})\right] \leq \\ &\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right]\left(\mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b}) | b \in \mathcal{B}^*\left(t_b^\Phi, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b, a_{\text{-}b}\right)\right]\right) + \\ &\gamma \mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b}) | b \in \mathcal{B}^*\left(t_b^\Phi, a_{\text{-}b}\right)\right]. \end{aligned} \tag{13}$$

Asymptotic uninfluenceability of the transaction costs implies that for every $\epsilon_2 > 0$ the first term in Equation (13) is less or equal than ϵ_2 and for every $\epsilon_3 > 0$ the second term can be chosen to be less or equal than ϵ_3 in sufficiently large markets by choosing $\gamma \leq \frac{\epsilon_3}{\mathbb{E}_{\text{-}b}\left[\Phi_b(t_b^\Phi, a_{\text{-}b}) | A^*\left(b, t_b^\Phi\right)\right]}$. If ϵ_1, ϵ_2 and ϵ_3 are chosen such that their sum is less or equal than ϵ, plugging Equations (11) and (13) into Equation (10) yields that in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] \geq -(\epsilon_1 + \epsilon_2 + \epsilon_3) \geq -\epsilon$, which completes the proof. The statement for sellers can be proven analogously. □

A.9 Proof of Theorem 8

Proof. **Best responses are close to truthfulness.** Consider a buyer b with private type t_b. We will show that there exists a constant $\kappa > 0$, such that $t_b - (1 + \phi_b)\, a_b \leq \kappa q(n, m)$, with $q(m, n) = \max\left\{\frac{1}{n}\left(1 + \frac{m}{n}\right), \frac{1}{m}\left(1 + \frac{n}{m}\right)\right\} = O(\max(m, n)^{-1})$, from which the statement follows. It is proven in Appendix B.3, that a best response a_b necessarily satisfies the first order condition in Equation (39), which implies the following bound:

$$t_b - (1 + \phi_b)\, a_b \leq \frac{(1 + \phi_b)\, k \mathbb{P}_{\text{-}b}\left[a_{m-1,n}^{(m)} \leq a_b \leq a_{m-1,n}^{(m+1)}\right]}{(m-1)\mathbb{P}_{\text{-}b}\left[a_{m-2,n}^{(m-1)} \leq a_b \leq a_{m-2,n}^{(m)}\right] f_{B,b}(a_b)}. \tag{14}$$

It can be proven analogous to Rustichini et al. (1994, Appendix) that

$$\frac{\mathbb{P}_{\text{-}b}\left[a_{m-1,n}^{(m)} \leq a_b \leq a_{m-1,n}^{(m+1)}\right]}{\mathbb{P}_{\text{-}b}\left[a_{m-2,n}^{(m-1)} \leq a_b \leq a_{m-2,n}^{(m)}\right]} \leq 2\left[F_{B,b}(a_b) + \frac{n}{m}\frac{(1 - F_{B,b}(a_b))\, F_{S,b}(a_b)}{1 - F_{S,b}(a_b)}\right]. \tag{15}$$

Defining $\tau_b \equiv 2\max_{x \in [\underline{a}_{S,b}, \overline{a}_{B,b}]}\left\{\frac{F_{B,b}(x)}{f_{B,b}(x)}, \frac{\left(1 - F_{B,b}(x)\right)F_{S,b}(x)}{f_{B,b}(x)\left(1 - F_{S,b}(x)\right)}\right\}$ yields $t_b - (1 + \phi_b)\, a_b \leq \frac{\tau_b k(1 + \phi_b)}{m-1}\left[1 + \frac{n}{m}\right]$. To obtain the bounds in the theorem, note that $\frac{n}{n-1}$ and $\frac{m}{m+1}$ are both less than 2. Setting $\kappa \equiv 2\tau_b k$ proves the statement for buyers. The statement for sellers can be proven analogous.

Truthfulness is an ϵ-best response. First, we estimate the difference in utility when a buyer switches from a bid a_b^1 to a smaller bid a_b^2, i.e., $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^1, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^2, a_{\text{-}b})\right]$. We consider all six cases for the realizations of $a_{\text{-}b}^{(m)}$ and $a_{\text{-}b}^{(m+1)}$ with respect to $a_b^1 > a_b^2$.

		$u_b(t_b, a_b^1, a_{\text{-}b})$	$u_b(t_b, a_b^2, a_{\text{-}b})$
I	$a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)}$	$t_b\text{-}(1+\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	$t_b\text{-}(1+\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
II	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$t_b\text{-}(1+\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	$t_b\text{-}(1+\phi_b)\left(ka_b^2+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
III	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b\text{-}(1+\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	0
IV	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$t_b\text{-}(1+\phi_b)\left(ka_b^1+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	$t_b\text{-}(1+\phi_b)\left(ka_b^2+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
V	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b\text{-}(1+\phi_b)\left(ka_b^1+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	0
VI	$a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^1 \geq a_b^2$	0	0

		$u_b(t_b, a_b^1, a_{\text{-}b}) - u_b(t_b, a_b^2, a_{\text{-}b})$
I	$a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)}$	0
II	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$\text{-}k(1+\phi_b)\left(a_{\text{-}b}^{(m+1)}\text{-}a_b^2\right)$
III	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b\text{-}(1+\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
IV	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$\text{-}k(1+\phi_b)\left(a_b^1\text{-}a_b^2\right)$
V	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b\text{-}(1+\phi_b)\left(ka_b^1+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
VI	$a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^1 \geq a_b^2$	0

We want to lower bound $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^1, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^2, a_{\text{-}b})\right]$. It is therefore sufficient to lower bound the expression in **II** and **IV**, since they are negative and neglect the positive difference in the other cases. In order to prove truthfulness is close to optimal, consider $a_b^1 = t_b^\Phi$ and $a_b^2 = a_b$ a best response. We show that for any $\epsilon > 0$ it holds in sufficiently large finite markets that the difference in expected utility is bounded from below by $-\epsilon$. Because best responses are ϵ-close to truthfulness in sufficiently large finite markets, in such markets $t_b^\Phi - a_b \leq \delta$ holds for all $\delta > 0$. The difference in **II** and **IV** is thus lower bounded by $-k(1+\phi_b)\delta$, and $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^\Phi, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] \leq -k(1+\phi_b)\delta\left(\mathbb{P}\left[\mathbf{II}\right] + \mathbb{P}\left[\mathbf{IV}\right]\right) \leq -2k(1+\phi_b)\delta$. If for a given $\epsilon > 0$, $\delta > 0$ is chosen such that $\delta \leq \frac{\epsilon}{2k(1+\phi_b)}$, it holds in sufficiently large finite markets that t_b^Φ is ϵ-close to a best response a_b. In infinite markets, the expected utility is equal to $\mathbb{E}[u_b(t_b, a_b, a_{\text{-}b})] = t_b - (1+\phi_b)P^*$, if $a_b \geq P^*$, and zero otherwise. If $t_b^\Phi \geq P^*$, then the expected utility is equal to $t_b - (1+\phi_b)P^* > 0$, and therefore a best response. If $t_b^\Phi \leq P^*$, then the expected utility is equal to 0. Because every action $a_b > t_b^\Phi$ is dominated, t_b^Φ is again a best response. Therefore truthfully reporting t_b^Φ is a best response. The statement for sellers can be proven analogously. □

A.10 Proof of Theorem 9

Proof. Consider a buyer b with a gross value t_b, such that $t_b^\Phi > P_b^\infty$. First, we show that in sufficiently large markets an action $a_b^1 < P_b^\infty$ is not a best response. We show that there exists an action $a_b^2 > P_b^\infty$ such that in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^2, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^1, a_{\text{-}b})\right] > 0$, which implies that a_b^1 is not a best response. Because the net value is by assumption continuous and strictly increasing in the gross value, there exists a gross value $t_b' < t_b$, such that $t_b^\Phi > t_b^{\Phi\prime} > P_b^\infty$. Denote the difference between t_b^Φ and $t_b^{\Phi\prime}$ by $\delta > 0$. It holds that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^{\Phi\prime}, a_{\text{-}b})\right] = \mathbb{E}_{\text{-}b}\left[u_b(t_b', t_b^{\Phi\prime}, a_{\text{-}b})\right] + \delta \geq \delta$, because the net value is assumed to be ex-post individually rational. Note that this inequality holds for every market size. It therefore suffices to show that for $a_b^1 < P_b^\infty$ it holds in sufficiently large markets that

$\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] < \delta$. We can upper bound the expected utility by neglecting the expected market price and the expected fee and get that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^1, a_{\text{-}b})\right] \leq t_b \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^1, a_{\text{-}b})\right]$. Proposition 3 implies that for any $\gamma > 0$ it holds in sufficiently large markets that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] \leq \gamma$. If we choose $\gamma < \frac{\delta}{t_b}$, the statement follows. We therefore consider an action a_b that is ϵ-distant to the critical value, that is, there exists $\epsilon > 0$ such that $a_b - P_b^\infty \geq \epsilon$. We will prove that in sufficiently large markets it holds that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty + \epsilon/2, a_{\text{-}b})\right] < 0$, which proves that a_b is not a best response in sufficiently large markets. Therefore, best responses must be ϵ-close, but above the critical value in sufficiently large markets. Using the law of total expectation, the expected difference in transaction cost can be lower bounded by

$$\begin{aligned}
&\mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b\left(P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] = \\
&\mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b, a_{\text{-}b}\right)\right] \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] - \\
&\mathbb{E}_{\text{-}b}\left[\Phi_b(P_b^\infty + \epsilon/2, a_{\text{-}b}) | b \in \mathcal{B}^*\left(P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(P_b^\infty + \epsilon/2, a_{\text{-}b})\right] \geq \\
&\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(P_b^\infty + \epsilon/2, a_{\text{-}b})\right] \left(\mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b}) | b \in \mathcal{B}^*\left(a_b, a_{\text{-}b}\right)\right] - \right. \\
&\left.\mathbb{E}_{\text{-}b}\left[\Phi_b(P_b^\infty + \epsilon/2, a_{\text{-}b}) | b \in \mathcal{B}^*\left(P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right]\right)
\end{aligned} \tag{16}$$

The inequality on the last line holds because the trading probability is monotone, which implies $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] \geq \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(P_b^\infty + \epsilon/2, a_{\text{-}b})\right]$. It follows from Proposition 3 that for every γ it holds in sufficiently large markets that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(P_b^\infty + \epsilon/2, a_{\text{-}b})\right] \geq 1 - \gamma$. Combining this with the assumption of influenceability of the transaction costs yields that there exists $\delta > 0$ such that it holds in sufficiently large markets that $\mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b\left(P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] \geq (1-\gamma)\delta$. Using Equation (31) from Lemma 18, it therefore holds in sufficiently large markets that $\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] \leq t_b\gamma - (1-\gamma)\delta$. If we now choose $\gamma < \delta/t_b + \delta$, the difference is strictly smaller than 0, which proves that a_b is not a best response in sufficiently large markets. The statement for sellers can be proven analogously. $\square$

A.11 Proof of Theorem 10

Proof. To prove that best responses are in an ϵ-neighbourhood of the critical value in sufficiently large markets, consider a buyer b with gross value t_b, such that $t_b^\Phi > P_b^\infty$. It follows analogous to Appendix A.10 that in sufficiently large markets an action $a_b^1 < P_b^\infty$ is not a best response. We therefore consider an action $a_b > P_b^\infty$. That is, there exists $\epsilon > 0$ such that $a_b - P_b^\infty \geq \epsilon$. We will prove that in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty + \epsilon/2, a_{\text{-}b})\right] < 0$, which proves that a_b is not a best response in such markets. Therefore, best responses must be ϵ-close, but above the critical value in sufficiently large markets. For two bids $a_b^1 > a_b^2$ Lemma 18 implies in the presence of a spread fee that $\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^2, a_{\text{-}b}\right)\right] \leq \left(t_b - \phi_b a_b^1\right) \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^1, a_{\text{-}b})\right] - \left(t_b - \phi_b a_b^2\right) \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right]$. Set $a_b^1 = a_b$ and $a_b^2 = P_b^\infty + \epsilon/2$. It follows from Proposition 3 that for any $\gamma > 0$ in sufficiently large markets $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right], \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(P_b^\infty + \epsilon/2, a_{\text{-}b})\right] \geq 1 - \gamma$ and therefore also $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] \leq \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(P_b^\infty + \epsilon/2, a_{\text{-}b})\right] + \gamma$. Thus, it holds in sufficiently large markets that $\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] \leq -\phi_b(1-\gamma)(a_b - (P_b^\infty + \epsilon/2)) + \gamma\left(t_b - \phi_b a_b\right)$.

By assumption, it holds that $a_b - (P_b^\infty + \epsilon/2) \geq \epsilon/2$, which yields
$\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] \leq -\phi_b(1-\gamma)\frac{\epsilon}{2} + \gamma\left(t_b - \phi_b a_b\right) \leq -\phi_b(1-\gamma)\frac{\epsilon}{2} + \gamma t_b$.
If γ is chosen such that $\gamma < \frac{\phi_b \epsilon}{2t_b + \phi_b \epsilon}$ holds, then in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, P_b^\infty + \epsilon/2, a_{\text{-}b}\right)\right] < 0$, which implies that a_b is not a best response in such markets.

Next, we prove that for sufficiently small $\epsilon > 0$, there exist beliefs, such that the critical value is not an ϵ-best response in sufficiently large finite markets. Consider a buyer b with gross value $t_b^\Phi > P_b^\infty$ in a sequence of market environment with $m(l) = l$, $n(l) = l$, $T = [0,1]$ and uniformly distributed beliefs over actions for both buyers and sellers. In this case, the critical value P_b^∞ is equal to $\frac{1}{2}$. By assumption, there exists $\epsilon > 0$, such that $t_b = P_b^\infty + \epsilon$ for $\epsilon > 0$. We will show that in sufficiently large markets $\mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty + \epsilon/4, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty, a_{\text{-}b})\right] > 0$, which proves that P_b^∞ is not a best response. In order to estimate the difference in expected utility for two bids $a_b^1 > a_b^2$, we use a table similar to the one in Appendix A.9:

		$u_b(t_b, a_b^1, a_{\text{-}b})$	$u_b(t_b, a_b^2, a_{\text{-}b})$
I	$a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)}$	$t_b\text{-}\phi_b a_b^1\text{-}(1\text{-}\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	$t_b\text{-}\phi_b a_b^2\text{-}(1\text{-}\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
II	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$t_b\text{-}\phi_b a_b^1\text{-}(1\text{-}\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	$t_b\text{-}\phi_b a_b^2\text{-}(1\text{-}\phi_b)\left(ka_b^2+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
III	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b\text{-}\phi_b a_b^1\text{-}(1\text{-}\phi_b)\left(ka_{\text{-}b}^{(m+1)}+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	0
IV	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$t_b\text{-}\phi_b a_b^1\text{-}(1\text{-}\phi_b)\left(ka_b^1+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	$t_b\text{-}\phi_b a_b^2\text{-}(1\text{-}\phi_b)\left(ka_b^2+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$
V	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b\text{-}\phi_b a_b^1\text{-}(1\text{-}\phi_b)\left(ka_b^1+(1\text{-}k)a_{\text{-}b}^{(m)}\right)$	0
VI	$a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^1 \geq a_b^2$	0	0

Analogously, we consider the difference in utilities:

		$u_b(t_b, a_b^1, a_{\text{-}b}) - u_b(t_b, a_b^2, a_{\text{-}b})$
I	$a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)}$	$-\phi_b\left(a_b^1 - a_b^2\right)$
II	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$-\phi_b\left(a_b^1 - a_b^2\right) - k(1-\phi_b)\left(a_{\text{-}b}^{(m+1)} - a_b^2\right)$
III	$a_b^1 \geq a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b - \phi_b a_b^1 - (1-\phi_b)\left(ka_{\text{-}b}^{(m+1)} + (1\text{-}k)a_{\text{-}b}^{(m)}\right)$
IV	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_b^2 \geq a_{\text{-}b}^{(m)}$	$-\phi_b\left(a_b^1 - a_b^2\right) - k\left((1-\phi_b)\left(a_b^1 - a_b^2\right)\right)$
V	$a_{\text{-}b}^{(m+1)} \geq a_b^1 \geq a_{\text{-}b}^{(m)} \geq a_b^2$	$t_b - \phi_b a_b^1 - (1-\phi_b\left(ka_b^1 + (1\text{-}k)a_{\text{-}b}^{(m)}\right)$
VI	$a_{\text{-}b}^{(m+1)} \geq a_{\text{-}b}^{(m)} \geq a_b^1 \geq a_b^2$	0

In order to obtain a lower bound on the expected difference in utility, we bound all five non-zero terms from below. We set $a_b^1 = P_b^\infty + \epsilon/4$ and $a_b^2 = P_b^\infty$, which implies that there difference is equal to $\epsilon/4$. The expressions in **I**, **II** and **IV** are therefore greater or equal than $-\epsilon/4$. For **III** and **V**, the lower bound $t_b - (P_b^\infty + \epsilon/4) = \frac{3\epsilon}{4}$ holds, because $t_b = P_b^\infty + \epsilon$. Combining these bounds with the probabilities of each event, the following inequality holds:

$$
\begin{aligned}
&\mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty + \epsilon/4, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty, a_{\text{-}b})\right] \geq \\
&-\frac{\epsilon}{4}\mathbb{P}_{\text{-}b}\left[P_b^\infty \geq a_{\text{-}b}^{(m)}\right] + \frac{3\epsilon}{4}\mathbb{P}_{\text{-}b}\left[P_b^\infty + \epsilon/4 \geq a_{\text{-}b}^{(m)} \geq P_b^\infty\right] = \\
&-\frac{\epsilon}{2}\mathbb{P}_{\text{-}b}\left[P_b^\infty \geq a_{\text{-}b}^{(m)}\right] + \frac{3\epsilon}{4}\left(\mathbb{P}_{\text{-}b}\left[a_{\text{-}b}^{(m)} \leq P_b^\infty + \epsilon/4\right] - \mathbb{P}\left[a_{\text{-}b}^{(m)} \leq P_b^\infty\right]\right)
\end{aligned}
\tag{17}
$$

By definition $a_{\text{-}b}^{(m)}$ is the m'th smallest submission in a set of $m-1$ bids and n asks. Since buyer b assumes that those are uniformly distributed and that there are $m(l) = l$ and $n(l) = l$ many buyers and sellers, it follows from order statistics that $a_{\text{-}b}^{(m)} \sim Beta(l,l)$. This distribution is symmetric on $[0,1]$ for every l and therefore at the critical value $P_b^\infty = \frac{1}{2}$, it holds that $\mathbb{P}_{\text{-}b}\left[a_{\text{-}b}^{(m)} \leq P_b^\infty\right] = \frac{1}{2}$. Furthermore, it follows from Proposition 3 that for any $\gamma > 0$ it holds in sufficiently large markets that $\mathbb{P}\left[a_{\text{-}b}^{(m)} \leq P_b^\infty + \epsilon/4\right] \geq 1-\gamma$. It follows that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty + \epsilon/4, a_{\text{-}b})\right] - \mathbb{E}_{\text{-}b}\left[u_b(t_b, P_b^\infty, a_{\text{-}b})\right] \geq -\frac{\epsilon}{8} + \frac{3\epsilon}{4}\left(\frac{1}{2}-\gamma\right)$, which is positive if γ is chosen to be smaller than $\frac{1}{3}$. The statement for sellers can be proven analogously. □

A.12 Proof of Theorem 11

Proof. **Common prior beliefs.** Consider any common prior belief system. It follows from Proposition 14 that the only market performance, that can be achieved with purely influenceable transaction costs, is of the form $(L, 0, 1-L)$ for some loss $L \in [0,1]$ depending on the belief system. Not charging transaction costs, which is a special case of uninfluenceable transaction costs, leads to the fully efficient market outcome $(1,0,0)$. Therefore, not charging any transaction cost weakly dominates any purely influenceable transaction cost for the class of common prior beliefs. It follows from Proposition 12 that the set of achievable market performances is the same for all uninfluenceable transaction costs and independent from the class of beliefs. Corollary 13 then implies that each uninfluenceable transaction cost can be scaled to be optimal. This scaling therefore dominates all uninfluenceable transaction costs, including the special case of no transaction costs. As the latter dominates all purely influenceable transaction costs, transitivity implies that the optimal scaling of any uninfluenceable transaction costs dominates all other transaction costs.

Heterogeneous prior beliefs. We show that for some belief systems, spread fees strictly dominate any uninfluenceable transaction cost. Consider a belief system, such that traders report their gross value.[31] Consider the optimal scaling of any uninfluenceable transaction costs with market performance $(\hat{W}, \hat{R}, \hat{L})$. We distinguish between the following two cases: (i) $\hat{L} = 0$ and (ii) $\hat{L} > 0$. For (i), it follows from Proposition 12 that the unique achievable market performance with $\hat{L} = 0$ is $(1,0,0)$. By assumption, there exists $R \in (0,1]$, such that $U(1-R, R, 0) > U(1,0,0)$. For (ii), by assumption, a Pareto improvement of welfare and revenue leads to an increase in utility, which implies that $U(\hat{W}, \hat{R}+\hat{L}, 0) > U(\hat{W}, \hat{R}, \hat{L})$. Therefore, it suffices to show that for the fixed belief system and any $R \in (0,1]$, there exist spread fees that lead to the market performance $(1-R, R, 0)$. For some choice of R this market performance maximizes the continuous objective function U on the space of all performance triples Δ. Moreover, this market performance is strictly better than the best one achievable by uninfluenceable transaction costs, $(\hat{W}, \hat{R}, \hat{L})$. Consider symmetric spread fees $\phi = \phi_b = \phi_s$. Because traders report their gross values, it follows that the revenue generated from symmetric spread fees is equal to $R = \phi \cdot G^{gross}$ and $W = (1-\phi) \cdot G^{gross}$. As G^{gross} is normalized to 1, it holds that for a symmetric spread fee $\phi \in [0,1]$ and the above constructed beliefs, the market performance is equal to $(\phi, 1-\phi, 0)$. In particular spread fees can thus yield higher utility than any

[31] One such belief system is that all traders belief a critical value equal to their gross value, $P_i^\infty = t_i$, and price-guess. A second one would be that all traders believe that bidding the critical value is not individually rational, and report truthfully.

uninfluenceable transaction cost, which finishes the proof. $\square$

A.13 Proof of Proposition 12

Proof. By Proposition 1, net values are strictly increasing and continuous as a function of the gross value. Thus, net demand $D^{net}(\cdot)$ and net supply $S^{net}(\cdot)$ are continuous. Further, net demand is strictly decreasing on $A_B^{net} = [\underline{a}_B, \overline{a}_B]$ and net supply is strictly increasing on $A_S^{net} = [\underline{a}_S, \overline{a}_S]$. If $\underline{a}_S < \overline{a}_B$, then there exists a unique price P^*_{net}, which lies in the interval $[\underline{a}_S, \overline{a}_B]$ equating net demand and supply, and leading to a strictly positive trading volume Q^*_{net}. Otherwise, if $\underline{a}_S \geq \overline{a}_B$, then the trading volume is equal to zero. Denote by $T_B^* = t_B(\mathcal{B}^*)$ and $T_S^* = t_S(\mathcal{S}^*)$ the set of gross values of all buyers and sellers involved in trade. Because the net values are strictly increasing in the gross value, the following holds: $T_B^* = [\underline{t_b}, \bar{t}]$ and $T_S^* = [\underline{t}, \overline{t_s}]$, that is, the buyers with the highest and the seller with the lowest gross values are involved in trade. $\underline{t_b}$ and $\overline{t_s}$ can be expressed as continuous functions of the trading volume. We show that all market metrics can be expressed via gross demand $D^{gross}(\cdot)$ and gross supply $S^{gross}(\cdot)$, as well as the net trading volume via $\underline{t_b}$ and $\overline{t_s}$:

Gross gains of trade. Note that $G^{real} = \int_{\mathcal{B}^*} (t_b - P^*)\, d\mu_B(b) + \int_{\mathcal{S}^*} (P^* - t_s)\, d\mu_S(s) = \int_{\mathcal{B}^*} t_b \mu_B(b) + \int_{\mathcal{S}^*} t_s d\mu_S(s)$, where the second inequality follows from the fact that $\mu_B(\mathcal{B}^*) = \mu_S(\mathcal{S}^*)$. Moreover, $D^{gross}(P) = \mu_B(\{b \in B : t_b \geq P\}) = \int_T 1_{t_b \geq P} d\mu_B^t(t_b) = \int_T 1_{t_b \geq P} f_B^t(t_b) dt_b$, where $f_B^t(\cdot)$ is the density of μ_B^t. Thus, $\int_{P^*_{gross}}^{\bar{t}} D^{gross}(P)dP = \int_{P^*_{gross}}^{\bar{t}} \left(\int_T 1_{t_b \geq P} f_B^t(t_b) dt_b\right) dP$. Because the function $1_{t_b \geq P} f_B^t(t_b)$ is non-negative and integrable, we can use Fubini's theorem to get $\int_{P^*_{gross}}^{\bar{t}} \left(\int_T 1_{t_b \geq P} f_B^t(t_b) dt_b\right) dP = \int_T \left(\int_{P^*_{gross}}^{\bar{t}} 1_{t_b \geq P} dP\right) f_B^t(t_b) dt_b$. Next, we evaluate the inner integral $\int_{P^*_{gross}}^{\bar{t}} 1_{t_b \geq P} dP$ for fixed t_b. If $t_b \leq P^*_{gross}$, then the integral is equal to zero. If $t_b > P^*_{gross}$, then the integral is equal to $t_b - P^*_{gross}$. The inner integral is therefore equal to $(t_b - P^*_{gross}) 1_{t_b \geq P^*_{gross}}$, which yields $\int_T \left(\int_{P^*_{gross}}^{\bar{t}} 1_{t_b \geq P} dP\right) f_B^t(t_b) dt_b = \int_{t_b \geq P^*_{gross}} (t_b - P^*_{gross}) f_B^t(t) dt_b = \int_{t_b \geq P^*_{gross}} (t_b - P^*_{gross}) d\mu_B^t(t_b)$. For the gross gains of trade, we suppose that traders report the gross value. As there is no tie-breaking, it therefore holds that $\mathcal{B}^* = \{b \in B : t_b \geq P^*_{gross}\}$. This implies that $\int_{t_b \geq P^*_{gross}} (t_b - P^*_{gross}) d\mu_B^t(t_b) = \int_{\{t_b \in T : b \in B^*\}} (t_b - P^*_{gross}) d\mu_B^t(t_b) = \int_{\mathcal{B}^*} \left(t_b - P^*_{gross}\right) d\mu_B(b)$, where the last equality holds, because μ_B^t is the pushforward-measure of μ_B via the map $b \mapsto t_b$. Analogous reasoning yields that $\int_{\underline{t}}^{P^*_{gross}} S^{gross}(P)dP = \int_{\mathcal{S}^*} (P^* - t_s)\, d\mu_S(s)$, and thus

$$G^{gross} = \int_{\underline{t}}^{P^*_{gross}} S^{gross}(P)dP + \int_{P^*_{gross}}^{\bar{t}} D^{gross}(P)dP. \tag{18}$$

Loss. The loss L is given by integrating the gains of trade of all traders, who would have traded, if all traders reported their gross value, but who do not trade given their net value. Recall that $\underline{t_b}$ denotes the lowest gross value, with which a buyer is involved in trade given net demand and supply. Therefore, $L = \int_{\{b \in B : P^*_{gross} \leq t_b \leq \underline{t_b}\}} (t_b - P^*_{net}) d\mu_B(b) + \int_{\{s \in S : \overline{t_s} \leq t_s \leq P^*_{gross}\}} (P^*_{net} - t_s) d\mu_S(s)$. Note that the mass of traders, who would have traded given their gross value, but do not trade given their net value, is equal for both buyers and sellers, which implies

$$\begin{aligned}
&\int_{\{b\in B:P^*_{gross}\le t_b\le \underline{t_b}\}}(t_b-P^*_{net})d\mu_B(b)+\int_{\{s\in S:\overline{t_s}\le t_s\le P^*_{gross}\}}(P^*_{net}-t_s)d\mu_S(s)\\
=&\int_{\{b\in B:P^*_{gross}\le t_b\le \underline{t_b}\}}t_b d\mu_B(b)-\int_{\{s\in S:\overline{t_s}\le t_s\le P^*_{gross}\}}t_s d\mu_S(s) \qquad (19)\\
=&\int_{\{b\in B:P^*_{gross}\le t_b\le \underline{t_b}\}}(t_b-P^*_{gross})d\mu_B(b)+\int_{\{s\in S:\overline{t_s}\le t_s\le P^*_{gross}\}}(P^*_{gross}-t_s)d\mu_S(s)
\end{aligned}$$

Fubini's Theorem implies that $\int_{P^*_{gross}}^{\underline{t_b}} D^{gross}(P)dP = \int_T \left(\int_{P^*_{gross}}^{\underline{t_b}} 1_{t_b\ge P}dP\right) f^t_B(t_b)dt_b$. Again, we evaluate the inner integral $\int_{P^*_{gross}}^{\underline{t_b}} 1_{t_b\ge P}dP$ for fixed t_b. If $t_b \le P^*_{gross}$, then the integral is equal to zero. If $P^*_{gross}\le t_b\le \underline{t_b}$, the integral is equal to $t_b - P^*_{gross}$. If $t_b \ge \underline{t_b}$, the integral is equal to $\underline{t_b} - P^*_{gross}$. Combining these three observations, we get that

$$\int_{P^*_{gross}}^{\underline{t_b}} D^{gross}(P)dP = \int_{P^*_{gross}\le t_b\le \underline{t_b}}(t_b-P^*_{gross})f^t_B(t_b)dt_b+\left(\underline{t_b}-P^*_{gross}\right)\mu^t_B(\{t_b\in T: t_b\ge \underline{t_b}\}). \qquad (20)$$

Finally, it holds that $\mu^t_B(\{t_b\in T: t_b\ge \underline{t_b}\}) = Q^*_{net}$, which implies that

$$\int_{P^*_{gross}}^{\underline{t_b}} D^{gross}(P)dP = \int_{b\in B:P^*_{gross}\le t_b\le \underline{t_b}}(t_b-P^*_{gross})d\mu_B(b)+\left(\underline{t_b}-P^*_{gross}\right)Q^*_{net}. \qquad (21)$$

Analogous reasoning yields that

$$\int^{P^*_{gross}} S^{gross}(P)dP = \int_{\{s\in S:\le t_s\le P^*_{gross}\}}(P^*_{gross}-t_s)d\mu_S(s)+\left(P^*_{gross}-\overline{t_s}\right)Q^*_{net}. \qquad (22)$$

Combining Equation (21) and Equation (22) with Equation (19) yields

$$L = \int_{\overline{t_s}}^{P^*_{gross}} S^{gross}(P)dP + \int_{P^*_{gross}}^{\underline{t_b}} D^{gross}(P)dP - \left(\underline{t_b}-\overline{t_s}\right)Q^*_{net}. \qquad (23)$$

Net gains of trade. By definition, the net gains of trade are given by the formula $G^{net} = G^{gross} - L$. Equations (18) and (23) therefore imply

$$G^{net} = \int_{\underline{t}}^{\overline{t_s}} S^{gross}(P)dP + \left(\underline{t_b}-\overline{t_s}\right)Q^*_{net} + \int_{\underline{t_b}}^{\bar{t}} D^{gross}(P)dP. \qquad (24)$$

Revenue. Note that for regular uninfluenceable transaction costs $R = \int_{\mathcal{B}^*}\Phi_b(a_b, a_{\text{-}b})d\mu_B(b) + \int_{\mathcal{S}^*}\Phi_s(a_s, a_{\text{-}s})d\mu_S(s) = (\Phi_b(P^*_{net})+\Phi_s(P^*_{net}))\cdot Q^*_{net}$. To prove Equation (25), it thus suffices to show that $\Phi_b(P^*_{net}) = \underline{t_b} - P^*_{net}$ and $\Phi_s(P^*_{net}) = P^*_{net} - \overline{t_s}$. Consider a buyer with type $\underline{t_b}$, that is, they are the buyer with the lowest gross value that is involved in trade. Therefore, their net value was equal to the market price P^*_{net}. For regular transaction costs, the net value is such that a trader's utility is equal to zero, if the market price is equal to it, see Proposition 1. Therefore, $u_b(\underline{t_b}, P^*_{net}, P^*_{net}) = 0$, that is $\underline{t_b} - P^*_{net} - \Phi_b(P^*_{net}) = 0$, which implies that $\Phi_b(P^*_{net}) = \underline{t_b} - P^*_{net}$. Analogous reasoning yields that $\Phi_s(P^*_{net}) = P^*_{net} - \overline{t_s}$, which implies

$$R = \left(\underline{t_b} - \overline{t_s}\right) Q^*_{net}. \tag{25}$$

Welfare. The welfare W satisfies the formula $W = G^{net} - R$. Equations (24) and (25) therefore implies

$$W = \int_{\underline{t}}^{\overline{t_s}} S^{gross}(P)dP + \int_{\underline{t_b}}^{\overline{t}} D^{gross}(P)dP. \tag{26}$$

The characterization of market metrics yields that for any regular uninfluenceable transaction cost, the market performance (W, R, L) is fully specified by gross demand, gross supply, and the trading volume Q^*_{net}. Moreover, the map $Q^*_{net} \mapsto (W(Q^*_{net}), R(Q^*_{net}), L(Q^*_{net}))$ is continuous on $[0, Q^*_{gross}]$ with $(W(0), R(0), L(0)) = (0, 0, 1)$ and $(W(Q^*_{gross}), R(Q^*_{gross}), L(Q^*_{gross})) = (1, 0, 0)$. Thus, for any regular uninfluenceable transaction cost, the set of achievable market performances is a subset of the continuous curve $Q^*_{net} \mapsto (W(Q^*_{net}), R(Q^*_{net}), L(Q^*_{net}))$ in the space of all market performances. To show that the set of achievable market performances is equal to this curve, it suffices to prove that for any $V \in [0, Q^*_{gross}]$, there exists a scaling, such that the trading volume is equal to V.

Observation 17. *For any $P \in T$ it holds that $D^{net}_{\gamma_B}(P)$ and $S^{net}_{\gamma_S}(P)$ are continuous and decreasing in γ_B, γ_S. It holds that $D^{net}_0(P) = D^{gross}(P)$ and $S^{net}_0(P) = S^{gross}(P)$ and for sufficiently large $\gamma_B, \gamma_S > 0$ it holds that $D^{net}_{\gamma_B}(P) = S^{net}_{\gamma_S}(P) = 0$.*

Proof of Observation 18. It holds that $D^{net}_{\gamma_B}(P) = \mu_B\left(\{b \in \mathcal{B} : t_b^{\Phi_{\gamma_B}} \geq P\}\right)$. It follows from Proposition 1 that $t_b^{\Phi_{\gamma_B}}$ is the unique solution to the equation $t_b - t_b^{\Phi_{\gamma_B}} - \gamma_B \cdot \Phi_b(t_b^{\Phi_{\gamma_B}}, t_b^{\Phi_{\gamma_B}}) = 0$. The net value for buyers is therefore continuous and strictly decreasing in γ_B, with $t_b^{\Phi_0} = t_b$. Let $t_b^{\Phi_{\gamma_B}, -1}(P)$ denote the gross value, such that the net value corresponding to Φ_{γ_B} is equal to P. It then holds for $P \in T$ that $D^{net}_{\gamma_B}(P) = \mu^t_B\left([t_b^{\Phi_{\gamma_B}, -1}(P), \bar{t}]\right)$. The function $x \mapsto \mu^t_B\left([x, \bar{t}]\right) = \int_x^{\bar{t}} f^t_B(y)dy$ is continuous and strictly decreasing. It therefore follows from the analytical properties of the net value that the map $\gamma_B \mapsto \mu^t_B\left([t_b^{\Phi_{\gamma_B}, -1}(P), \bar{t}]\right)$ is continuous and strictly decreasing. It holds that $t_b^{\Phi_0, -1}(P) = P$, which proves that $D^{net}_0(P) = D^{gross}(P)$. For sufficiently large $\gamma_B > 0$, it holds that $t_b^{\Phi_{\gamma_B}, -1}(P) > \bar{t}$, which implies that the net demand is zero. The statements regarding net supply can be proven analogously. □

Next, fix some $V \in [0, Q^*_{gross}]$. It follows from Observation 17 and the Intermediate Value Theorem that there exists a scaling γ_B and γ_S, such that $D^{net}_{\gamma_B}(P^*_{gross}) = V$ and $S^{net}_{\gamma_S}(P^*_{gross}) = V$. Thus, P^*_{gross} is the market price for the scaled transaction costs with trading volume V, which completes the proof. □

A.14 Proof of Proposition 14

Proof. The realized gains of trade G^{real} and therefore also the loss $L = G^{gross} - G^{real}$ are fully determined by the set of traders involved in trade, which depends on the belief system via price-guessing, but not on the purely heterogeneous transaction cost or its scaling. Therefore, for a given belief system, all achievable market performances have the same loss L.

Common prior beliefs. Suppose that all traders believe in the same critical value P^{∞}. For purely influenceable transaction costs, it follows from Appendix A.1 that the net value equals the gross value. Independent of the transaction cost and scaling, traders act as price-guessers, therefore the achievable market outcome is unique. That is, traders bid the critical value P^{∞}, if it is individually rational, else they report their gross value. Realized demand $D(P)$ is then equal to gross demand $D^{gross}(P)$, if $P \leq P^{\infty}$, and zero else. Realized supply $\mathcal{S}(P)$ is then equal to gross supply $S^{gross}(P)$, if $P \geq P^{\infty}$, and zero else. It follows from the definition of the k-DA that the realized market price is equal to the critical value P^{∞}. All buyers and sellers involved in trade submit an action equal to the market price. Because the transaction cost is purely influenceable, all traders involved in trade pay no transaction cost, as their action is equal to the market price. Hence, the revenue R generated for the platform is equal to zero. The unique achievable market performance for a common prior belief system is thus equal to $(1 - L, 0, L)$ for some loss L, which depends on how far the realized market price is from the true market price. Now, suppose that the belief system is calibrated, that is, the critical value P^{∞} coincides with the gross market price P^{*}_{gross}—the intersection of gross demand and supply. The realized market price is thus equal to the gross market price. Also, the allocation is equal to $\mathcal{B}^{*} = \{b \in \mathcal{B} : a_b \geq P^{\infty}\}$ and $\mathcal{S}^{*} = \{s \in \mathcal{S} : a_s \leq P^{\infty}\}$ and therefore coincides with the allocation that would occur if all traders reported their gross value instead of price-guessing. Thus, the realized gains of trade coincide with the gross gains of trade, even though the action distributions are different. This implies that the loss $L = G^{gross} - G^{real}$ is equal to zero, that is, for calibrated belief systems, the unique achievable market performance is fully efficient, $(W, R, L) = (1, 0, 0)$.

Heterogeneous prior beliefs. Suppose that for a fixed purely heterogeneous transaction cost, the market performance is equal to (W, R, L). As the loss is constant from arguments made at the beginning of the proof, scaling the transaction cost only changes the distribution of welfare W and revenue R, while their sum $G^{real} = W + R$ remains constant. For a transaction cost Φ and scaling $\gamma = (\gamma_B, \gamma_S)$, the revenue is equal to $R = \gamma_B \cdot \int_{\mathcal{B}^{*}} \Phi_b(a_b, a_{\text{-}b}) d\mu_B(b) + \gamma_S \cdot \int_{\mathcal{S}^{*}} \Phi_s(a_s, a_{\text{-}s}) d\mu_S(s)$, and thus continuous in the scaling parameters. It follows that the set of achievable market performances is a line-segment with constant loss and, in case the revenue is equal to zero for any scaling, it is a singleton. Finally, we prove that for any $L \geq 0$, there exists a belief system that leads to loss L.[32] For $a, b \geq 0$, consider a heterogeneous prior belief system such that $P^{\infty}(t_b) = t_b^{\Phi} - a$ and $P^{\infty}(t_s) = t_s^{\Phi} + b$. The remainder of the proof uses the following equivalence: Note that the latter corresponds to the net values of constant transaction costs $\Phi_b = a$ and $\Phi_s = b$. Price-guessing then corresponds

[32] While this can even be proven for a common prior belief system, we provide a proof for a natural heterogeneous prior belief system.

to truthfulness in the presence of those constant transaction costs. For such a strategy-profile, it was proven in Appendix A.13 that the realized gains of trade G^{real} and therefore also the loss $L = G^{gross} - G^{real}$ are a continuous function of the trading volume Q^*, with $L(Q^*_{gross}) = 0$ and $L(0) = 1$. Moreover, the same proof shows that for a suitable choice of $a, b \geq 0$ (in Proposition 12 via the scaling of any regular uninfluenceable transaction cost including constant ones), any trading volume in $[0, Q^*_{gross}]$ and any loss $L \in [0, 1]$ can be achieved. □

B Supporting and auxiliary results

B.1 Explicit formulas

In this section we derive explicit formulas for some of the concepts introduced in the model in Section 3 that are used in the proofs. We will sometimes differentiate between finite markets with m buyers and n sellers and infinite markets with market ratio R. Throughout this section, consider a buyer b with gross value t_b and bid a_b, and a seller s with gross value t_s and ask a_s. Let a denote an action distribution. Recall that in a finite market, $a^{(k)}$ denotes the k'th smallest element in the set of all taken actions.

B.1.1 Involvement in trade

Finite markets. If $a_b < a^{(m)}_{\text{-}b}$, then it is strictly smaller than the $m+1$'st smallest element in the set of all actions a (including a_b) and buyer b is not involved in trade, because their bid is below the market price. If $a_b > a^{(m)}_{\text{-}b}$, then it is at least the $m+1$'st largest element and therefore sufficient to be involved in trade. If $a_b = a^{(m)}_{\text{-}b}$, then the buyer might be subject to tie-breaking. The statement for sellers are analogous. If $a_s > a^{(m)}_{\text{-}s}$, then it is at least the $m+1$'st smallest element in the set of all actions (including a_s) and seller s is not involved in trade, because their ask was above the market price. If $a_s < a^{(m)}_{\text{-}s}$, then it is at most the m'th smallest action and therefore sufficient to be involved in trade. If $a_s = a^{(m)}_{\text{-}s}$s, then the seller might be subject to tie-breaking.

Infinite markets. If there exists no demand excess, then a buyer is involved in trade, if $a_b \geq P^*(a)$. If $a_b < P^*(a)$, then the buyer is not involved in trade. If there exists demand excess, it is generated by bids at $P^*(a)$. If $a_b > P^*(a)$, then the buyer is involved in trade. If $a_b = P^*(a)$, then the buyer might be subject to tie-breaking. The statement for sellers are analogous. If there exists no supply excess, then the seller is involved in trade, if $a_s \leq P^*(a)$. If $a_s > P^*(a)$, then the seller is not involved in trade. If there exists supply excess, it is generated by asks at $P^*(a)$. If $a_s < P^*(a)$, then the seller is involved in trade. If $a_s = P^*(a)$, then the seller might be subject to tie-breaking.

B.1.2 Trading probabilities given beliefs

We can now express the probability of trade, given the beliefs of a trader.

Finite markets. Given the belief that actions are random variables with continuous distribution, tie-breaking is a probability zero event in finite markets. It follows that $\mathbb{P}_{a_{\text{-}b}}[b \in \mathcal{B}^*(a_b, a_{\text{-}b})] =$

$\mathbb{P}_{a_{\text{-}b}}\left[a_b \geq a_{\text{-}b}^{(m)}\right]$ and $\mathbb{P}_{a_{\text{-}s}}\left[s \in \mathcal{S}^*(a_s, a_{\text{-}s})\right] = \mathbb{P}_{a_{\text{-}s}}\left[a_s \leq a_{\text{-}s}^{(m)}\right]$. Explicit formulas for such probabilities are derived in a more general context below (see Equations (42) and (43)). The statement for sellers are analogous.

Infinite markets. If there exists no demand excess at P^*, then $\mathbb{P}_{a_{\text{-}b}}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] = 1$ if $a_b \geq P^*(a)$, and zero otherwise. Suppose that there is strictly positive demand excess. That is $\mu_B(\mathcal{B}_{\geq}(P^*(a))) = Q^*(a) + x$ and $\mu_B(\mathcal{B}_{>}(P^*(a))) = Q^*(a) - y$ for $x > 0$ and $y \geq 0$. Then, $\mathbb{P}_{a_{\text{-}b}}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] = 1$ if $a_b > P^*(a)$, $\mathbb{P}_{a_{\text{-}b}}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] = \frac{y}{x+y}$ if $a_b = P^*(a)$, and zero else. If there exists no supply excess, then then $\mathbb{P}_{a_{\text{-}s}}\left[s \in \mathcal{S}^*(a_s, a_{\text{-}s})\right] = 1$ if $a_s \leq P^*(a)$, and zero otherwise. Suppose that there is strictly positive supply excess. That is $\mu_S(\mathcal{S}_{\leq}(P^*(a))) = Q^*(a) + x$ and $\mu_S(\mathcal{S}_{<}(P^*(a))) = Q^*(a) - y$ for $x > 0$ and $y \geq 0$. Then, $\mathbb{P}_{a_{\text{-}s}}\left[s \in \mathcal{S}^*(a_s, a_{\text{-}s})\right] = 1$ if $a_s < P^*(a)$, $\mathbb{P}_{a_{\text{-}s}}\left[s \in \mathcal{S}^*(a_s, a_{\text{-}s})\right] = \frac{y}{x+y}$ if $a_s = P^*(a)$, and zero else.

B.1.3 Market Price

Finite markets. Recall that $P^*(a) = ka^{(m)} + (1-k)a^{(m+1)}$, see Rustichini et al. (1994); Jantschgi et al. (2022). Interpreting the market price as a function of a single action yields

$$P^*\left(a_b, a_{\text{-}b}\right) = \begin{cases} (1\text{-}k)a_{\text{-}b}^{(m)} + ka_b & \text{if } a_{\text{-}b}^{(m)} \leq a_b \leq a_{\text{-}b}^{(m+1)}, \\ (1\text{-}k)a_{\text{-}b}^{(m)} + ka_{\text{-}b}^{(m+1)} & \text{else.} \end{cases} \tag{27}$$

$$P^*\left(a_s, a_{\text{-}s}\right) = \begin{cases} (1\text{-}k)a_s + ka_{\text{-}s}^{(m)} & \text{if } a_{\text{-}s}^{(m\text{-}1)} \leq a_s \leq a_{\text{-}s}^{(m)}, \\ (1\text{-}k)a_{\text{-}s}^{(m\text{-}1)} + ka_{\text{-}s}^{(m)} & \text{else.} \end{cases} \tag{28}$$

Note that $P^*\left(a_b, a_{\text{-}b}\right)$ depends only on $a_{\text{-}b}^{(m)}$ and $a_{\text{-}b}^{(m+1)}$, and $P^*\left(a_s, a_{\text{-}s}\right)$ depends only on $a_{\text{-}s}^{(m\text{-}1)}$ and $a_{\text{-}s}^{(m)}$. In some proofs, this dependence will be of importance and we will, for example, write $P^*\left(a_b, a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)}\right)$ instead of $P^*\left(a_b, a_{\text{-}b}\right)$. In addition, for a trader i, we will in some proofs consider $\tilde{P}^*\left(a_i, a_{\text{-}i}\right)$, which is equal to the market price, if i is involved in trade, and zero otherwise.

Infinite markets. In an infinite market, a single trader cannot influence the market price. It therefore holds for a trader i and for all actions a_i and a'_i that $P^*(a_i, a_{\text{-}i}) = P^*(a'_i, a_{\text{-}i})$. By abuse of notation, we will in some proofs write $P^*(a_{\text{-}i})$.

B.1.4 Utility functions

For a buyer the utility of being involved in trade is equal to the difference between their gross value and the market price minus the additional transaction cost: $u_b\left(t_b, a_b, a_{\text{-}b}\right) = t_b\text{-}P^*\left(a_b, a_{\text{-}b}\right)\text{-}\Phi_b\left(a_b, a_{\text{-}b}\right)$. In finite markets, it holds that

$$\begin{aligned} \mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] &= \int_{\{a_b \geq a_{\text{-}b}^{(m)}\}} \left(t_b\text{-}P^*\left(a_b, a_{\text{-}b}\right)\text{-}\Phi_b\left(a_b, a_{\text{-}b}\right)\right) d\mu_b\left(a_{\text{-}b}\right) = \\ t_b \cdot \mathbb{P}_{\text{-}b}[b \in \mathcal{B}^*(a_b, a_{\text{-}b})] &- \int_{[\underline{a}_{S,b}, \overline{a}_{S,b}]^2} \tilde{P}^*\left(a_b, a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)}\right) d\mu_b(a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)}) - \mathbb{E}_{\text{-}b}\left[\Phi_b(a_b, a_{\text{-}b})\right], \end{aligned} \tag{29}$$

where $\mu_b(a_{\text{-}b})$ denotes the distribution of $a_{\text{-}b}$ according to the beliefs of buyer b. Note that both $a_{\text{-}b}^{(m)}$ and $a_{\text{-}b}^{(m+1)}$ have support in $[\underline{a}_{S,b}, \overline{a}_{S,b}]$. That is because $a_{\text{-}b}$ consists of $m-1$ bids and n asks. So there must be at least one ask below or equal to $a_{\text{-}b}^{(m)}$. In infinite markets, we have that $\mathbb{E}_{\text{-}b}[u_b(t_b, a_b, a_{\text{-}b})] = (t_b - P^*(a_b, a_{\text{-}b}) - \Phi_b(a_b, a_{\text{-}b}))\,\mathbb{P}_{\text{-}b}[b \in \mathcal{B}^*(a_b, a_{\text{-}b})]$.

For a seller the utility of being involved in trade is equal to the difference between the market price and their gross value minus the additional transaction cost: $u_s(t_s, a_s, a_{\text{-}s}) = P^*(a_s, a_{\text{-}s})\text{-}t_s\text{-}\Phi_s(a_s, a_{\text{-}s})$. In finite markets, it holds that

$$\begin{gathered}\mathbb{E}_{\text{-}s}[u_s(t_s, a_s, a_{\text{-}s})] = \int\limits_{\{a_s \leq a_{\text{-}s}^{(m)}\}} (P^*(a_s, a_{\text{-}s})\text{-}t_s - \Phi_s(a_s, a_{\text{-}s}))\, d\mu_s(a_{\text{-}s}) = \\ \int\limits_{[\underline{a}_{B,s}, \overline{a}_{B,s}]^2} \tilde{P}^*\left(a_s, a_{\text{-}s}^{(m\text{-}1)}, a_{\text{-}s}^{(m)}\right) d\mu_s(a_{\text{-}s}^{(m\text{-}1)}, a_{\text{-}s}^{(m)}) - t_s \cdot \mathbb{P}_{\text{-}s}[s \in \mathcal{S}^*(a_s, a_{\text{-}s})] - \mathbb{E}_{\text{-}s}[\Phi_s(a_s, a_{\text{-}s})]\,,\end{gathered} \tag{30}$$

where $\mu_s(a_{\text{-}s})$ denotes the distribution of $a_{\text{-}s}$ according to the beliefs of a seller s. Note that both $a_{\text{-}s}^{(m\text{-}1)}$ and $a_{\text{-}s}^{(m)}$ have support in $[\underline{a}_{B,s}, \overline{a}_{B,s}]$. In infinite markets, it holds that $\mathbb{E}_{\text{-}s}[u_s(t_s, a_s, a_{\text{-}s})] = (P^*(a_s, a_{\text{-}s}) - t_s - \Phi_s(a_s, a_{\text{-}s}))\,\mathbb{P}_{\text{-}s}[s \in \mathcal{S}^*(a_s, a_{\text{-}s})]$.

Lemma 18. *For bids $a_b^1 > a_b^2$ and for asks $a_s^1 < a_s^2$ it holds that*

$$\begin{gathered}\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^2, a_{\text{-}b}\right)\right] \leq \\ t_b\left(\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^1, a_{\text{-}b})\right] - \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right]\right) - \left(\mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^2, a_{\text{-}b}\right)\right]\right).\end{gathered} \tag{31}$$

$$\begin{gathered}\mathbb{E}_{\text{-}s}\left[u_s\left(t_s, a_s^1, a_{\text{-}s}\right)\right] - \mathbb{E}_{\text{-}s}\left[u_s\left(t_s, a_s^2, a_{\text{-}s}\right)\right] \\ \leq 2\overline{a}_{B,s}\left(1 - \mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s, a_{\text{-}s})\right]\right) - t_s\left(\mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s^1, a_{\text{-}s})\right] - \mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s^2, a_{\text{-}s})\right]\right) \\ - \left(\mathbb{E}_{\text{-}s}\left[\Phi_s\left(a_s^1, a_{\text{-}s}\right)\right] - \mathbb{E}_{\text{-}s}\left[\Phi_s\left(a_s^2, a_{\text{-}s}\right)\right]\right).\end{gathered} \tag{32}$$

Proof. Recall that $\tilde{P}^*$ denotes the market price, if a trader is involved in trade, and zero otherwise. For a buyer b with private type t_b, Equation (29) yields that

$$\begin{gathered}\mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[u_b\left(t_b, a_b^2, a_{\text{-}b}\right)\right] = t_b\left(\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^1, a_{\text{-}b})\right] - \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b^2, a_{\text{-}b})\right]\right) - \\ \int_{[\underline{a}_{S,b}, \overline{a}_{S,b}]^2} \left(\tilde{P}^*\left(a_b^1, a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)}\right) - \tilde{P}^*\left(a_b^2, a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)}\right)\right) d\mu(a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)}) - \\ \left(\mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^1, a_{\text{-}b}\right)\right] - \mathbb{E}_{\text{-}b}\left[\Phi_b\left(a_b^2, a_{\text{-}b}\right)\right]\right).\end{gathered} \tag{33}$$

Note that the integral in the difference above is non-negative, because $\tilde{P}^*(a_b, a_{\text{-}b}^{(m)}, a_{\text{-}b}^{(m+1)})$ is increasing in a_b for fixed $a_{\text{-}b}^{(m)}$ and $a_{\text{-}b}^{(m+1)}$. Equation (31) follows by neglecting the term corresponding to the change in expected market price. For a seller s with private type t_s, Equation (30) yields

$$\begin{gathered}\mathbb{E}_{\text{-}s}\left[u_s\left(t_s, a_s^1, a_{\text{-}s}\right)\right] - \mathbb{E}_{\text{-}s}\left[u_s\left(t_s, a_s^2, a_{\text{-}s}\right)\right] = \\ \int_{[\underline{a}_{B,s}, \overline{a}_{B,s}]^2} \left(\tilde{P}^*\left(a_s^1, a_{\text{-}s}^{(m\text{-}1)}, a_{\text{-}s}^{(m)}\right) - \tilde{P}^*\left(a_s^2, a_{\text{-}s}^{(m\text{-}1)}, a_{\text{-}s}^{(m)}\right)\right) d\mu(a_{\text{-}s}^{(m\text{-}1)}, a_{\text{-}s}^{(m)}) - \\ t_s\left(\mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s^1, a_{\text{-}s})\right] - \mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s^2, a_{\text{-}s})\right]\right) - \left(\mathbb{E}_{\text{-}s}\left[\Phi_s\left(a_s^1, a_{\text{-}s}\right)\right] - \mathbb{E}_{\text{-}s}\left[\Phi_s\left(a_s^2, a_{\text{-}s}\right)\right]\right).\end{gathered} \tag{34}$$

$t_s\left(\mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s^1, a_{\text{-}s})\right] - \mathbb{P}_{\text{-}s}\left[s \in \mathcal{S}^*(a_s^2, a_{\text{-}s})\right]\right) \geq 0$ holds, because the trading probability is decreasing for a seller in their ask. To see that the integral in Equation (34) is bounded from above by $2t_s\left(1-\mathbb{P}_{\text{-}s}\left[s\in\mathcal{S}^*(a_s^2,a_{\text{-}s})\right]\right)$, we split up the integral into all six possible cases for the realizations of $a_{\text{-}s}^{(m)}$ and $a_{\text{-}s}^{(m\text{-}1)}$ with respect to $a_s^1<a_s^2$. which is shown in the following table.[33]

		$\tilde{P^*}\left(a_s^1,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)$	$\tilde{P^*}\left(a_s^2,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)$
I	$a_{\text{-}s}^{(m)}\geq a_{\text{-}s}^{(m\text{-}1)}\geq a_s^2\geq a_s^1$	$ka_{\text{-}s}^{(m)}$+(1-k)$a_{\text{-}s}^{(m\text{-}1)}$	$ka_{\text{-}s}^{(m)}$+(1-k)$a_{\text{-}s}^{(m\text{-}1)}$
II	$a_{\text{-}s}^{(m)}\geq a_s^2\geq a_{\text{-}s}^{(m\text{-}1)}\geq a_s^1$	$ka_{\text{-}s}^{(m)}$+(1-k)$a_{\text{-}s}^{(m\text{-}1)}$	$ka_{\text{-}s}^{(m)}$+(1-k)a_s^2
III	$a_s^2\geq a_{\text{-}s}^{(m)}\geq a_{\text{-}s}^{(m\text{-}1)}\geq a_s^1$	$ka_{\text{-}s}^{(m)}$+(1-k)$a_{\text{-}s}^{(m\text{-}1)}$	0
IV	$a_{\text{-}s}^{(m)}\geq a_s^2\geq a_s^1\geq a_{\text{-}s}^{(m\text{-}1)}$	$ka_{\text{-}s}^{(m)}$+(1-k)a_s^1	$ka_{\text{-}s}^{(m)}$+(1-k)a_s^2
V	$a_s^2\geq a_{\text{-}s}^{(m)}\geq a_s^1\geq a_{\text{-}s}^{(m\text{-}1)}$	$ka_{\text{-}s}^{(m)}$+(1-k)a_s^1	0
VI	$a_s^2\geq a_s^1\geq a_{\text{-}s}^{(m)}\geq a_{\text{-}s}^{(m\text{-}1)}$	0	0

For **I**, **II**, **IV** and **VI**, the difference between $\tilde{P^*}\left(a_s^1,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)$ and $\tilde{P^*}\left(a_s^2,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)$ is less or equal than 0. It follows that

$$\begin{aligned}\int_{[\underline{a}_{B,s},\overline{a}_{B,s}]^2}&\left(\tilde{P^*}\left(a_s^1,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)-\tilde{P^*}\left(a_s^2,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)\right)d\mu(a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)})\leq\\&\int_{\mathbf{III}}(ka_{\text{-}s}^{(m)}+(1-k)a_{\text{-}s}^{(m\text{-}1)})d\mu_s^*(a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)})\\&+\int_{\mathbf{V}}(ka_{\text{-}s}^{(m)}+(1-k)a_s^1)d\mu(a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}).\end{aligned}\tag{35}$$

Because both integrands in Equation (41) are less or equal than $\overline{a}_{S,s}$, it follows that

$$\begin{aligned}&\int_{[\underline{a}_{B,s},\overline{a}_{B,s}]^2}\left(\tilde{P^*}\left(a_s^1,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)-\tilde{P^*}\left(a_s^2,a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)}\right)\right)d\mu(a_{\text{-}s}^{(m\text{-}1)},a_{\text{-}s}^{(m)})\\&\leq\overline{a}_{S,s}\mathbb{P}[\mathbf{III}]+\overline{a}_{S,s}\mathbb{P}[\mathbf{V}]\leq2\overline{a}_{S,s}\mathbb{P}[a_{\text{-}s}^{(m)}<a_s^2]=2\overline{a}_{S,s}\left(1-\mathbb{P}_{\text{-}s}\left[(s,a_s^2\in\mathcal{S}^*\right]\right),\end{aligned}\tag{36}$$

which finishes the proof. □

B.2 In-the-market and out-of-the-market traders

We sometimes focus on *in-the-market* traders with gross values t_i such that $t_i^\Phi\prec P_i^\infty$. Traders with such gross values are able to submit individually rational actions that make them likely to be involved in trade when the market is sufficiently large. By contrast, for an *out-of-the-market* trader, that is, one with gross value $t_i^\Phi\succ P_i^\infty$, the probability of trade, when acting individually rationally, vanishes in large markets. Observe that bidding the critical value P_i^∞ is individually rational for in-the-market traders but not for out-of-the-market traders.

Proposition 19 (For out-of-the-market traders, truthfulness is close to optimal)**.** *If bidding the critical value P_i^∞ is not individually rational for trader i, then for every $\epsilon>0$, in sufficiently large markets, truthfulness is an ϵ-best response.*

Proof. Consider a buyer b with gross value t_b, such that $t_b^\Phi<P_b^\infty$. A best response a_b with $a_b\leq t_b^\Phi$ must exist. That is because if there is a best response a_b with $a_b>t_b^\Phi$, the expected utilities must be

[33]Different to $\tilde{P^*}_b\left(a_b,y,z\right)$ it holds that $\tilde{P^*}_s\left(a_s,y,z\right)$ is not increasing in a_s for fixed y and z.

equal, as the net value dominates all larger actions, proving that t_b^Φ is a best response as well. By the monotonicity of the trading probability, it then holds that $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(a_b, a_{\text{-}b})\right] \leq \mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right]$. For all $\gamma > 0$, it holds by Proposition 3 that in sufficiently large markets $\mathbb{P}_{\text{-}b}\left[b \in \mathcal{B}^*(t_b^\Phi, a_{\text{-}b})\right] \leq \gamma$. The expected utility is upper bounded by neglecting the payment of market price and fee, that is the gross value times the probability of trade: $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right] \leq t_b\gamma$. Choose $\gamma \leq \frac{\epsilon}{t_b}$. This implies that in sufficiently large markets, the expected utility of a best response is upper bounded by ϵ. The expected utility of truthfulness is non-negative by assumption. This implies that truthfulness is an ϵ-best response. The statement for sellers can be proven analogously. $\square$

B.3 Strategic incentives for price and spread fees

This section contains a detailed discussion of the opposing strategic incentives for price and spread fees in finite markets: (i) Utility when trading, versus (ii) probability of trading.[34] Recall that a trader i believes that actions are distributed in intervals $A_{B,i} = [\underline{a}_{B,i}, \overline{a}_{B,i}]$ and $A_{S,i} = [\underline{a}_{S,i}, \overline{a}_{B,i}]$ with the assumption that $\overline{a}_{S,i} \geq \overline{a}_{B,i} > t_i^\Phi > \underline{a}_{S,i} \geq \underline{a}_{B,i}$. Consider a buyer b with action a_b. We can omit the analysis of $a_b > \overline{a}_{B,b}$ and $a_b < \underline{a}_{S,b}$; for the first, such an action is by assumption not individually rational and strictly dominated by t_b^Φ, for the second, any action below $\underline{a}_{S,b}$ has probability of trade equal to 0, because no seller is believed to submit an action below it. Therefore, the expected utility at such a bid is equal to 0. We therefore consider $a_b \in [\underline{a}_{S,b}, \overline{a}_{B,b}]$. As the market price depends only on a_b, $a_{\text{-}b}^{(m)}$ and $a_{\text{-}b}^{(m+1)}$. For ease of notation, let $y = a_{\text{-}b}^{(m)}$ and $z = a_{\text{-}b}^{(m+1)}$ and denote by $e(y, z)$ the joint density of y and z given the beliefs of buyer b.

Price fees. The expected utility $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right]$ of a buyer is of the form

$$\int_{a_b}^{\overline{a}_{S,i}} \int_{\underline{a}_{S,b}}^{a_b} (t_b\text{-}\,(1{+}\phi_b)\,(ka_b{+}\,(1\text{-}k)\,y))\,e(y,z)dydz + \int_{\underline{a}_{S,b}}^{a_b} \int_{\underline{a}_{S,b}}^{z} (t_b\text{-}\,(1{+}\phi_b)\,(kz{+}\,(1\text{-}k)\,y))\,e(y,z)dydz. \tag{37}$$

The expected utility is continuously differentiable as a function of a_b over the interval $[\underline{a}_{S,b}, \overline{a}_{S,b}]$. Straightforward computation using Leibniz's rule for differentiation under the integral sign yields $\frac{d\mathbb{E}_{\text{-}b}[u_b(t_b,a_b,a_{\text{-}b})]}{da_b} = (t_b - (1+\phi_b)\,a_b)\,f_y(a_b) - (1+\phi_b)\,k\mathbb{P}_{\text{-}b}\left[y \leq a_b \leq z\right]$, where $f_y(a_b)$ denotes the density function of y. If $a_b \in (\underline{a}_{S,b}, \overline{a}_{S,b})$ maximizes the expected utility, then the first order condition $\frac{d\mathbb{E}_{\text{-}b}[u_b(t_b,a_b,a_{\text{-}b})]}{da_b} = 0$ holds. $f_y(a_b)$ is equal to $\frac{d\mathbb{P}_{\text{-}b}[y \leq a_b]}{da_b}$. A formula for $\mathbb{P}_{\text{-}b}[y \leq a_b]$ is stated below in the section on first order conditions. Therefore, we can explicitly state the first order condition in terms of distribution and density functions, see Equation (39) below. The first order condition for a seller can be derived in analogy, see Equation (40) below.

Spread fees. The expected utility $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b, a_{\text{-}b})\right]$ of a buyer is of the form

$$\int_{a_b}^{\overline{a}_{S,b}} \int_{\underline{a}_{S,b}}^{a_b} (t_b\text{-}\phi_b a_b\text{-}\,(1\text{-}\phi_b)\,(ka_b{+}\,(1\text{-}k)\,y))\,e(y,z)dydz + \int_{\underline{a}_{S,b}}^{a_b} \int_{\underline{a}_{S,b}}^{z} (t_b\text{-}\phi_b a_b\text{-}\,(1\text{-}\phi_b)\,(kz{+}\,(1\text{-}k)\,y))\,e(y,z)dydz. \tag{38}$$

[34] This section is closely related to methods used in Rustichini et al. (1994) to analyze strategic incentives in k-DAs without transaction costs.

The expected utility is continuously differentiable as a function of a_b over the interval $[\underline{a}_{S,b}, \overline{a}_{S,b}]$. Straightforward computation using Leibniz's rule for differentiation under the integral sign yields $\frac{d\mathbb{E}_{\text{-}b}[u_b(t_b,a_b,a_{\text{-}b}]}{da_b} = (t_b - a_b)\, f_y(a_b) - \phi_b \mathbb{P}_{\text{-}b}\,[y \leq a_b] - (1-\phi_b)\, k\mathbb{P}_{\text{-}b}\,[y \leq a_b \leq x]$ where $f_y(a_b)$ denotes the density function of y. If $a_b \in (\underline{a}_{S,b}, \overline{a}_{S,b})$ maximizes the expected utility, then the first order condition $\frac{d\mathbb{E}_{\text{-}b}[u_b(t_b,a_b,a_{\text{-}b})]}{da_b} = 0$. holds. $f_y(a_b)$ is equal to $\frac{d\mathbb{P}_{\text{-}b}[y \leq a_b]}{da_b}$. A formula for $\mathbb{P}_{\text{-}b}[y \leq a_b]$ is stated below. Therefore, we can explicitly state the first order condition in terms of distribution and density functions, see Equation (39) below. The first order condition for a seller can be derived in analogy, see Equation (40) below.

First Order Conditions. Define $a_{i,j}$ as an action distribution for i buyers and j sellers. In this notation, a as defined in Section 3 corresponds to $a_{m,n}$ and for any buyer b and seller s, $a_{\text{-}b}$ and $a_{\text{-}s}$ correspond to $a_{m-1,n}$ and $a_{m,n-1}$. Denote again by $a^{(l)}_{i,j}$ its l'th smallest element.
We say that an action a_b satisfies the *buyer's first order condition* for gross value t_b if

$$\left.\begin{matrix} (t_b\text{-}\,(1{+}\phi_b)\,a_b) \\[1em] (t_b\text{-}a_b) \end{matrix}\right\} \cdot \left(n\mathbb{P}_{\text{-}b}\left[a^{(m\text{-}1)}_{m\text{-}1,n\text{-}1} \leq a_b \leq a^{(m)}_{m\text{-}1,n\text{-}1}\right] f_{S,b}(a_b) + (m\text{-}1)\,\mathbb{P}_{\text{-}b}\left[a^{(m\text{-}1)}_{m\text{-}2,n} \leq a_b \leq a^{(m)}_{m\text{-}2,n}\right] f_{B,b}(a_b)\right) = \begin{cases} (1+\phi_b)\, k\mathbb{P}_{\text{-}b}\left[a^{(m)}_{m\text{-}1,n\text{-}1} \leq a_b \leq a^{(m+1)}_{m\text{-}1,n}\right] & \text{for price fees} \\ \phi_b \mathbb{P}_{\text{-}b}\left[a^{(m)}_{m,n\text{-}1} \leq a_b\right] + (1\text{-}\phi_b)\, k\mathbb{P}_{\text{-}b}\left[a^{(m)}_{m\text{-}1,n} \leq a_b \leq a^{(m+1)}_{m\text{-}1,n}\right] & \text{for spread fees} \end{cases} . \tag{39}$$

We say that an action a_s satisfies the *seller's first order condition* for gross value t_s if

$$\left.\begin{matrix} ((1\text{-}\phi_s)\,a_s\text{-}t_s) \\[1em] (a_s\text{-}t_s) \end{matrix}\right\} \cdot \left((n\text{-}1)\,\mathbb{P}_{\text{-}s}\left[a^{(m\text{-}1)}_{m,n\text{-}2} \leq a_s \leq a^{(m)}_{m,n\text{-}2}\right] f_{S,s}(a) + m\mathbb{P}_{\text{-}s}\left[a^{(m\text{-}1)}_{m\text{-}1,n\text{-}1} \leq a_s \leq a^{(m)}_{m\text{-}1,n\text{-}1}\right] f_{B,s}(a)\right) = \begin{cases} (1\text{-}\phi_s)\,(1\text{-}k)\,\mathbb{P}_{\text{-}s}\left[a^{(m\text{-}1)}_{m,n\text{-}1} \leq a_s \leq a^{(m)}_{m,n\text{-}1}\right] & \text{for price fees} \\ \phi_s \mathbb{P}_{\text{-}s}\left[a^{(m)}_{m,n\text{-}1} \geq a_s\right] + (1\text{-}\phi_s)\,(1\text{-}k)\,\mathbb{P}_{\text{-}s}\left[a^{(m\text{-}1)}_{m,n\text{-}1} \leq a_s \leq a^{(m)}_{m,n\text{-}1}\right] & \text{for spread fees} \end{cases} . \tag{40}$$

Interpretation of the first order condition. The first order condition balances between the probability of trade and the utility when trading. In particular, an incremental increase Δa_b in a buyer's bid has two opposing effects: If the bid a_b does not include the buyer amongst those who trade, then by increasing it to $a_b + \Delta a_b$, the buyer may surpass other actions and be involved in trade. If the bid a_b is sufficient to include the buyer in trade, then increasing their bid by Δa_b may lead to an increase in market price and their fee. In Equation (39), the left-hand side of the equation describes the gain from increasing one's trading probability. The sum in brackets times Δa_b is the probability that the buyer enters the set of buyers who trade as they incrementally raise their bid by Δa_b. The first term in the sum is the marginal probability of acquiring an item by passing a seller's offer and the second term is the marginal probability of acquiring an item by passing another buyer's bid. For a price fee the profit from such a trade is between $t_b - (1 + \phi_b)a_b$ and $t_b - (1 + \phi_b)a_b - (1 + \phi_b)\Delta a_b$. Therefore, the marginal expected profit for a buyer who raises their bid is $t_b - (1 + \phi_b)a_b$ times the term in the brackets. In analogy, for a spread fee the marginal expected profit for a buyer who raises their bid is $t_b - \phi_b a_b$ times the term in the brackets. Next, in

Equation (39), the right-hand side of the equation describes the buyer's marginal execpted loss from increasing their bid above a_b. $\mathbb{P}_{\text{-}b}\left[a^{(m)}_{m\text{-}1,n} \leq a_b \leq a^{(m+1)}_{m\text{-}1,n}\right]$ is the probability that a buyer who increases their bid by Δa_b increases the market price by $k(1+\phi_b)\Delta a_b$ for a price fee and by $k(1-\phi_b)\Delta a_b$ for a spread fee. Additionally, for a spread fee $\mathbb{P}_{\text{-}b}\left[a^{(m)}_{m\text{-}1,n} \leq a_b\right]$ is the probability that a buyer who increases their bid by Δa_b increases the part of the charged fee depending on their bid by $\phi_b \Delta a_b$. The interpretation for a seller is symmetric and thus omitted.

Probabilities in the first order conditions. In this section we derive explicit formulas for the probabilities arising in the first order conditions in Equations (39) and (40), that are also used in the proof of Theorem 8 in Appendix A.9. Instead of deriving expressions for all different probabilities, note that for general n, m, l all of them can be expressed as one of the following three probabilities for different n, m, l: (i) $\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \leq a_i \leq a^{(l+1)}_{m,n}\right]$, (ii) $\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \leq a_i\right]$ and (iii) $\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \geq a_i\right]$. For (i) it is the probability that action a_i lies between the l'th and $l+1$'st smallest element in a set of m bids and n asks. The probability that another buyer submits an action smaller or equal a_i is $F^a_{B,i}(a_i)$. The probability that a buyer submits an action greater or equal a_i is therefore $1 - F^a_{B,i}(a_i)$. Replace $F^a_{B,i}$ by $F^a_{S,i}$ for sellers. The event that exactly l bids and asks are below a_i can be split up in the following way: Suppose that i buyers and j sellers bid and offer less or equal than a_i. Note that $i+j$ is equal to l. Assuming that there are m buyers and n sellers in total, this means that exactly $m-i$ buyers and $n-j$ sellers bid and offer more than a_i. Selecting i buyers and j sellers, the probability that exactly $i+j = l$ bids and offers are below or equal to a_i is $F^a_{B,i}(a_i)^i \, F^a_{S,i}(a_i)^j \left(1\text{-}F^a_{B,i}(a_i)\right)^{m\text{-}i} \left(1\text{-}F^a_{S,i}(a_i)\right)^{n\text{-}j}$, because the actions of traders are assumed to be independent. There are $\binom{m}{i}$ possibilities to choose i buyers and $\binom{n}{j}$ possibilities to choose j sellers. Therefore, the total probability that exactly l traders submit below a_i is equal to

$$\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \leq a_i \leq a^{(l+1)}_{m,n}\right] = \sum_{\substack{i+j=l \\ 0 \leq i \leq m \\ 0 \leq j \leq n}} \binom{m}{i}\binom{n}{j} F^a_{B,i}(a_i)^i \, F^a_{S,i}(a_i)^j \left(1\text{-}F^a_{B,i}(a_i)\right)^{m\text{-}i} \left(1\text{-}F^a_{S,i}(a_i)\right)^{n\text{-}j}. \tag{41}$$

For (ii) it is the probability that a_i is greater than the l'th action. That is, for some $k \in [l, m+n]$ the number of offers below a_i is exactly equal to k. Summing over k yields that

$$\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \leq a_i\right] = \sum_{k=l}^{n+m} \sum_{\substack{i+j=k \\ 0 \leq i \leq m \\ 0 \leq j \leq n}} \binom{m}{i}\binom{n}{j} F^a_{B,i}(a_i)^i \, F^a_{S,i}(a_i)^j \left(1\text{-}F^a_{B,i}(a_i)\right)^{m\text{-}i} \left(1\text{-}F^a_{S,i}(a_i)\right)^{n\text{-}j}. \tag{42}$$

For (iii), because distributions are assumed to be atomless $\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} = a_i\right] = 0$. Thus,

$$\mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \geq a_i\right] = 1 - \mathbb{P}_{\text{-}i}\left[a^{(l)}_{m,n} \leq a_i\right]. \tag{43}$$

B.4 Aggregate uncertainty

Consider an infinite market with regular transaction costs. Recall that regular transaction costs in infinite markets only depend on a traders action and the market price. Regular uninfluenceable transaction costs are functions of the market price, that is $\Phi_i(a_i, P^*) = \Phi_i(P^*)$. Examples include constant and price fees. Regular transaction costs are influenceable in infinite markets iff the map $a_i \mapsto \Phi_i(a_i, P^*)$ is strictly increasing for buyers and strictly decreasing for sellers. Spread fees are again an example of influenceable transaction costs. In this section suppose that trader i is uncertain about the uninfluenceable market price P^*. We assume P^* is a random variable and that the distribution is absolutely continuous with probability density function $f_{P^*,i}$ that is continuous and strictly positive on its support $[\underline{P_i^*}, \overline{P_i^*}]$ with $\underline{P_i^*} < \overline{P_i^*}$. Denote by $F_{P^*,i}$ the corresponding distribution function. Additionally, trader i also holds individual beliefs about the tie-breaking probability $p_i \in [0,1]$, if a traders' action is equal to P^*. Trader i may be more or less certain about their beliefs, which, for some degree $\delta > 0$, we measure by *δ-aggregate uncertainty* as follows: given $\delta > 0$, there exists a price P_i^*, such that $\mathbb{P}_i\left[P^* \in [P_i^* - \delta, P_i^* + \delta]\right] \geq 1 - \delta$.

Predictability of trade. For a buyer b with action a_b the probability of trading is equal to $1 - F_{P^*,b}(a_b)$. For a seller with ask a_s, it is equal to $F_{P^*,s}(a_s)$. δ-aggregate uncertainty is directly related to the predictability of trade. P_i^* corresponds to the critical value. If trader i submits an action that is strictly less (more) aggressive than the critical value, then for sufficiently small $\delta > 0$, the probability of trading is at least $1 - \delta$ (at most δ). Therefore Proposition 3 directly extends to settings with small uncertainty.

Existence of best responses. Proposition 5 extends to markets with aggregate uncertainty. The same proof method as in Appendix A.6 works. That is, the expected utility is continuous as a function of the action a_i of trader i. As best responses are necessarily located in the compact space $[\underline{P_i^*}, \overline{P_i^*}]$, the existence of a maximum follows from the Extreme Value theorem.

Uninfluenceable transaction costs. In the presence of aggregate uncertainty, Theorem 7 and Theorem 8 can be strengthened, as truthfulness is the unique best response.

Proposition 20. *Consider an uninfluenceable transaction cost and δ-uncertainty. For every $\delta > 0$, truthfulness is the unique best response.*

Proof. Consider a buyer b with gross value t_b and action a_b. The expected utility is $\mathbb{E}_b[u_b(t_b, a_b, P^*)] = \int_{\underline{P^*}}^{a_b} (t_b - x - \Phi_b(x))\, f_{P^*}(x)dx$, because tie-breaking is a probability zero event. Recall from Proposition 1 that $t_b - t_b^\Phi - \Phi_b(t_b^\Phi) = 0$. By assumption, the map $x \mapsto x + \Phi_b(x)$ is strictly increasing. Therefore, for $x \in [\underline{P^*}, t_b^\Phi)$, the integrand is strictly greater than zero. For $x \in (t_b^\Phi, \overline{P^*}]$, the integrand is strictly negative. Hence, the expected utility is maximized at the unique point $a_b = t_b^\Phi$. The function $a_b \mapsto \mathbb{E}_b[u_b(t_b, a_b, P^*)]$ is continuous, increasing on $[\underline{P^*}_b, t_b^\Phi]$ and decreasing on $[t_b^\Phi, \overline{P^*}_b]$. ϵ-best responses therefore approximate t_b^Φ. As truthfulness is the unique best response a_b, it holds that $E = \frac{\mathbb{P}_P^*[b \in \mathcal{B}^*(a_b, P^*)]}{\mathbb{P}_P^*[b \in \mathcal{B}^*(t_b^\Phi, P^*)]} = \frac{\mathbb{P}_P^*[b \in \mathcal{B}^*(t_b^\Phi, P^*)]}{\mathbb{P}_P^*[b \in \mathcal{B}^*(t_b^\Phi, P^*)]} = 1$. $\square$

Influenceable transaction costs. Theorem 9 and Theorem 10 also extend to markets with sufficiently small aggregate uncertainty.

Proposition 21. *Consider an influenceable transaction cost, δ-uncertainty, and assume that for trader i bidding the critical value P_i^∞ is strictly individually rational. Then, if δ is sufficiently small, best responses approximate price-guessing*

Proof. Consider a buyer b with gross value t_b and action a_b. Suppose that $t_b^\Phi > P_b^*$. Tie-breaking is a probability zero event. The expected utility is equal to $\mathbb{E}_b[u_b(t_b, a_b, P^*)] = \int_{\underline{P^*}}^{a_b} (t_b - x - \Phi_b(a_b, x))\, f_{P^*}(x)dx$. The expected utility is continuous in a_b on $[\underline{P^*}, \overline{P^*}]$ and attains a maximum by the Extreme Value Theorem, which proves the existence of a best response. First, we show that an action $a_b^1 < P_b^*$ is not a best response. We show that there exists an action $a_b^2 > P_b^*$ such that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, a_b^2, a_{\text{-}b})\right] - \mathbb{E}_b\left[u_b(t_b, a_b^1, a_{\text{-}b})\right] > 0$, which implies that a_b^1 is not a best response. Because the net value is by assumption continuous and strictly increasing in the gross value, there exists a gross value $t_b' < t_b$, such that $t_b^\Phi > t_b^{\Phi\prime} > P_b^\infty$. Denote the difference between t_b^Φ and $t_b^{\Phi\prime}$ by $\delta > 0$. It holds that $\mathbb{E}_{\text{-}b}\left[u_b(t_b, t_b^{\Phi\prime}, a_{\text{-}b})\right] = \mathbb{E}_{\text{-}b}\left[u_b(t_b', t_b^{\Phi\prime}, a_{\text{-}b})\right] + \delta \geq \delta$, because the net value is assumed to be ex-post individually rational. We therefore consider an action a_b with $a_b - P_b^* \geq \epsilon$ for some $\epsilon > 0$. We will show that if the aggregate uncertainty δ is sufficiently small, then a_b is not a best response, proving that best responses must be ϵ-close to P_b^*. More specifically, we prove that a buyer can increase their expected utility when switching to $P_b^* + \epsilon/2$.

For $\delta < \epsilon/2$ it holds that

$$\begin{gathered}
\mathbb{E}_b[u_b(t_b, a_b, P^*)] - \mathbb{E}_b[u_b(t_b, P_b^* + \epsilon/2, P^*)] = \\
\int_{\underline{P^*}}^{a_b} (t_b - x - \Phi_b(a_b, x))\, d\mu_{P^*}(x) - \int_{\underline{P^*}}^{P_b^*+\epsilon/2} (t_b - x - \Phi_b(P_b^\infty + \epsilon/2, x))\, d\mu_{P^*}(x) = \\
\int_{P_b^*+\epsilon/2}^{a_b} (t_b - x)\, d\mu_{P^*}(x) - \left(\int_{\underline{P^*}}^{P_b^*+\epsilon/2} (\Phi_b(a_b, x) - \Phi_b(\epsilon/2, x))\, d\mu_{P^*}(x) + \int_{P_b^*+\epsilon/2}^{a_b} \Phi_b(a_b, x) d\mu_{P^*}(x) \right).
\end{gathered} \tag{44}$$

Note that for any two actions $a_b^1 \geq a_b^2$ there exists a constant $\gamma > 0$, such that for all $P \in [\underline{P^*}, a_b^2]$ it holds that $\Phi_b(a_b^1, P) - \Phi_b(a_b^2, P) \geq \gamma$. That is because the map $a_b \mapsto \Phi_b(a_b, P)$ is strictly increasing on $[\underline{P^*}, a_b]$. Therefore, for fixed actions a_b^1 and a_b^2 the continuous function $P \mapsto \Phi_b(a_b^1, P) - \Phi_b(a_b^2, P)$ is strictly positive on the compact interval $[\underline{P^*}, a_b^2]$ and attains a strictly positive minimum by the Extreme Value theorem. Consider the constant $\gamma > 0$ that corresponds to $a_b^1 = a_b$ and $a_b^2 = P_b^* + \epsilon/2$. Together with δ-aggregate uncertainty, we get that $\int_{\underline{P^*}}^{P_b^*+\epsilon/2} (\Phi_b(a_b, x) - \Phi_b(P_b^* + \epsilon/2, x))\, d\mu_{P^*}(x) \geq (1 - \delta)\gamma$. Moreover it holds that $\int_{P_b^*+\epsilon/2}^{a_b} (t_b - x)\, d\mu_{P^*}(x) \leq \delta t_b$ and $\int_{P_b^*+\epsilon/2}^{a_b} \Phi_b(a_b, x) d\mu_{P^*}(x) \geq 0$. Thus $\mathbb{E}_b[u_b(t_b, a_b, P^*)] - \mathbb{E}_b[u_b(t_b, P_b^* + \epsilon/2, P^*)] \leq t_b\delta - (1 - \delta)\gamma$. If $\delta < \frac{\gamma}{t_b + \gamma}$, then the difference in expected utility is strictly negative, proving that a_b is not a best response. This implies that best responses are ϵ-close to P_b^* if δ is sufficiently small. □