Financial Transaction Taxes and the Informational Efficiency of Financial Markets: A Structural Estimation ⋆

Marco Cipriani[1,*], Antonio Guarino[2], Andreas Uthemann[3]

Abstract

We develop a new methodology to estimate the impact of a financial transaction tax (FTT) on financial market outcomes. In our sequential trading model, there are price-elastic noise and informed traders. We estimate the model through maximum likelihood for a sample of 60 NYSE stocks in 2017. We quantify the effect of introducing an FTT given the parameter estimates. An FTT increases the proportion of informed trading, improves information aggregation, but lowers trading volume and welfare. For some less liquid stocks, however, an FTT blocks private information aggregation.

Keywords: Financial Transaction Tax, Market Microstructure, Structural Estimation.
JEL: G14, D82, C13

⋆ Ron Kaniel was the editor for this article. We thank the editor, an anonymous referee, Douglas Gale, Piero Gottardi, Toru Kitagawa, Nicola Pavoni, Jean-Charles Rochet, Peter Norman Sørensen and seminar participants at several institutions and conferences for comments and suggestions. We thank Hugo Freeman, Aaron Plesset and Federico Tagliati for outstanding research assistance. The views expressed in this paper are solely those of the authors and do not necessarily represent those of the Bank of Canada, the Federal Reserve Bank of New York, or the Federal Reserve System. Andreas Uthemann thanks the Economic and Social Research Council (ESRC) for funding provided via the Systemic Risk Centre [grant numbers ES/R009724/1 and ES/K002309/1].
* Corresponding author: `marco.cipriani@ny.frb.org`
[1] Money and Payments Studies Function, Federal Reserve Bank of New York
[2] Department of Economics, University College London
[3] Bank of Canada and Systemic Risk Centre, London School of Economics

1. Introduction

In August 2012, France introduced a 0.2% financial transaction tax (FTT) on the purchase of shares of companies with a market capitalization above €1 billion. A similar tax was introduced in Italy in March 2013, and in Spain in 2021. A proposal by the European Commission for the introduction of a 0.1% *ad valorem* FTT for the entire European Union has triggered lengthy debates among member states; as of 2021, a decision by the European parliament has yet to be made. Proposals to introduce the tax have also been advanced in many other countries, including the United States, especially after the 2008 financial crisis.[4]

The debate on the merits of FTTs dates back to Keynes's *General Theory*, which proposed the use of these taxes to reduce stock market volatility and speculation. Tobin (1978) followed up on Keynes's tax proposal and suggested a 1% tax on all foreign exchange transactions (the so-called "Tobin tax") to reduce capital flow and exchange rate volatility. Later, Stiglitz (1989) and Summers and Summers (1989) advocated the use of FTTs to avoid the build-up of asset bubbles and reduce market volatility. The economic argument underlying these proposals is that financial markets suffer from too much "noise trading" activity not based on fundamentals with adverse consequences for price volatility and informational efficiency. An FTT would mainly affect such traders with short investment horizons, leaving long-term investors unaffected. As a result, the effect of the tax would be socially beneficial. The opposite view is that an FTT, by introducing a friction into the trading process, slows down

[4] Transaction taxes are not a recent policy innovation. The United Kingdom's Stamp Duty, implemented in 1694 to finance the war against France, is the first financial transaction tax ever imposed; although reformed several times, it is still in force. Sweden imposed an FTT from 1984 until 1990. The United States levied a transaction tax from 1914 to 1966. In the US, there is a tiny fee of $21.80 per million dollars of securities transactions supports the operation costs of the Securities and Exchange Commission ("Section 31 fee"). Moreover, exchanges and brokers sometimes require the payment of trading fees, which are similar to an FTT from an economic point of view; however, they are typically per share, rather than *ad valorem* and their proceeds are not collected by the government. Importantly, when they are imposed, these fees are levied at rates that are negligible compared to those of an FTT.

price discovery, thereby lowering liquidity and increasing price volatility. This view has been advocated by, for instance, Edwards (1993) and Schwert and Seguin (1993), who observe that informed traders stabilize markets by offsetting the effects of noise traders. Similarly, Kupiec (1996) shows that an FTT increases price volatility and lowers market liquidity.

A recent theoretical literature has studied the mechanisms through which an FTT impacts different classes of market participants and generally has found that the overall effect of the tax on market outcomes is theoretically indeterminate. Dávila and Parlatore (2021) show that, in a CARA-Normal set up, the impact of the tax on price informativeness depends on the relative elasticity of informed traders and noise traders (hedgers).[5] Sørensen (2017) studies the impact of an FTT on market composition and welfare in a one-period version of the Glosten and Milgrom (1985) model. In Sørensen (2017) noise traders trade because they value the asset differently from the market maker (and the informed traders); these differences in asset valuations arise from liquidity or hedging reasons. The impact of an FTT on welfare is theoretically ambiguous, since it can lower the cost of trading for noise traders (if it crowds out informed traders and reduces the bid-ask spread) or increase it (if the total cost of trading due to the spread and the tax is higher than without a tax). The net result depends on the relative price elasticity of noise and informed traders.[6]

[5] Dupont and Lee (2007) show that, in a static model of a competitive specialist market, an FTT decreases (increases) market depth and liquidity when informational asymmetry is high (low). Subrahmanyam (1998) shows that, in a strategic model, an FTT increases market liquidity with a monopolist-informed trader but decreases with multiple informed traders. Song and Zhang (2005) propose a model of noise trading in which the FTT has a compositional effect on trading activity: a tax discourages noise trading and fundamental trading to a different extent depending on market conditions; in an economy with low (high) volatility and low (high) noise traders' participation, an increase in the FTT reduces (increases) volatility. Other papers study the effect of the FTT on different types of traders using agent-based model, see, e.g., Mannaro et al. (2008) and Pellizzari and Westerhoff (2009).

[6] Glosten and Putniņš (2019) also study welfare in a Glosten and Milgrom (1985) model, but they do not study the effect of a tax. They illustrate how the welfare loss due to informed trading (and a positive spread) depends on how market characteristics, such as the quality of private information, affect market composition. The welfare implication of an FTT is also studied by Citanna et al. (2006), in a general equilibrium model with incomplete markets.

Unfortunately, the existing empirical studies, mainly event studies or difference-in-differences analyses, only partially resolve the theoretical ambiguity highlighted by the theoretical literature (see, e.g., the discussion in Habermeier and Kirilenko (2003)). There is strong evidence on the negative impact of the FTT on trading volumes. The best known case is the Swedish transaction tax of 1986, which led to a migration of 60% of trading volume in the eleven most traded Swedish shares from Stockholm to London (see Umlauf (1993)); Colliard and Hoffmann (2017) find a similar impact on volume for the imposition of an FTT in France. The results on other market outcomes, for instance, price volatility, however, are more ambiguous: while some empirical studies find that an FTT reduces price volatility (Umlauf (1993); Jones and Seguin (1997)), others find either a non significant or a negative effect (e.g., Colliard and Hoffmann (2017) and Deng et al. (2018)).[7]

In contrast to the existing empirical literature, in this paper, we estimate the impact of an FTT through structural estimation. In particular, we develop a market-microstructure model of trading in financial markets *à la* Glosten and Milgrom (1985), featuring both informed and noise traders that is amenable to structural estimation. Informed traders trade as they possess private information whereas noise traders trade because they have a private valuation of the asset. Both types of traders have price-elastic demands. As a result, the introduction of an FTT affects their behavior; its impact on informational efficiency, liquidity, volatility, and welfare depends on the model's parameter values.

We resolve this indeterminacy by estimating the model using transaction data on a sample of 60 stocks traded on the New York Stock Exchange (NYSE) in 2017, a period in which no FTT was levied. To study how the introduction of an FTT would impact different market segments, we divide the market into quartiles by market capitalization, and randomly

[7] The impact of an FTT on volatility and market efficiency has also been studied in the laboratory (e.g., Noussair et al. (1998), and Cipriani and Guarino (2008b)).

sample 15 stocks from each quartile. Among other parameters, we estimate the proportion of informed traders in the market, and the price sensitivity of both informed and noise traders. We simulate the model without an FTT and with different tax rates; this allows us to compare trading volume, bid-ask spread, price volatility, informational efficiency, and welfare under different tax regimes.

The structural estimation approach has several advantages compared to the standard diff-in-diff approach. It enables us to measure the effects of the introduction of an FTT through counterfactual policy experiments. We estimate the impact of the FTT on trading volume, volatility, bid ask-spread, and market composition (i.e., we address the debate on whether the tax affects noise traders' activity more than informed traders'); importantly, we are able to estimate the impact as a function of the rate. Moreover, we gauge the impact of a tax beyond observable market data, such as the price levels or intraday price volatility. In particular, we can recover market participants' beliefs and preferences from the structural parameter estimates; we use these estimates also to construct direct measures of informational efficiency and welfare, with and without an FTT, something that could not be achieved without a structural estimation.

In previous structural estimations of market microstructure models — e.g., Easley et al. (1996, 1997) and the following voluminous literature on the PIN (probability of informed trading) — noise traders trade for exogenous reasons (e.g., liquidity shocks) independently of the price. Moreover, informed traders receive a perfectly informative signal, so that they buy or sell independently of the price level. Since traders' behavior does not depend on the price, these models are not suitable to study an FTT. Recently, Cipriani and Guarino (2014) have studied herd behavior through a structural estimation of a market microstructure model in which informed traders (but not noise traders) are price elastic as they receive a signal of finite precision.[8] Here, we build on Cipriani and Guarino (2014) and introduce price-elastic noise traders, who receive a shock to their asset valuation, using the approach developed by Glosten and Putniņš (2019) to study welfare in a Glosten and Milgrom (1985) model; such a shock may be interpreted as the result of hedging motives. The presence of both price-elastic informed traders and price-elastic noise traders allows us to measure the compositional effect of the tax.

In our theoretical model, an FTT can completely shut down informed traders' activity or even overall trading activity; given our parameter estimates, such drastic effects rarely occur at the tax rates at which an FTT is usually levied. Across all stocks and within each quartile, however, the tax has a strong negative impact on trading volume, a result consistent with that found in the existing empirical literature.

Furthermore, the price sensitivity of informed traders (due to the precision of their signal) is generally lower than that of noise traders (due to their private values). Because of this, an FTT increases the fraction of trading activity by informed traders, widens the bid-ask spread, and increases price volatility. An important focus of our analysis is informational efficiency, measured as the distance of the price from the fundamental value during the day; for almost all stocks, the tax improves informational efficiency. Finally, by increasing the spread and lowering noise traders' market participation, an FTT reduces welfare.

For few stocks, mainly in the first market capitalization quartile, we estimate the precision of informed traders' signal to be very low. For these stocks the impact of the FTT is generally the opposite of the one described above—e.g., volatility decreases and so does the market informational efficiency—pointing to a significant source of heterogeneity in the tax impact. Moreover, in our theoretical model, for these stocks, an FTT may cause informed traders to

[8] For theoretical studies of herd behavior in financial markets, see Gervais (1997), Avery and Zemsky (1998), and Cipriani and Guarino (2008a).

stop trading altogether, similarly to an informational cascade (Bikhchandani et al. (1992); Welch (1992); Gale (1996); Hirshleifer and Hong Teoh (2003); Smith and Sørensen (2000)). As a result, the price may not converge to the fundamental asset value. For these stocks, even with a relatively small FTT of 5 bps, the probability of the price not converging to the fundamental value is significant.

The rest of the paper is organized as follows. Section 2 describes the model and its equilibrium predictions. Section 3 explains the effect of an FTT. Section 4 describes the data. Section 5 presents the parameter estimates. Section 6 reports the results on the impact of an FTT. Section 7 concludes. The Appendix contains further estimation results and other supplementary material.

2. The Model

Our model builds on Cipriani and Guarino (2014). Informed and noise traders trade an asset with a market maker over multiple days. In contrast to Cipriani and Guarino (2014), noise traders are price elastic. In the following subsections, we briefly describe the model and refer the reader to Cipriani and Guarino (2014) for a more complete description; wherever there is a significant difference between the two models, we describe it in detail.

2.1. The asset

The fundamental value of the asset on day $d = 1, 2, 3, ...$ is denoted by V^d. With probability α, an information event occurs and the asset value changes from the previous day, that is, $V^d \neq v^{d-1}$. On information event days, with probability δ, the asset value increases to $v_H^d = v^{d-1} + \lambda_H v^{d-1}$ where $\lambda_H > 0$ (good-event day); with probability $1 - \delta$, it decreases to $v_L^d = v^{d-1} + \lambda_L v^{d-1}$ where $-1 < \lambda_L < 0$ (bad-event day).[9] We assume that

[9] Note that, in contrast to Cipriani and Guarino (2014), the change in the asset value is modeled as a multiplicative change. In the existing literature, e.g., Easley et al. (1997), the typical assumption is that

$(1-\delta)\lambda_L = -\delta\lambda_H$, which implies that the value (and therefore, as we shall see, the daily closing price) is a martingale.

2.2. The market structure

Each day d, trading happens at discrete times $t = 1, 2, 3, \ldots$. At each time t of day d, a trader can exchange (buy or sell) one unit of the asset with the market maker or decide not to trade. We denote the action of a trader at time t on day d by x_t^d and the history of trades and prices until time $t-1$ of day d by h_t^d.

The market maker sets the ask and bid prices at which traders can buy or sell. We denote the ask price at time t of day d by a_t^d and the bid price by b_t^d. As in Glosten and Milgrom (1985), the market maker sets the ask and bid prices equal to the expected value of the asset conditional on the history of trades, h_t^d, and the action at time t. Note that the market maker does not know whether an information event has occurred on a given day.

2.3. The traders

There are two types of traders, informed and noise traders. A trader's type is private information. On information event days, at each time t an informed trader is chosen to trade with probability μ and a noise trader is chosen to trade with probability $1-\mu$. On no-event days, all traders are noise traders.

Informed traders. Informed traders are risk neutral. An informed trader active at time t on day d receives a private signal S_t^d about the asset value v^d, distributed according to the

the change in fundamental value between days is additive. Neither choice is relevant for the estimation of the model's parameters. However, a multiplicative change in fundamental value is preferable when studying an FTT since it makes the impact of an *ad-valorem* tax (as most FTTs are) proportional to the change in asset value. Additionally, a multiplicative change is consistent with the observation that the variance of asset returns does not decrease as the price of the asset increases (as would be the case with additive changes).

following value-contingent linear density functions

$$f^H(s_t^d|V^d = v_H^d) = 1 + \tau(2s_t^d - 1), \tag{1}$$
$$f^L(s_t^d|V^d = v_L^d) = 1 - \tau(2s_t^d - 1),$$

where the parameter $\tau \in (0, 2]$ measures the informativeness of the signal. For $0 < \tau \leq 1$, the support of the densities is $[0, 1]$; for $\tau > 1$, the support shrinks to $[1 - \frac{1}{\tau}, 1]$ for f^H and to $[0, \frac{1}{\tau}]$ for f^L. An informed trader knows that an information event has occurred; depending on the signal realization and the precision τ, they may still be unsure about the direction of the change. As $\tau \to 0$, signals become uninformative. As τ increases, so does the signal informativeness. Following Smith and Sørensen (2000), beliefs are bounded when $\tau \in (0, 1)$ and unbounded when $\tau \in [1, 2]$. With bounded beliefs, no signal realization perfectly reveals the asset value. In contrast, with unbounded beliefs some high (low) signal realizations are possible only when the asset value is high (low), and, therefore, the signal can be perfectly informative; a signal $s_t^d \geq 1/\tau$ perfectly reveals that the asset has a high value on day d, and a signal $s_t^d \leq (\tau - 1)/\tau$ perfectly reveals that the value is low. As τ tends to 2, almost all informed traders perfectly know the asset value.

Noise traders. Noise traders choose to buy or sell the asset not because they are informed about its fundamental value, but because their private valuation of the asset differs from that of the market maker, e.g., because of hedging reasons. In contrast to conventional empirical market microstructure models of sequential trading, we allow noise traders to be price-elastic. Whereas the conventional price-inelastic noise traders always have a private valuation of the asset so large or so small that they want to buy or sell independently of the price, price-elastic noise traders' private valuations are such that their decision to trade depends on the price they face.

We model noise traders by following an approach similar to that of Glosten and Putniņš

(2015). Specifically, with probability $0 < \varepsilon < 1$, noise traders receive a pseudo signal (or shock) n_t^d, which is distributed uniformly on the interval $[0,1]$, independently of the asset's fundamental value. With probability $(1-\varepsilon)$ noise traders receive no pseudo signal.

Since the pseudo signal is independent of the asset's fundamental value (that is, of whether we are in a good, bad or no-event day), it conveys no information. Noise traders, however, update their asset valuation *as if* the pseudo signal were informative. In particular, upon receiving it, a noise trader computes their asset valuation *as if* the pseudo signal were distributed according to the following value-contingent pseudo densities:

$$\begin{aligned}\tilde{g}^H(n_t^d|V^d = v_H^d) &= 1 + \nu(2n_t^d - 1), \\ \tilde{g}(n_t^d|V^d = v^{d-1}) &= \frac{\nu}{2-\nu}, \\ \tilde{g}^L(n_t^d|V^d = v_L^d) &= 1 - \nu(2n_t^d - 1),\end{aligned} \tag{2}$$

with the supports of the three pseudo densities being, respectively, $[\frac{\nu-1}{\nu}, 1]$, $[\frac{\nu-1}{\nu}, \frac{1}{\nu}]$, and $[0, \frac{1}{\nu}]$ and with $1 \leq \nu \leq 2$ (for $\nu = 2$, $\tilde{\Pr}(n = 0.5|V^d = v^{d-1}) = 1$).[10] As in the case of informed traders, the value-contingent pseudo densities are linear; their slopes depend on the parameter ν.

With probability $1-\varepsilon$, when noise traders do not receive a pseudo signal, their asset valuation is equal to that of the market maker; therefore, they optimally decide not to trade whenever the bid-ask spread is positive. With probability ε, when they receive a pseudo signal, they update their valuation for the asset to $E(V^d|h_t^d, n_t^d)$ analogously to how informed traders update valuations upon observing informative signals (event though, in the case of noise traders, the pseudo signal is not informative).

[10] We use the symbol $\tilde{g}$ to emphasize that these are not the true distributions of the random variable n_t^d; for the same reason, we refer to them as "pseudo densities."

As shown by Glosten and Putniņš (2015), one can interpret the pseudo signal n_t^d in two different ways. Pseudo signals may be a way of modeling bounded rationality: under this interpretation, noise traders believe that their pseudo signal is informative about the asset fundamental value, whereas it is not.

There is, however, another interpretation: the wedge between a noise trader's and the market maker's asset valuation created by the pseudo signal n_t^d may be due to hedging reasons of rational noise traders. Indeed, Glosten and Putniņš (2015) prove, in a two-state economy where traders are risk averse and have hedging reasons to trade, that there is a distribution of traders' risk aversion such that their asset valuation is distributed identically to the asset valuation generated by the pseudo signal. In Appendix D.2, we prove an equivalent result for our three-state economy. In Glosten and Putniņš (2015) and in our proof, the hedging reasons arise because risk-averse noise traders are either endowed with or short of one unit of the asset; a similar proof could be constructed assuming that they own an asset or an income stream that is negatively or positively correlated with the risky asset.[11]

The equivalence between private value due to hedging reasons and pseudo signal is intuitive: because of the mechanism of Bayesian updating, the wedge between the market maker's and the noise traders' valuation generated by the pseudo signal shrinks as uncertainty about the fundamental decreases during the day (because of trading activity); this is consistent with the wedge being generated by hedging reasons, which become smaller as the asset value is learned.[12]

[11] In Glosten and Putniņš (2015), Glosten and Putniņš (2019), and in our proof, the hedging reasons arise because noise traders are either endowed with or short of one unit of the asset; a similar proof could be constructed assuming that they own an asset or an income stream that is negatively or positively correlated with the risky asset. Note that in Glosten and Putniņš (2019) the same distribution of asset valuations is obtained when traders' risk aversion coefficient is the same but there is heterogeneity in wealth; we offer a similar proof for our three-state economy case in Appendix D.2.

[12] In Appendix D.1, we also show that the distribution of private valuations generated by the pseudo signal is independent of the fundamental value of the asset conditional on the history of trades; we also show that the wedge created by noise traders' pseudo signals with respect to the market maker's valua-

Some final remarks are in order. First, as we mentioned above, the pseudo signal is not informative since it is uniformly distributed on $[0, 1]$ independently of the type of day. Second, since noise traders' asset private valuation is pinned down by linear pseudo density functions, we can solve the model analytically, which is important for the maximum likelihood estimation. Third, the support of the pseudo density $\tilde{g}(\cdot|v^{d-1})$ on a no-event day is the intersection of the supports of $\tilde{g}^H(\cdot|v_H^d)$ and $\tilde{g}^L(\cdot|v_L^d)$ for good and bad-event days. This, along with the restriction $\nu \geq 1$, guarantees that, for any history, there are always noise traders whose asset valuation is either equal to v_H^d or to v_L^d; these traders want to buy and sell irrespective of any price in the range $[v_L^d, v_H^d]$, thus avoiding market breakdowns.[13] Finally, at least for some signal realizations, if $\nu < 2$, noise traders' valuation of the asset is always between v_L^d and v_H^d, so that at any time t noise traders' demand is, indeed, price-elastic; for $\nu = 2$, noise traders act as if, with probability one, the day were a bad-event day (when $n_t^d < 0.5$) or a good-event day (when $n_t^d > 0.5$).

Modelling noise traders' private value shocks as pseudo signals is a simple and flexible way of introducing price-elastic noise traders into our asset market. This encompasses the standard empirical market microstructure models, where noise traders are price inelastic: whereas in those models price inelasticity is assumed, in our analysis the elasticity of noise traders' demand is estimated.[14] Allowing noise traders to be price-elastic creates an important trade off when analyzing the effect of a transaction tax: the tax discourages both

tion can be written as a non-i.i.d., time-varying multiplicative shock.

[13] Because of asymmetric information, the market would break down in the absence of noise traders. Note that from an empirical viewpoint, this is a natural assumption, since we do not observe such market breakdowns.

[14] As mentioned above, for $\nu = 2$, noise traders buy or sell with fixed probabilities as in the previous market microstructure literature. It is important to notice though, that in this literature, noise traders do not change their behavior for *any* price. In our model, instead, even for $\nu = 2$, noise traders remain price-elastic, since, e.g., they do not buy when the price exceeds v_H, which seems a desirable characteristic. This may happen when the ask price is close to v_H and, in addition, the trader has to pay an FTT. We will discuss this issue in more detail when we introduce an FTT into our model.

informed and uninformed traders from trading and affects the market composition (that is, the relative proportion of informed and uninformed traders) as well as informational efficiency, welfare, and other market outcomes.

Comparing the estimates of τ and ν informs us about the relative sensitivity of informed and noise traders to the price. In particular, when $\tau > \nu$, the same realization of an informed trader's signal and of a noise trader's pseudo signal moves the informed trader's expectation more than the noise trader's asset valuation; as a result, there is a level of the FTT or of the bid-ask spread such that whereas the informed trader trades, the noise trader does not; in the remainder of the paper, we will refer to this parameter configuration as as situation in which informed traders are "less price elastic" as their behavior is less likely to be impacted by the imposition of the tax or widening of the bid-ask spread.

2.4. Equilibrium

Similarly to Cipriani and Guarino (2014), we can characterize the Perfect Bayesian Equilibrium of the economy through a unique set of thresholds that, at each time t, pin down informed and noise traders' decisions as a function of their (pseudo) signals. We denote informed traders' thresholds at time t of day d as σ_t^d and β_t^d, and noise traders' thresholds as κ_t^d and γ_t^d. An informed trader sells for any signal lower than σ_t^d, buys for any signal greater than β_t^d, and does not trade for any signal in between. Similarly, a noise trader sells for any pseudo signal lower than κ_t^d, buys for any pseudo signal greater than γ_t^d, and does not trade for any pseudo signal in between.

The thresholds are defined as the solution to a system of equations setting the traders' asset valuations given their (pseudo) signals equal to the bid or the ask price. For instance, the equation for an informed trader's buy threshold β_t^d is

$$E\left(V^d | h_t^d, \beta_t^d\right) = a_t^d. \tag{3}$$

As in Cipriani and Guarino (2014), the equilibrium thresholds at time t can be written as an explicit function of beliefs at time $t-1$ and the model's parameter.[15] Since the thresholds pin down the traders' strategies, they also determine the probability of a trade. For instance, given β_t^d and κ_t^d, the probability of a buy order on a good-event day at time t is

$$\Pr\left(x_t^d = \text{buy}|h_t^d, v_H^d, \Phi\right) = \qquad (4)$$
$$\mu\left[1 - F^H\left(\beta_t^d|v_H^d\right)\right] + (1-\mu)\varepsilon\left(1-\kappa_t^d\right),$$

where $F^H(\cdot|v_H^d)$ is the cumulative distribution function of $f^H(\cdot|v_H^d)$. To understand this expression, recall that a trader active at time t is an informed trader with probability μ and a noise trader with probability $1-\mu$. An informed trader buys if his signal is above the buy threshold β_t^d, which happens with probability $1 - F^H(\beta_t^d|v_H^d)$. A noise trader receives a pseudo signal shock with probability ε, in which case he buys if their pseudo signal is larger than κ_t^d, which happens with probability $1-\kappa_t^d$ (as pseudo signals are uniformly distributed). We use a similar methodology to compute the probability of any trade after any history recursively; given the probability of an action, we can compute agents' posterior beliefs via Bayes's rule and, therefore, the action threshold for the following periods. In this way, we obtain an analytic recursive expression for the model's likelihood function.[16]

3. The Financial Transaction Tax

We now consider the introduction of a financial transaction tax (FTT) levied *ad valorem*, as are most of the proposed and implemented FTTs. Specifically, whenever a trader buys

[15] In particular, they are a function of α, δ, μ, τ, ν, and ε, but not of λ_H and λ_L. Intuitively, the reason is that λ_H and λ_L have the same impact on traders' and market maker's valuations and cancel out from the thresholds' equations.

[16] The likelihood function is described in detail in Appendix B.

the asset, they pay a tax $\rho\, a_t^d$; and whenever they sell, they pay $\rho\, b_t^d$. As a result, an informed trader with signal s_t^d finds it optimal to buy when $E\left(V^d|h_t^d, s_t^d\right) > a_t^d(1+\rho)$, and to sell when $E\left(V^d|h_t^d, s_t^d\right) < b_t^d(1-\rho)$; they choose not to trade when $b_t^d(1-\rho) \leq E\left(V^d|h_t^d, s_t^d\right) \leq a_t^d(1+\rho)$. Similarly, a noise trader with pseudo signal n_t^d buys if $E\left(V^d|h_t^d, n_t^d\right) > a_t^d(1+\rho)$, sells if $E\left(V^d|h_t^d, n_t^d\right) < b_t^d(1-\rho)$ and chooses not to trade if $b_t^d(1-\rho) \leq E\left(V^d|h_t^d, n_t^d\right) \leq a_t^d(1+\rho)$.

As in the case of no FTT, the equilibrium can be characterized in terms of buy and sell thresholds for informed and noise traders. As we discussed in Section 2.4, the thresholds are the solutions to a system of equations that set the traders' valuations given their (pseudo) signal equal to the bid or ask prices. For instance, the equations for informed traders' buy and sell thresholds are

$$\begin{aligned} E\left(V^d|h_t^d, \beta_t^d\right) &= a_t^d(1+\rho), \\ E\left(V^d|h_t^d, \sigma_t^d\right) &= b_t^d(1-\rho). \end{aligned} \tag{5}$$

We refer the reader to Appendix B for the analytical derivation of the thresholds for informed and noise traders. Here we only note that, in equilibrium, informed traders who receive less informative signals and noise traders who receive weaker pseudo signals may find it optimal to abstain from trade to avoid paying the tax; therefore, the tax modifies the equilibrium bid and ask prices, since the proportion of informed and noise traders on either side of the market is different than if no tax were imposed.

If the tax is high enough, traders may be unwilling to participate in the market altogether. With a tax rate such that noise traders do participate, informed traders may still be unwilling to trade. The next proposition states these results formally.

Proposition 1. *There exists a tax rate*

$$\bar{\rho}^N = \left(\frac{2\delta}{1-|2\delta-1|}\right)\lambda_H \tag{6}$$

such that if $\rho > \bar{\rho}^N$*, in equilibrium, traders do not trade at time* $t = 1$ *(for any (pseudo) signal realization), that is, the market does not open. Suppose* $\rho < \bar{\rho}^N$*, then there exists a tax rate*

$$\bar{\rho}^I = \left(\frac{2\min\{\tau,1\}\,\delta}{1-\min\{\tau,1\}|2\delta-1|}\right)\lambda_H \leq \bar{\rho}^N, \tag{7}$$

such that if $\rho > \bar{\rho}^I$*, in equilibrium, informed traders do not trade at any time* t *(for any signal realization).*

We refer the reader to Appendix C for the proof of this and all following propositions. Here we note that, in the proofs, the thresholds $\bar{\rho}^N$ and $\bar{\rho}^I$ are computed considering the first trading time of a day. If traders do not find it profitable to trade at that time, they never find it optimal, since no information is revealed by the trading history and the prices set by the market maker remain constant. Note that if $\tau \geq 1$, $\bar{\rho}^I = \bar{\rho}^N$, since after an extreme signal (or pseudo signal) both informed traders and noise traders learn the value of the asset; in contrast, for $\tau < 1$, $\bar{\rho}^I$ is increasing in τ, that is, in how strongly the signal affects informed traders' evaluations.[17] Note also that the levels of the tax rates $\bar{\rho}^I$ and $\bar{\rho}^N$ increase with λ_H and δ: higher λ_H and δ mean that there are higher gains from trade and a higher tax rate is needed to discourage trading.[18]

Finally we can compute a threshold for the tax rate, such that both types of traders are

[17] For $\tau \geq 1$, there is a measure of informed traders receiving a perfectly revealing signal (or at least, a signal overwhelming any history of trades, when $\tau = 1$). Therefore, for a tax to shut informed traders out of the market, the tax payment must be higher than the gains from trade for a trader who knows the realization of the asset value. The same argument applies to noise traders, since $\nu \geq 1$.

[18] Remember that, because the fundamental value is a martingale, $|\lambda_L| = (\delta/(1-\delta))\,\lambda_H$; therefore, for a given λ_H, the higher is δ, the higher is $|\lambda_L|$, and therefore the gains from trade.

active on both sides of the market at time $t = 1$:

Corollary 1. *There exists a tax rate*

$$\bar{\rho} = \left(\frac{2 \min\{\tau, 1\}\, \delta}{1 + \min\{\tau, 1\} |2\delta - 1|} \right) \lambda_H < \bar{\rho}^I, \tag{8}$$

such that if $\rho < \bar{\rho}$, in equilibrium, at time $t = 1$, both noise traders and informed traders buy and sell for some (pseudo) signal realizations.

Note that $\bar{\rho}$ is smaller than $\bar{\rho}^I$, because when $\rho < \bar{\rho}$, informed traders both buy and sell (depending on their signal realization). When studying the asymptotic impact of an FTT we will assume that $\rho < \bar{\rho}$, that is, that the tax rate is low enough that, at least at time $t = 1$, informed traders both buy and sell.

3.1. The Asymptotic Impact of an FTT

In the absence of a tax, bid and ask prices and agents' beliefs converge almost surely to the true asset value (Cipriani and Guarino (2014)). With an FTT this is no longer true.

On event days, when informed traders' beliefs are bounded ($\tau < 1$), the FTT prevents the market maker from learning, even asymptotically, whether an information event is good or bad. The reason is that as the price converges toward either v_H^d or v_L^d, private information becomes less valuable to informed traders and there will be a time after which they all find it optimal not to trade.[19] We state this in the next proposition:

Proposition 2. *Consider a tax rate $\rho < \bar{\rho}$. If $\tau < 1$, in equilibrium, on an event day, for almost all histories there exists a time T, such that, for any $t > T$, informed traders decide*

[19] The idea that a tax eventually stops all trading activity is developed by Cipriani and Guarino (2008b) in a Glosten and Milgrom (1985) model and, in a different setup, by Lee (1998).

not to trade for any signal realization s_t^d*, that is, as* $t \longrightarrow \infty$,

$$\Pr(X_t^d = \text{no trade}|h_t^d, s_t^d) = 1. \tag{9}$$

Since, as Proposition 2 shows, after a long enough history of trades, informed traders stop acting upon their signal, the market maker's beliefs about whether the event was good never converge to either 0 or 1.[20] Moreover, because of noise traders' activity and because private information is only of finite precision, the market maker's belief about the event being good may converge to a level below δ on a good-event day; analogously, it may converge to a level above δ on a bad-event day. That is, not only learning about the event type is never complete, the market maker may end up assigning a probability higher than 50% to the event that did not occur. In the following proposition we provide bounds on the probabilities that the market maker's belief about a good event remains stuck above 0.5 on a bad-event day, or stuck below 0.5 on a good-event day.

Proposition 3. *Consider a tax rate* $\rho < \bar{\rho}$*. If* $\tau < 1$*, in equilibrium, there exists* $0 < \delta_l < \delta_h < 1$ *such that the probability that the belief* $\Pr(V^d = v_H^d|h_t^d)$ *converges to a value lower than* δ_l *on a good-event day* ($V^d = v_H^d$) *is bounded above by* $\frac{\delta_l(\delta_h-\delta)}{\delta(\delta_h-\delta_l)}$ *and the probability that the belief* $\Pr(V^d = v_H^d|h_t^d)$ *converges to a value higher than* δ_h *on a bad-event day* ($V^d = v_L^d$) *is bounded above by* $\frac{(1-\delta_h)(\delta-\delta_l)}{(1-\delta)(\delta_h-\delta_l)}$.

The logic of the proof is that, e.g., conditional on a bad event, the likelihood ratio $\frac{\Pr(V^d=v_H^d|h_t^d)}{\Pr(V^d=v_L^d|h_t^d)}$ is a martingale and, therefore, equal in expectation to its unconditional value

[20] Note that even when all informed traders stop trading because of the tax, a measure of noise traders continues to trade, and therefore, the probability of a given trade depends on whether an event has occurred; this is why, even with a tax, the market maker learns whether an event has occurred (see Lemma 1 in Appendix C). In other words, an informational cascade, a situation in which the probability of any action is independent of the fundamental value and, therefore, "[...] no new information reaches the market" (Avery and Zemsky, 1998, p. 733) never arises.

$(\frac{\delta}{1-\delta})$. This fact, along with the observation that informed trading ceases after $\Pr(V^d = v_H^d | V^d \neq v^{d-1}, h_t^d)$ has reached either the high or the low threshold allows us to pin down the probability of correct and wrong convergence of the market maker's belief, $\Pr(V^d = v_H^d | h_t^d)$ (see Appendix C for the derivation and a discussion).

4. The Data

Our sample consists of 60 stocks traded on the New York Stock Exchange (NYSE) in 2017. The dataset is constructed using the NYSE Trade and Quote (TAQ) and the Center for Research in Security Prices (CRSP) databases. We start by considering all common stocks that are listed in the CRSP database on December 30, 2016 and meet all standard criteria as detailed in Hasbrouck (2009) and Holden and Jacobsen (2014). We restrict attention to active stocks that have the NYSE as primary listing and whose closing price on December 30, 2016 exceeded $1. To study how the tax impacts different market segments, we divide this sample into quartiles by market capitalization, and randomly sample 15 stocks from each quartile (see Appendix A.1 for a list of the stocks and summary statistics). Our sample represents roughly 5% of the total market capitalization of the NYSE. The data comprises all trades and quotes during regular trading hours, that is, from 9:30am to 4:00pm, in the period from January 3, 2017 to December 29, 2017. To classify a transaction as a buy or a sell, we start by signing it with the standard Lee and Ready (1991) algorithm as implemented by an updated version of Holden and Jacobsen (2014).[21] We aggregate trades that occur within 500 microseconds (0.0005 seconds) of each other using the modal trade of that set.[22] This aggregation procedure is meant to capture the fact that it takes time for trading information

[21] Holden and Jacobsen (2014) consider datasets with time stamps at millisecond precision. We use a more recent TAQ version with microsecond precision. The authors have updated their methodology to sign trades for this case (for link to the code, see https://kelley.iu.edu/cholden/instructions.pdf).

[22] If the mode is not unique, we sign the trade as either a buy or as a sell with equal probability.

to reach market participants (see, e.g., Aquilina et al. (2021) for a discussion of reaction times in equity markets).[23]

We use the established convention of inserting no-trades between two transactions if the elapsed time between them exceeds a particular time interval (see, e.g., Easley et al. (1997)). Following Chung et al. (2005) and Cipriani and Guarino (2014) we chose as our interval the ratio between the total trading time in a day and the average number of transactions across trading days; in order to have the same interval across stocks (which allows us to compare our results) and to avoid missing relevant no trade periods in heavily traded stocks, we use the average number of transactions in the most traded stock (and therefore the smallest interval).[24] The interval resulting from this computation is 1 second: if there is no transaction for more than 1 second, we insert a number of no-trades equal to the number of seconds without trading activity. In the Appendix A.3, we show for the median stocks of each quartile that changing the no-trade interval around 1 second does not materially affect the results; although some parameters change (notably ε gets smaller as the no-trade interval increases), economically relevant composite parameters are largely unaffected; in particular, the proportion of informed trading (PIN) and the probability of observing one trade in a 1 second interval during a no-event day are stable (we present these measures in the next section; see Easley et al. (1997) for a discussion).

		α	δ	μ	τ	ν	ε
Q1	median	0.814	0.597	0.002	1.895	1.382	0.028
	mean	0.743	0.598	0.020	1.443	1.368	0.045
	std dev	0.183	0.137	0.044	0.830	0.105	0.050
	std err	0.032	0.048	0.000	0.006	0.013	0.001
Q2	median	0.822	0.603	0.002	1.919	1.353	0.110
	mean	0.788	0.565	0.012	1.797	1.404	0.158
	std dev	0.196	0.189	0.038	0.490	0.160	0.123
	std err	0.026	0.045	0.000	0.011	0.007	0.003
Q3	median	0.923	0.487	0.002	1.871	1.359	0.182
	mean	0.863	0.485	0.002	1.847	1.370	0.185
	std dev	0.120	0.145	0.0003	0.131	0.027	0.096
	std err	0.016	0.064	0.000	0.020	0.011	0.002
Q4	median	0.912	0.500	0.002	1.823	1.390	0.322
	mean	0.886	0.468	0.002	1.831	1.414	0.368
	std dev	0.082	0.240	0.0004	0.074	0.054	0.158
	std err	0.018	0.035	0.000	0.018	0.007	0.003
$\tau < 1$ stocks	median	0.585	0.672	0.076	0.097	1.305	0.040
	mean	0.558	0.570	0.087	0.100	1.432	0.062
	std dev	0.295	0.206	0.067	0.070	0.338	0.051
	std err	0.060	0.083	0.008	0.012	0.060	0.001
all stocks	median	0.867	0.584	0.002	1.894	1.378	0.143
	mean	0.820	0.529	0.009	1.730	1.389	0.189
	std dev	0.159	0.186	0.029	0.504	0.100	0.161
	std err	0.020	0.048	0.000	0.012	0.008	0.002

Table 1: **Parameter estimates** Summary statistics of the parameter estimates across all stocks, for each market capitalization quartile, and for bounded beliefs ($\tau < 1$) stocks. "std err" refers to the median standard error. Standard errors are estimated by bootstrap.

5. Results

5.1. Estimates

In this section, we show aggregate statistics for the parameter estimates of the stocks in our sample. The parameters are obtained by estimating the model presented in Section 2 (i.e., with no FTT) for each stock in our dataset through maximum likelihood.[25] Table 1 presents the median, mean, and standard deviation of the parameter estimates, across all stocks and separately for each quartile; we also present the median of the parameters' standard errors, estimated by bootstrap. In Table 3, we report the results of a series of Wald and Likelihood Ratio tests on the parameter estimates. Parameter estimates and standard errors at the individual stock level are reported in Appendix A.2.

The probability of an information event (α) is high for all quartiles. Across all stocks, the mean probability is 82% and increases from 74% of the first quartile to 89% of the fourth. In other words, more information events are observed in the stocks with the highest market capitalization. This result is similar to that of Easley et al. (1996), who find a positive correlation between their estimate of the probability of an information event and trading volume (market capitalization and trading volume are positively correlated, with a correlation of 0.32 in our sample of stocks). However, there is heterogeneity within quartiles, with the standard deviation ranging from 8% in the fourth quartile to 20% in the second

[23] Easley et al. (2012, 2016) also aggregate trades by "time bars" in their study of high frequency trading.

[24] If we had used the average number of transactions across all stocks, we might have missed relevant no-trades in the most traded stocks (which could have potentially affected the estimates of all parameters); in contrast, by using the average number of trades in the most traded stock, we are just adding additional no-trades in less-traded stocks (which only leads to a lower estimate of ε and does not affect economically relevant parameters).

[25] In the estimation, we use the Nelder-Mead algorithm as implemented in Julia's Optim package (Mogensen and Riseth, 2018). Before estimating the model with actual transaction data, we simulated data for several sets of parameters and verified that we were able to recover them with our estimation algorithm; the optimization routine converges to the same parameter set starting from a large set of initial conditions.

quartile. As Table 3 shows, for all stocks we reject the null that α equals 1 or 0; that is, for all stocks, there is event uncertainty.

The mean and median probability of a good event (δ) are slightly above 50%. While there is heterogeneity across quartiles, we do not observe any obvious relationship between δ and market capitalization.

On event days, the proportion of informed traders (μ) is, on average, 0.9%. This percentage is higher for the first quartile (2%) and lower for the last two quartiles (less than 1%). The median is stable across quartiles and always lower than 1%. There is heterogeneity in μ across stocks, with an overall standard deviation of 2.9%; the first quartile has the highest standard deviation.

The probability that a noise trader receives a private value shock (ε) monotonically and strongly increases with quartiles. The mean estimates increase from 5% in the first quartile to 16% in the second, 19% in the third and to 37% in the fourth; the median follows a similar pattern. The increase in ε (but not in μ) with market capitalization implies that the increase in trading activity for larger stocks is due to non-information based trading.

The mean informativeness of the private signal (τ) is 1.4 for the first quartile and approximately 1.8 for the others. The median across all stocks is 1.9. As Table 3 shows, for all stocks we reject the null that τ equals 2, that is, in contrast to what Easley et al. (1997) and several empirical market microstructure paper impose as an assumption, during event days, informed traders do not receive a perfectly informative signal. For most of our stocks, beliefs are unbounded since $\tau > 1$. Only for 5 stocks is τ significantly less than one (see Table 3), which implies that a large proportion of informed traders receives an incorrect signal;[26] all these stocks but one are in the first quartile. The aggregate statistics of the parameter

[26] Given the signal density functions, for $\tau < 1$ the probability of an incorrect signal, i.e., less than 0.5 on a good day and greater than 0.5 on a bad day, is given by $0.5 - 0.25\,\tau$. For $\tau = 0.1$, the average τ across the 5 stocks, this probability is 0.475.

estimates for this group of 5 stocks are included in Table 1.

The parameter ν (which measures the impact of noise traders' shocks on their valuations) is fairly constant across quartiles, with mean and median values ranging from 1.35 to 1.41. These estimates imply that, on average, almost half of noise traders who receive a shock are price elastic.[27]

To understand the implication of these parameters for trade informativeness, we compute the standard measure of PIN, the Probability of Informed Trade (Easley et al. (1996)):

$$\text{PIN} = \frac{\alpha\mu}{\alpha\mu + \varepsilon(1 - \alpha\mu)}, \tag{10}$$

where the numerator represents the probability of a trade by an informed trader and the denominator the approximate probability of any trade at the beginning of a day.[28]

Table 2 shows PINs by quartile and for stocks with $\tau < 1$. Across all stocks, the median PIN is 1.5%. The PIN is monotonically decreasing with the quartiles, from approximately 9% for the first quartile to 1.8% for the second, 1.3% for the third and less than 1% for the fourth. In other words, the order flow is more informative in stocks with lower market capitalization. This is in line with previous results on PIN (e.g., Easley et al. (1996)).

---

[27] Noise traders are price elastic when they receive signals belonging to the intersection of the supports of g^H and g^L. The proportion of price-elastic noise traders is given by the area corresponding to such intersection, measured with the uniform distribution, that is, with $\frac{2-\nu}{\nu}$ (for, e.g., $\nu = 1.37$, this proportion is 0.46).

[28] In computing these statistics, we implicitly assume that all informed traders and noise traders trade (i.e., their asset valuation never falls within the bid-ask spread). We do so to use the same PIN formulas as in Easley et al. (1996) and in much of the subsequent literature. In Appendix A.7, we also report results on the probability of an informed trade computed taking into account the bid-ask spread, and we show how this probability is affected by different rates of the FTT; moreover, we compute the probability of informed trade with correct private information (PCIN), as introduced by Cipriani and Guarino (2014).

	Q1	Q2	Q3	Q4	$\tau < 1$ stocks	all stocks
PIN	0.090	0.018	0.013	0.006	0.456	0.015

Table 2: **PIN** Median PIN measure across all stocks, for each quartile, and for those stocks with $\tau < 1$.

5.2. A comparison with previous models

Our model nests two previous models of sequential trading with private information: when $\nu = 2$, we obtain the model estimated in Cipriani and Guarino (2014), with inelastic noise traders but elastic informed traders; when $\tau = \nu = 2$ we obtain the model estimated in Easley et al. (1997), where both noise and informed traders are inelastic. It is interesting to see whether these restrictions are rejected by the data. Table 3 reports the Likelihood Ratio tests for these two restrictions.

For all stocks, we can reject the restriction that $\nu = 2$; that is, the assumption of inelastic noise traders as in Cipriani and Guarino (2014) and most of the market microstructure literature is rejected by the data. Similarly, the assumption that both noise and informed traders are inelastic as in Easley et al. (1997) ($\tau = \nu = 2$) is rejected for all 60 stocks.[29] Allowing for price elasticity in informed and noise traders' behavior significantly improves the fit of the model.[30]

These results imply that the behavior of both informed and noise traders changes through-

[29] The tests show that our model fits the data better than both Easley et al. (1997) and Cipriani and Guarino (2014); this is reassuring on the goodness of fit of our model. Additionally, in Appendix A.4, we compare the trading activity (average number of buys and sells) in the simulated model with those in the actual NYSE data. For each market capitalization quartile, and across all stocks, the simulated average number of buys and sells is very close to that of our NYSE stocks.

[30] In Appendix D.3, we report the standard deviations and the range of the difference between noise traders private asset valuations and the price (obtained through a simulation exercise, described in detail in Section 6.4. Across all stocks, the average daily standard deviation is 62 bps. See Hendershott and Menkveld (2014) and Hollifield et al. (2006) for other empirical approaches to estimate the standard deviation of noise traders' private valuations; our estimates are comparable to those found by Hendershott and Menkveld (2014) for NYSE stocks over 1995-2005; also similarly to Hendershott and Menkveld (2014), our estimates of the standard deviation decrease by market capitalization.

out the day and depends on the price these traders face, including an FTT, if one is imposed. Since the FTT impacts both informed and noise traders, the implication for informational efficiency and welfare are theoretically ambiguous and depend on the model's parameter estimates. In Section 6, we will be able to quantify them given our parameter estimates.

H_0	$\tau = 2$	$\tau = 1$	$\nu = 2$	$\alpha = 0$	$\alpha = 1$	CG	EKO
H_1	$<$	$<$	$<$	$>$	$<$	$\neq$	$\neq$
Rejections of H_0:							
1st quartile	15	4	15	15	15	15	15
2nd quartile	15	1	15	15	15	15	15
3rd quartile	15	0	15	15	15	15	15
4th quartile	15	0	15	15	15	15	15
all stocks	60	5	60	60	60	60	60

Table 3: **Test of hypotheses.** Results of various hypothesis tests. The first five columns are Wald tests, the last two columns are LR tests. CG refers to Cipriani and Guarino (2014) and EKO refers to Easley et al. (1997). H_0 specifies the null hypothesis, H_1 the alternative. The number of rejections (i.e. number of stocks for which H_0 is rejected) is based on a confidence level of 1%.

6. The Impact of an FTT

Given the parameter estimates of each stock, we study the impact of an FTT in three ways: i) we compute $\bar{\rho}^N$, the tax rate above which the market does not open, and $\bar{\rho}^I$, the tax rate above which all trading is noise; ii) for different levels of the tax, we compute the probability of severe mispricing, that is, that the market maker's belief about the occurrence of a good event remains stuck at a high (low) level on a bad (good)-event day; and iii) for different levels of the tax, we simulate the price path and the trading flow and compute measures of volume, liquidity, volatility, informational inefficiency, and welfare.

To conduct this analysis, we need an estimate of λ_H, the percentage change in the asset value on a good-event day. Indeed, when a tax is present, λ_H is needed to calculate the equilibrium trading thresholds for our informed and price-elastic noise traders: the size of the tax in relation to the possible gains from holding the asset, parameterized by λ_H (see footnote 18), affects these thresholds.

6.1. Calibration of λ_H

We cannot estimate λ_H as part of the maximum likelihood estimation, since in the absence of the tax trading decisions are independent of λ_H (see Section 2.4 and Appendix B). Instead, we estimate λ_H through the standard deviation of daily price changes. In a market with no FTT, the asset price converges to the asset fundamental value; if all price movements were due to private information, the standard deviation of the daily price percentage changes, σ, would be equal to $\lambda_H \sqrt{\frac{\alpha\delta}{1-\delta}}$ (see Appendix A.5).

Hence, for each stock, we could estimate σ from daily price changes and then compute λ_H using the estimates of α and δ. However, unlike in our theoretical model, observed price changes (and therefore their standard deviation) may also be due to events that became public information during or after the trading day (and which are not traders' private information and are not revealed by the order flow). Therefore, we need to decompose the price variance into public and private component and consider the private component only. If we did not do so, our estimates of λ_H would be inflated and as a result, the impact of a tax would be underestimated.

To this purpose, we use the variance decomposition method developed by Hasbrouck (1991): we decompose the variance of the stock price percentage changes into a trade-correlated component, interpreted as the component driven by private information, and a trade-uncorrelated component, interpreted as the component driven by public information. In the interest of space, we discuss this procedure in Appendix A.5 and here we only discuss the results.

Table 4 reports the square root of Hasbrouck's R_w^2 statistics along with the standard deviation of daily log-price changes (σ) and the estimated λ_H.[31] For each of these measures,

[31] We prefer to report the square root of Hasbrouck's statistics because we are decomposing the standard deviation of percentage changes (whereas Hasbrouck decomposes their variance).

		R_w	σ	λ_H	$\bar{\rho}^N$	$\bar{\rho}^I$	$\bar{\rho}$
Q1	median	0.409	0.028	0.011	0.0198	0.0115	0.0087
	mean	0.400	0.043	0.015	0.0231	0.0139	0.0094
Q2	median	0.371	0.023	0.007	0.0124	0.0122	0.0062
	mean	0.368	0.025	0.010	0.0160	0.0141	0.0070
Q3	median	0.348	0.013	0.005	0.0056	0.0056	0.0032
	mean	0.362	0.016	0.007	0.0077	0.0077	0.0046
Q4	median	0.427	0.011	0.005	0.0068	0.0068	0.0030
	mean	0.430	0.016	0.010	0.0116	0.0116	0.0041
$\tau < 1$	median	0.409	0.032	0.023	0.0258	0.0012	0.0011
	mean	0.413	0.048	0.022	0.0361	0.0029	0.0026
all stocks	median	0.387	0.018	0.006	0.0100	0.0087	0.0047
	mean	0.391	0.027	0.011	0.0146	0.0118	0.0063

Table 4: **Tax thresholds** Results of Hasbrouck decomposition, R_w and σ, standard deviation of log-price changes (daily), the resulting λ_H, and the corresponding tax thresholds across all stocks, for each market capitalization quartile, and for those stocks with $\tau < 1$.

we report the mean and median across all stocks, by quartile and, separately, for the stocks with $\tau < 1$. We find that the median value of R_w, that is, the fraction of the standard deviation of the log-price changes that is due to private information, is 39%. As a result, if the standard deviation of daily log-price changes is 1.8% (the median standard deviation in our sample), the standard deviation of daily log-price changes due to private information is $1.8\% \times 0.39 = 0.7\%$.

Given these results and our parameter estimates, across all stocks the median estimate of λ_H is 0.6% (and that of λ_L is -0.8%).[32] Note also that the median estimate of λ_H is decreasing with the quartiles; the reason is that while R_w is approximately constant, stocks with higher market capitalization are less volatile.

[32] The median estimate of λ_L is higher (in absolute value) than that of λ_H because across stocks the median δ is greater than 0.5 and, by the martingale assumption, this implies a larger downward movement.

6.2. Threshold Tax Rates

Proposition 1 defines two tax rate thresholds, $\bar{\rho}^N$ and $\bar{\rho}^I$, such that for any tax rate $\rho > \bar{\rho}^N$ no trader would ever trade, and for any tax rate $\rho \in (\bar{\rho}^I, \bar{\rho}^N)$ only noise traders would ever trade. Given our estimates, we can compute these thresholds for each stock in our sample. Table 4 reports summary statistics on the thresholds for the entire sample and by quartile. The median tax rate that would shut down all trading activity, $\bar{\rho}^N$, is 100 bps. The median threshold is 198 bps in the first quartile and 68 bps in the fourth. This reflects the fact that the volatility due to private information is higher in the lower quartiles (the parameter λ_H decreases with the quartiles). Perhaps, *a priori*, one would expect that trading in smaller stocks is more likely to be shut by an FTT, but actually the opposite happens. In any case, even in the fourth quartile, the tax rate that shuts down trading activity is much higher than most FTTs actually implemented (e.g., in France or Italy, FTTs are in the range of $10 - 30$ bps).

The median tax rate that shuts information-based trading, $\bar{\rho}^I$, is 87 bps; $\bar{\rho}^I$ ranges from a median of 115 bps in the first quartile to a median of 68 bps in the fourth. Note that this is lower than $\bar{\rho}^N$ for the first two quartiles but equal to $\bar{\rho}^N$ for the last two. This reflects the fact that for the last two quartiles, beliefs are unbounded for all stocks ($\tau > 1$), and with unbounded beliefs $\bar{\rho}^I = \bar{\rho}^N$. Finally, though the median tax rate that prevents information-based trading is high for all quartiles, there are also stocks for which informed trading is eliminated by a very low tax rate: for instance, the threshold is just higher than 5 bps for one stock in each of the first two quartiles, that is, in stocks with relatively low market capitalization. The stocks for which informed trading activity is unprofitable even with a small tax rate are those with a very low estimate of τ.

6.3. Asymptotic Impact of the FTT

Out of the 60 stocks in our sample, 55 have $\tau > 1$ (unbounded beliefs). For these stocks, with any tax rate lower than $\bar{\rho}^I$, the market asymptotically learns the true asset realization; for a tax rate higher than $\bar{\rho}^I$, no learning occurs. For the remaining 5 stocks, instead, $\tau < 1$, that is, beliefs are bounded; as we proved in Section 3, with bounded beliefs, convergence of the price to the asset value does not occur for any FTT. Proposition 3 provides two thresholds, δ_l and δ_h, such that, in the presence of a tax, informed traders stop trading whenever the market maker's belief on the good event (δ_t^d) is either below the lower threshold or above the upper threshold. Given these thresholds, the proposition provides bounds on the probability that the market maker's belief on whether the event is good remains stuck at a high level when the event is bad, or at a low level when the event is good; when the market maker belief is stuck, so is the price. We refer to these cases as cases of "wrong convergence," since the belief (and the price) converge to a high level whereas the asset value is low or vice versa.

In Table 5, we report the thresholds and the probabilities of wrong convergence on a good-event and on a bad-event day for the sample of 5 stocks with bounded beliefs. Threshold and probabilities are computed at the median parameter values for these 5 stocks, for $\alpha_t^d = 1$ (see Proposition 3) and for an FTT of 5, 10 and 20 bps. The thresholds and probabilities for individual stocks are reported in the Appendix A.6.

bps	δ_l	δ_h	$\Pr(\delta_\infty^d = \delta_l \mid v_H^d)$	$\Pr(\delta_\infty^d = \delta_h \mid v_L^d)$
5	0.042	0.956	0.020	0.093
10	0.070	0.927	0.037	0.118
20	0.222	0.770	0.059	0.574

Table 5: **Belief convergence** Belief thresholds and probabilities of wrong convergence for median parameter values across stocks with estimated $\tau < 1$.

For a 5 bps tax, the lower threshold is 0.04 and the upper threshold is 0.96; that is, even for a small tax of 5bps, the price does not converge to a value very close to the fundamental one. For such a tax rate, the probability of wrong convergence is 0.02 on a good-event day and 0.09 on a bad-event day. With a 10-bps tax, the thresholds are 0.07 and 0.93 and the corresponding probabilities are 0.04 and 0.12. For a 20 bps tax, they are 0.22 and 0.77 with probabilities 0.06 and 0.57. The probability of wrong convergence increases with the tax rate and, for a given tax rate, is higher on bad-event days than on good-event days. The reason for this asymmetry between good and bad-event days is that across these 5 stocks, the median estimate of δ is higher than 0.5; that is, the trader's and the market maker's prior beliefs are closer to the true asset value on good days, lowering the likelihood of a wrong convergence.[33]

Overall, our analysis shows that even when an FTT does not prevent all informed traders from trading, in stocks where beliefs are bounded, the price may still end up not reflecting the information held by these traders; indeed, even when the tax rate is small, there is a sizeable probability of wrong convergence (and, therefore, of the price being stuck far away from the asset value).

6.4. Within Day Effects of an FTT

Whereas one can derive the asymptotic effects of an FTT analytically, it is difficult to obtain analytical results for the impact of the tax on intra-day trading activity. This is a difficulty that our model shares with any Glosten and Milgrom (1985) type of model.[34] Since the impact of a tax is hard to tease out analytically, we proceed by simulating the model. In

[33] The bounds computed analytically according to Proposition 3 are virtually identical to the proportions of wrong convergence we obtain by simulating the model with our parameter estimates. Recall that the reason Proposition 3 only provides bounds on the probabilities is that the thresholds may be reached before the occurrence of the event is completely learned, that is, before α_t^d has converged to 1. Simulation results, however, indicate that at the median parameters for $\tau < 1$, this does not occur.

[34] Indeed, because of this difficulty, Dupont and Lee (2007) and Sørensen (2017) study the impact of an FTT restricting their analysis to a static (one-period) version of the Glosten and Milgrom (1985) model.

particular, we take the median parameters for all the stocks in our sample and, separately, for all the stocks in each quartile, and for the five stocks whose $\tau < 1$; we simulate the model with these sets of parameter values with no FTT and with an FTT of 5 bps, 10 bps, and 20 bps. The simulations are run for $100,000$ days, each consisting of $25,660$ trading times (decisions), the median number of decisions in our sample. In the subsections below, we report the effect of the FTT on the trading volume, market composition, liquidity (the bid-ask spread), volatility, informational efficiency and welfare. For each variable we report the average computed across simulated days and its standard error.[35]

6.4.1. Impact of an FTT on Volume

	0bps	5bps	Δ	10bps	Δ	20bps	Δ
Q1	727	479	−0.34	424	−0.42	358	−0.51
Q2	2,809	2,376	−0.15	2,063	−0.27	1,601	−0.43
Q3	4,667	3,822	−0.18	3,136	−0.33	2,151	−0.54
Q4	8,254	7,276	−0.12	6,296	−0.24	4,582	−0.44
$\tau < 1$	1,656	789	−0.52	709	−0.57	635	−0.62
all stocks	3,661	3,151	−0.14	2,748	−0.25	2,125	−0.42

Table 6: **Volume** Average daily trades by market capitalization quartile for different levels of the tax rate. Δ represents the change relative to the no FTT (0bps) volume. Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 3.

Table 6 reports the average number of daily trades (buys or sells) for different tax rates. The FTT lowers the volume of trade substantially. Across all stocks, at the median parameters, even a 5 bps FTT reduces the volume of trade by approximately 14%. The effect of an FTT is much larger for the first quartile (a 34% reduction in volume for a 5 bps tax) than

[35] For each day, we set $v^{d-1} = \frac{100(1-\delta)}{\lambda_H}$, which guarantees that each day $v_H - v_L = 100$, that is, the price range equals 100. We chose to do so in order to present many of the simulation results (e.g., the bid and ask spread or the average distance of the price from the fundamental) as percentages.

for the others. The impact of the FTT is also very large in those stocks with $\tau < 1$.[36] The higher the tax rates, the higher the reduction in trading volume: for the median parameters across all stocks, a 20 bps FTT reduces the volume of trade by approximately 42%. These results are in line with the strong negative impact on trading volume found in the empirical literature on the FTT, in particular for less liquid stocks. For instance, Colliard and Hoffmann (2017) find an average reduction of 10% in trading volume in response to a 20 bps tax on French stocks; they find a stronger effect on less liquid stocks (-20%) and a very strong immediate effect for the average French stock (-60%) in the first month of the tax. Umlauf (1993) finds a reduction of trading volume of 60% on the Stockholm stock exchange in response to a 100 bps tax in Sweden. Deng et al. (2018) find a reduction of 55% in trading volume in Chinese stocks in response to a 30 bps stamp duty.

6.4.2. Impact of an FTT on Market Composition

Table 7 reports the proportion of trades (buys and sells) by informed traders on event days for different tax rates. Table 7 also reports the proportion of buys on good-event days and sells on bad-event days by informed traders.

Given that across all stocks and for each quartile, τ is bigger than ν, that is, noise traders are more price elastic than informed traders, one may have expected the tax to increase the proportion of informed trading substantially. It is indeed true that the proportion of informed buys on good-event days and of informed sells on bad-event days increases in the tax rate. However, when we look at the proportion of informed buys and sells together, we find that this increase is modest (and actually does not occur for the first quartile).

The reason is the following. As the price converges to the fundamental value, most

[36] For these stocks, not surprisingly, informed trading is highly reduced by the tax, due to the high price elasticity of informed traders. Moreover, given those stocks' parameter values, on no-event days, the market maker learns quickly that an event has not occurred, with the result that noise trading also becomes quickly unprofitable, leading to the overall reduction in volume in the presence of an FTT.

informed traders receive signals that agree with the history of trades; such signals move their asset valuation by little with respect to that of the market maker, and they may abstain from trading if they have to pay an FTT. This is *a fortiori* true for those noise traders who receive pseudo signals in agreement with the history of trades (since their pseudo signals are less precise). Recall, however, that noise traders' pseudo signals are independent of the asset value (their distribution is uniform). Therefore, even as the price converges to the fundamental value, a large proportion of noise traders receive pseudo signals at odds with the past history of trades. These pseudo signals move their asset valuation significantly away from that of the market maker; these noise traders, who are on the "wrong" side of the market, may find it optimal to trade even in the presence of a tax. As a result, the fraction of noise traders impacted by the tax is smaller, thereby muting the overall impact of the tax on market composition.

Finally, for the stocks with $\tau < 1$, the proportion of informed trading monotonically declines with the introduction of an FTT, from more than 50% without a tax to only 14% with a 20 bps tax. For these stocks, informed traders' signal is less precise than the noise traders' pseudo signal ($\tau < \nu$), and, therefore, informed traders are more affected by the tax.

	informed trading				*informed buying*				*informed selling*			
	0bps	5bps	10bps	20bps	0bps	5bps	10bps	20bps	0bps	5bps	10bps	20bps
Q1	0.081	0.060	0.056	0.051	0.192	0.193	0.196	0.206	0.196	0.195	0.199	0.207
Q2	0.021	0.022	0.022	0.023	0.045	0.050	0.054	0.062	0.045	0.050	0.054	0.061
Q3	0.014	0.014	0.015	0.015	0.027	0.031	0.035	0.043	0.027	0.031	0.035	0.043
Q4	0.007	0.008	0.008	0.009	0.014	0.015	0.017	0.021	0.014	0.015	0.017	0.021
$\tau < 1$	0.562	0.315	0.245	0.142	0.578	0.294	0.223	0.141	0.578	0.374	0.301	0.178
all stocks	0.017	0.018	0.018	0.018	0.035	0.039	0.042	0.048	0.035	0.038	0.041	0.047

Table 7: **Market composition** Average daily share of informed trading on event days (left panel), average daily share of informed buying on good-event days (middle panel) and average daily share of informed selling on bad-event days (right panel). Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 0.0005.

6.4.3. Impact of an FTT on the Bid-Ask Spread

Theory does not give us a prediction on the effect of the FTT on the spread, a standard measure of liquidity. Suppose that only informed traders were price elastic, as in Cipriani and Guarino (2014). In this case, the FTT would crowd out a fraction of informed traders; this would reduces adverse selection, and the market maker would set a lower spread to make zero expected profits. However, with fewer informed traders, learning would be slower and, therefore, the spread would converge to zero at a slower pace. Furthermore, in our sample of stocks, noise traders are price-elastic too; therefore, an FTT also affects their willingness to trade. As a result, the theoretical impact of the FTT on the average daily spread is ambiguous.

In Table 8, we show simulation results on the impact of the FTT on the average daily bid-ask spread for our sample of stocks.[37] First, for any tax rate, the spread decreases through the quartiles, reflecting the fact that in higher quartiles there is less information-based trading, which makes the stocks more liquid. Second, for the median parameters across all stocks, the spread increases, though only moderately, with the tax rate. This is consistent with the fact that as we showed above, the proportion of informed traders is higher with a higher FTT. We observe the same behavior for the last three quartiles. In the first quartile, the spread is slightly decreasing, consistent with the fact that the proportion of informed traders decreases with the tax. Finally, for the stocks with $\tau < 1$ the spread slightly increases for 5 bps and then decreases. A decrease in the spread reflects the fall in information-based trading. As we observed, though, a lower proportion of informed traders, especially with a relatively imprecise signal, also means a lower speed of learning, with a consequent higher spread. This is the reason the spread increases for a 5 bps FTT.

[37] In the simulations, the daily price range is always 100, e.g., 2 can be read as a spread of 2% of the maximum possible daily change in asset value due to private information. The price change due to private information is only part of the overall price change, which includes public information (see Section 6.1)

	0bps	5bps	10bps	20bps
Q1	2.212	2.193	2.164	2.111
Q2	1.131	1.214	1.271	1.328
Q3	0.888	0.979	1.049	1.099
Q4	0.516	0.570	0.629	0.715
$\tau < 1$	0.786	0.804	0.791	0.631
all stocks	0.989	1.061	1.117	1.180

Table 8: **Bid-ask spread.** Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 0.005.

6.4.4. Impact of an FTT on Volatility

In Table 9, we report the average annualized standard deviation of intraday log-price changes (i.e., log-price changes between trading times during a day), where we define the price at time t of day d as $p_t^d = \mathbb{E}\left(V^d | h_t^d\right)$. For the median parameters across all stocks and across each quartile except the first, volatility increases with the tax rate. In the first quartile and for the stocks with $\tau < 1$ volatility decreases with the tax rate. The reason for this difference is that in the first quartile and for $\tau < 1$, the FTT displaces informed traders to a greater extent than noise traders, thereby reducing the price impact of trades.[38] In the existing empirical literature, the results on the impact of an FTT on price volatility are heterogeneous: whereas Umlauf (1993) and Jones and Seguin (1997) have found a positive relationship between transaction costs and price volatility, more recent work by Deng et al. (2018), focusing on cross-listed Chinese stocks, finds that an FTT reduces price volatility, as Tobin (1978) had envisioned; Colliard and Hoffmann (2017), instead, find no significant effect of an FTT on volatility in France. Our results help to explain this conflicting evidence, since they show that the FTT can affect the market composition in different ways, and this in turn affects price volatility differently.

[38] While in some work, e.g., Stiglitz (1989) and Summers and Summers (1989), volatility is attributed to excessive noise trading, in our model price volatility is the result of private information.

	0bps	5bps	10bps	20bps
Q1	17.917	17.906	17.810	17.609
Q2	8.858	9.114	9.289	9.433
Q3	4.789	4.977	5.130	5.219
Q4	3.689	3.844	4.012	4.258
$\tau < 1$	28.023	25.907	23.203	16.078
all stocks	7.762	7.984	8.157	8.331

Table 9: **Volatility** Annualized intraday standard deviation of log prices changes. Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 0.01.

6.4.5. The Impact of an FTT on Informational Efficiency

In this section, we investigate how the FTT affects the ability of the market to aggregate private information. Informational efficiency is high when the market maker's expected value of the asset, $p_t^d = \mathbb{E}\left(V^d | h_t^d\right)$ (which we have defined as the price of the asset) is close to the asset value. In order to estimate the market's informational efficiency, we compute the average distance between the price and the asset value v^d, that is, $|p_t^d - v^d|$, throughout the trading day. A higher value of this average distance indicates a lower informational efficiency. The impact of an FTT on informational efficiency is theoretically ambiguous. One could imagine that, by lowering the incentive of informed traders to trade, an FTT would also lower informational efficiency. This however is not necessarily true: no trades, not only buys and sells, also convey information to the market maker. To complicate matters further, the FTT also impacts the behavior of noise traders, and the market maker endogenously changes the bid-ask spread; as a result, market composition changes and so does informational efficiency.

As reported in Table 10, our simulation results show that, for all tax rates, informational efficiency is lower in stocks with higher market capitalization. We know that informed trading monotonically decreases in the quartiles (see Tables 2 and 7); it is, therefore, not surprising that the price is closer to the fundamental value in the first quartiles.

Across all stocks and for the last three quartiles, a higher FTT improves informational

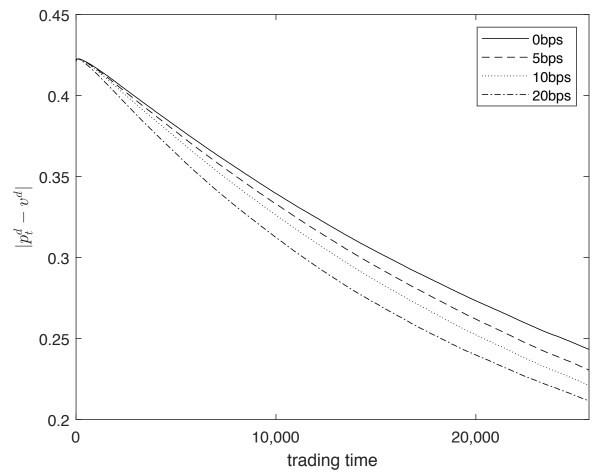

Figure 1: Distance between price and fundamental, $|p_t^d - v^d|$, for the median parameters across all stocks by trading time.

efficiency although the effect is not very large. The impact of the tax is different for the five stocks with $\tau < 1$ and for the first quartile: for these stocks, by reducing informed trading activity, the tax has a negative impact on informational efficiency.

	0bps	5bps	10bps	20bps
Q1	14.415	13.985	13.925	14.199
Q2	29.715	28.898	28.180	27.151
Q3	35.611	34.635	33.523	32.162
Q4	40.570	40.078	39.467	38.245
$\tau < 1$	10.331	12.834	15.912	21.909
all stocks	32.440	31.684	30.992	29.873

Table 10: **Informational inefficiency** Distance between price and fundamental, $|p_t^d - v^d|$, averaged over the trading day. Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 0.05.

It is also interesting to look at the evolution of informational efficiency over the course of the day. In Figure 1 we plot $|p_t^d - v^d|$ for the median parameters, averaged across all simulated days, by trading time. Consistently with what we saw in Table 10, an FTT increases informational efficiency. The higher the tax rate we consider, the higher the efficiency. Finally, informational efficiency increases monotonically during the day.

6.4.6. Impact of an FTT on Welfare

Any transaction between a noise trader and the market maker realizes a gain from trade equal to the absolute value of the difference between the noise trader's valuation of the asset (which includes a private value) $\mathbb{E}(V^d|h_t^d, n_t^d)$ and the market maker's, $\mathbb{E}(V^d|h_t^d)$.[39] Even in the absence of a tax, not all gains from trade are realized because the bid-ask spread prevents some noise traders with a gain from trade from trading (whenever their valuation falls into the spread). The introduction of the FTT increases the price noise traders pay, thereby, everything else being equal, reducing their trading activity and the market's allocative efficiency. The tax may also impact the bid-ask spread, thereby making its effect on allocative efficiency theoretically ambiguous; in our sample of stocks, however, we know that the spread increases in the tax rate (with the exception of the first quartile and for those stocks for which $\tau < 1$) and therefore, we can expect the FTT to have a negative impact on welfare.

In Table 11 (left panel) we report the fraction of noise traders receiving a pseudo signal (and therefore having a private value) who decide to trade. At the median parameters across all stocks, without an FTT, almost all these traders trade, since the probability of informed trading is low and a small spread is enough for the market maker to have zero expected profits: in the first quartile, where liquidity is lowest, 95% of them trade, and this proportion reaches 99% in the fourth quartile. An FTT decreases the fraction of noise traders trading, monotonically in the tax rate, across all stocks and for all quartiles: at the

[39] In contrast, transactions between the market maker and informed traders do not realize any gain from trade as informed traders do not have private values and value the asset the same as the market maker. Remember that in Section 2 we explained that noise traders' activity can stem from hedging reasons or bounded rationality. In the first case, the difference between the noise traders' and the market maker's valuation is clearly a gain from trade. If instead, noise traders are irrational, one could view their trading as a wash from a welfare perspective (their expected loss is the market marker's expected profit); one could also take a more libertarian view (see, e.g., Glosten (2009)) and consider their trading as welfare increasing (because, though irrationally, they do want to trade the asset).

median parameters across all stocks, the proportion of noise traders with a pseudo signal who trade goes from 98% without a tax to 57% with a 20bps tax.

In the right panel of Table 11, we report the welfare loss for different levels of the tax rate. In computing the welfare loss, we take into account not only whether a trade occurs, but also the size of the foregone gain from trade (i.e., the difference between the trader's and the market maker's valuation). In particular, we measure the loss as the average ratio between the daily forgone gains from trade and the daily potential gains from trade. In the absence of an FTT, the loss in welfare due the bid-ask spread is negligible; only in the first quartile, the least liquid one, can one observe a very modest welfare loss (0.1%). An FTT lowers gains from trade for all quartiles. At the median parameters across all stocks, the welfare loss is 0.9% for a tax rate of 5 bps and reaches 8.3% for a tax rate of 20 bps. Across tax rates, the loss is generally bigger in the top two quartiles; that is, the tax has a greater impact on welfare in more liquid stocks.

Finally, note that for the stocks with $\tau < 1$ the impact of an FTT on welfare is very strong even with a small tax. The reason is that, as we discuss in Section 6.4.1, the tax heavily impacts noise traders' trading activity. Moreover, for these stocks, a tax rate of 5 bps has a stronger impact than a tax rate of 20 bps: a 20 bps tax displaces a larger proportion of informed traders (see Table 7), which lowers the spread (see Table 8), thereby mitigating the impact of the tax on noise traders.

Finally, an important question is the FTT effectiveness in raising revenues. In Table 12, we report the average ratio of the daily forgone gains from trade over the daily tax revenues, a measure of the deadweight loss caused by the tax. The cost of imposing an FTT increases in the tax rate. At the median parameters, a 5bps tax is associated with a 10% deadweight loss, whereas a 20bps is associated with a 36% deadweight loss.

	0bps	5bps	10bps	20bps
Q1	0.946	0.634	0.564	0.478
Q2	0.980	0.828	0.719	0.557
Q3	0.989	0.809	0.664	0.455
Q4	0.995	0.876	0.758	0.551
$\tau < 1$	0.982	0.555	0.554	0.572
all stocks	0.985	0.847	0.739	0.571

	0bps	5bps	10bps	20bps
Q1	0.001	0.012	0.026	0.054
Q2	0.000	0.009	0.027	0.079
Q3	0.000	0.017	0.055	0.161
Q4	0.000	0.010	0.038	0.135
$\tau < 1$	0.000	0.408	0.405	0.369
all stocks	0.000	0.009	0.027	0.083

Table 11: **Welfare** Average fraction of noise traders with a pseudo-signal who trade (left panel) and average daily ratio of foregone gains from trade to total available gains from trade for noise traders (right panel). Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 0.005 (left panel) and 0.0001 (right panel).

	5bps	10bps	20bps
Q1	0.311	0.363	0.454
Q2	0.117	0.203	0.375
Q3	0.124	0.251	0.531
Q4	0.074	0.163	0.394
$\tau < 1$	0.771	0.769	0.756
all stocks	0.100	0.183	0.359

Table 12: **Deadweight loss** Ratio of average daily foregone gains from trade for noise traders to average daily tax revenue (all days). Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. All standard errors of simulation means are less than 0.0005.

7. Conclusion

We presented a novel methodology to quantify the impact of an FTT on the informational efficiency of financial markets. We built a sequential trading model in which both informed traders and noise traders are price-elastic. Informed traders receive private information of heterogeneous quality, and noise traders have heterogeneous private values (for instance, stemming from liquidity or hedging reasons). Compared to previous work in market microstructure, price-elastic noise traders are a modelling innovation that is crucial to understand the differential impact of an FTT on market composition.

In our model, as in most of the theoretical work on the FTT, the impact of an FTT on market outcomes is theoretically ambiguous and depends on the model's parameters. In order to resolve this ambiguity, we estimated the model through maximum likelihood for a sample of 60 stocks traded on the NYSE, with different levels of market capitalization. The structural estimation allowed us to estimate the probability of (good or bad) informational events, the composition of the market in terms of informed and noise trading, as well as the traders' price sensitivity. For most of the stocks in our sample, noise traders are price elastic, more so than informed traders. As a result, the introduction of an FTT changes the composition of the market, making the order flow more informative, that is, an FTT has a positive effect on informational efficiency. Moreover, by reducing noise trading, an FTT has a negative effect on welfare; this effect is monotonically increasing in the tax rate.

References

Aquilina, M., Budish, E.B., O'Neill, P., 2021. Quantifying the high-frequency trading "arms race". The Quarterly Journal of Economics .

Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. American Economic Review 88, 724–748.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy 100, 992–1026.

Chung, K.H., Li, M., McInish, T.H., 2005. Information-based trading, price impact of trades, and trade autocorrelation. Journal of Banking & Finance 29, 1645–1669.

Cipriani, M., Guarino, A., 2008a. Herd behavior and contagion in financial markets. The BE Journal of Theoretical Economics 8.

Cipriani, M., Guarino, A., 2008b. Transaction costs and informational cascades in financial markets. Journal of Economic Behavior & Organization 68, 581–592.

Cipriani, M., Guarino, A., 2014. Estimating a structural model of herd behavior in financial markets. American Economic Review 104, 224–51.

Citanna, A., Polemarchakis, H.M., Tirelli, M., 2006. The taxation of trades in assets. Journal of Economic Theory 126, 299–313.

Colliard, J.E., Hoffmann, P., 2017. Financial transaction taxes, market composition, and liquidity. The Journal of Finance 72, 2685–2716.

Dávila, E., Parlatore, C., 2021. Trading costs and informational efficiency. The Journal of Finance 76, 1471–1539.

Deng, Y., Liu, X., Wei, S.J., 2018. One fundamental and two taxes: When does a tobin tax reduce financial price volatility? Journal of Financial Economics 130, 663–692.

Dupont, D.Y., Lee, G.S., 2007. Effects of securities transaction taxes on depth and bid-ask spread. Economic Theory 31, 393–400.

Easley, D., Kiefer, N.M., O'Hara, M., 1997. One day in the life of a very common stock. The Review of Financial Studies 10, 805–835.

Easley, D., Kiefer, N.M., O'Hara, M., Paperman, J.B., 1996. Liquidity, information, and infrequently traded stocks. The Journal of Finance 51, 1405–1436.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69–90.

Easley, D., de Prado, M.L., O'Hara, M., 2016. Discerning information from trade data. Journal of Financial Economics 120, 269–285.

Easley, D., López de Prado, M.M., O'Hara, M., 2012. Flow toxicity and liquidity in a high-frequency world. The Review of Financial Studies 25, 1457–1493.

Edwards, F.R., 1993. Taxing transactions in futures markets: objectives and effects. Journal of Financial Services Research 7, 75–91.

Gale, D., 1996. What have we learned from social learning? European Economic Review 40, 617–628.

Gervais, S., 1997. Market microstructure with uncertain information precision: A multi-period analysis. WP .

Glosten, L.R., 2009. Welfare costs of informed trade. WP .

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Glosten, L.R., Putniņš, T.J., 2015. Welfare costs of informed trade. WP .

Glosten, L.R., Putniņš, T.J., 2019. Welfare costs of informed trade. WP .

Habermeier, K., Kirilenko, A.A., 2003. Securities transaction taxes and financial markets. IMF Staff Papers 50, 165–180.

Hasbrouck, J., 1991. The summary informativeness of stock trades: An econometric analysis. The Review of Financial Studies 4, 571–595.

Hasbrouck, J., 2009. Trading costs and returns for us equities: Estimating effective costs from daily data. The Journal of Finance 64, 1445–1477.

Hendershott, T., Menkveld, A.J., 2014. Price pressures. Journal of Financial Economics 114, 405–423.

Hirshleifer, D., Hong Teoh, S., 2003. Herd behaviour and cascading in capital markets: A review and synthesis. European Financial Management 9, 25–66.

Holden, C.W., Jacobsen, S., 2014. Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. The Journal of Finance 69, 1747–1785.

Hollifield, B., Miller, R.A., Sandås, P., Slive, J., 2006. Estimating the gains from trade in limit-order markets. The Journal of Finance 61, 2753–2804.

Jones, C.M., Seguin, P.J., 1997. Transaction costs and price volatility: evidence from commission deregulation. American Economic Review 87, 728–737.

Kupiec, P.H., 1996. Noise traders, excess volatility, and a securities transactions tax. Journal of Financial Services Research 10, 115–129.

Lee, C.M., Ready, M.J., 1991. Inferring trade direction from intraday data. The Journal of Finance 46, 733–746.

Lee, I.H., 1998. Market crashes and informational avalanches. The Review of Economic Studies 65, 741–759.

Mannaro, K., Marchesi, M., Setzu, A., 2008. Using an artificial financial market for assessing the impact of tobin-like transaction taxes. Journal of Economic Behavior & Organization 67, 445–462.

Mogensen, P.K., Riseth, A.N., 2018. Optim: A mathematical optimization package for Julia. Journal of Open Source Software 3, 615.

Noussair, C., Robin, S., Ruffieux, B., 1998. The effect of transaction costs on double auction markets. Journal of Economic Behavior & Organization 36, 221–233.

Pellizzari, P., Westerhoff, F., 2009. Some effects of transaction taxes under different microstructures. Journal of Economic Behavior & Organization 72, 850–863.

Schwert, G.W., Seguin, P.J., 1993. Securities transaction taxes: an overview of costs, benefits and unresolved questions. Financial Analysts Journal 49, 27–35.

Smith, L., Sørensen, P., 2000. Pathological outcomes of observational learning. Econometrica 68, 371–398.

Song, F.M., Zhang, J., 2005. Securities transaction tax and market volatility. The Economic Journal 115, 1103–1120.

Sørensen, P., 2017. The financial transactions tax in markets with adverse selection. WP .

Stiglitz, J.E., 1989. Using tax policy to curb speculative short-term trading. Journal of Financial Services Research 3, 101–115.

Subrahmanyam, A., 1998. Transaction taxes and financial market equilibrium. The Journal of Business 71, 81–118.

Summers, L.H., Summers, V.P., 1989. When financial markets work too well: A cautious case for a securities transactions tax. Journal of Financial Services Research 3, 261–286.

Tobin, J., 1978. A proposal for international monetary reform. Eastern Economic Journal 4, 153–159.

Umlauf, S.R., 1993. Transaction taxes and the behavior of the swedish stock market. Journal of Financial Economics 33, 227–240.

Welch, I., 1992. Sequential sales, learning, and cascades. The Journal of Finance 47, 695–732.

Appendices

A. Data and additional empirical results

A.1. Sample of stocks

company name	ticker	price	spread	volume	market cap
GRAHAM CORP	GHM	22	5	8,479	215
VISHAY PRECISION GROUP INC	VPG	19	53	5,002	230
VERSO CORP	VRS	7	85	128,174	237
NATRL GROCERS BY VIT COTTAGE INC	NGVC	12	25	18,159	267
GENESIS HEALTHCARE INC	GEN	4	24	31,933	319
TRECORA RESOURCES	TREC	14	72	9,824	339
RUBICON PROJECT INC	RUBI	7	13	76,351	364
CHANNELADVISOR CORP	ECOM	14	35	20,181	370
ROADRUNNER TRANS SYSTEMS INC	RRTS	10	10	60,963	398
LUMBER LIQUIDATORS HOLDINGS INC	LL	16	6	151,964	429
NACCO INDUSTRIES INC	NC	91	17	2,846	471
MARCUS CORP	MCS	32	16	18,360	596
CBIZ INC	CBZ	14	37	27,120	735
CIVITAS SOLUTIONS INC	CIVI	20	25	34,438	741
MANITOWOC CO INC	MTW	6	17	373,371	830
Q1 mean		19	28	360,109	470
Q1 median		14	20	106,047	384
DYNEGY INC NEW DEL	DYN	8	12	793,990	992
FORTRESS TRANS AND INFR INV LLC	FTAI	13	23	22,093	1,007
CALLAWAY GOLF CO	ELY	11	9	128,602	1,031
BARRACUDA NETWORKS INC	CUDA	21	5	85,205	1,121
HERC HOLDINGS INC	HRI	40	10	80,134	1,137
TUTOR PERINI CORP	TPC	28	18	85,459	1,378
AZZ INC	AZZ	64	8	30,417	1,662
PARTY CITY HOLDCO INC	PRTY	14	35	100,568	1,697
SELECT MEDICAL HOLDINGS CORP	SEM	13	38	407,678	1,753
HERTZ GLOBAL HOLDINGS INC	HTZ	22	5	461,409	1,789
ARCH COAL INC	ARCH	78	5	140,574	1,919
KBR INC	KBR	17	6	269,895	2,381
COMMERCIAL METALS CO	CMC	22	5	310,960	2,517
ESTERLINE TECHNOLOGIES CORP	ESL	89	11	30,241	2,644
KNIGHT TRANSPORTATION INC	KNX	33	15	407,602	2,645
Q2 mean		33	14	489,201	1,763
Q2 median		22	10	384,707	1,725

AMERICAN EAGLE OUTFITTERS INC NE	AEO	15	7	1,165,330	2,759
LEGG MASON INC	LM	30	3	552,197	3,021
ALLETE INC	ALE	64	2	61,127	3,175
ONE GAS INC	OGS	64	6	39,385	3,342
ENERSYS	ENS	78	1	60,850	3,392
RADIAN GROUP INC	RDN	18	6	549,523	3,855
VEEVA SYSTEMS INC	VEEV	41	2	224,677	4,119
TEGNA INC	TGNA	21	5	356,932	4,586
SONOCO PRODUCTS CO	SON	53	2	95,252	5,262
PERKINELMER INC	PKI	52	2	142,101	5,713
PINNACLE FOODS INC	PF	53	2	555,105	6,310
JACOBS ENGINEERING GROUP INC	JEC	57	2	219,421	6,887
NISOURCE INC	NI	22	5	618,225	7,145
SPECTRUM BRANDS HOLDINGS INC	SPB	122	1	105,448	7,299
ROLLINS INC	ROL	34	3	83,008	7,358
Q3 mean		48	3	860,786	4,948
Q3 median		52	2	555,105	4,586
RITE AID CORP	RAD	8	12	3,502,339	8,670
TORCHMARK CORP	TMK	74	1	106,012	8,752
CBRE GROUP INC	CBG	31	3	501,436	10,621
PRINCIPAL FINANCIAL GROUP INC	PFG	58	2	316,898	16,645
AMERISOURCEBERGEN CORP	ABC	78	4	407,548	17,210
EVERSOURCE ENERGY	ES	55	2	301,283	17,519
AMERIPRISE FINANCIAL INC	AMP	111	1	221,164	17,534
MOODYS CORP	MCO	94	1	203,561	18,024
INTERNATIONAL PAPER CO	IP	53	2	563,930	21,819
HP INC	HPQ	15	7	2,031,599	25,309
DEERE & CO	DE	103	1	584,182	32,651
MARSH & MCLENNAN COS INC	MMC	68	1	486,171	34,849
BECTON DICKINSON & CO	BDX	166	1	201,132	35,301
PEPSICO INC	PEP	105	1	859,812	150,059
CHEVRON CORP NEW	CVX	118	1	1,345,075	222,190
Q4 mean		76	3	2,583,087	42,477
Q4 median		74	1	1,409,032	18,024
All stocks mean		44	12	1,071,144	12,393
All stocks median		31	5	477,748	2,702

Table A.1: List of sample stocks in order of market capitalization. Price ($), volume (monthly; in 100 shares), bid-ask spread (in bps) and market capitalization (in $MM) as of 30 December 2016 (source: CRSP).

Table A.2 shows that the number of buys and sells increases with market capitalization. In all quartiles, the ratio between buys and sells is very close to one.

		market cap	buys	sells	% buys	trades	no trades	% trades
Q1	mean	465	549	565	49	1,115	23,068	4
	median	398	343	355	49	698	23,168	3
Q2	mean	1,670	1,899	1,895	50	3,794	22,431	14
	median	1,697	1,334	1,250	50	2,584	22,644	10
Q3	mean	5,049	2,134	2,150	50	4,285	22,248	16
	median	4,586	2,020	1,998	50	4,019	22,297	15
Q4	mean	42,477	5,108	5,146	50	10,254	20,759	31
	median	18,024	4,037	3,949	50	8,047	21,363	27
all stocks	mean	12,415	2,423	2,439	50	4,862	22,126	16
	median	2,702	1,593	1,582	50	3,175	22,485	12

Table A.2: **Daily trades** Daily trading activity by market capitalization (in $MM) quartiles .

A.2. Parameter Estimates

	ticker	α	δ	μ	τ	ν	ε
Q1	GHM	0.835	0.596	0.002	1.895	1.379	0.013
		(0.0228)	(0.0323)	(0.0001)	(0.0056)	(0.0075)	(0.0001)
	VPG	0.885	0.784	0.002	1.856	1.340	0.011
		(0.0053)	(0.0125)	(0.0001)	(0.0067)	(0.0072)	(0.0003)
	VRS	0.868	0.721	0.076	0.097	1.552	0.040
		(0.0395)	(0.1771)	(0.0083)	(0.0078)	(0.0706)	(0.0011)
	NGVC	0.442	0.448	0.002	1.947	1.423	0.032
		(0.0914)	(0.0215)	(0.0001)	(0.0032)	(0.0049)	(0.0006)
	GEN	0.943	0.638	0.002	1.925	1.427	0.038
		(0.0115)	(0.0780)	(0.0001)	(0.0120)	(0.0096)	(0.0009)
	TREC	0.651	0.698	0.002	1.989	1.451	0.012
		(0.0623)	(0.0302)	(0.0001)	(0.0016)	(0.0127)	(0.0002)
	RUBI	0.824	0.542	0.002	1.950	1.329	0.096
		(0.0371)	(0.0565)	(0.0001)	(0.0092)	(0.0165)	(0.0051)
	ECOM	0.735	0.524	0.003	1.894	1.432	0.022
		(0.0316)	(0.0564)	(0.0001)	(0.0062)	(0.0143)	(0.0003)
	RRTS	0.814	0.245	0.028	0.167	1.252	0.028
		(0.0535)	(0.0796)	(0.0032)	(0.0291)	(0.0596)	(0.0008)
	LL	0.335	0.487	0.162	0.028	1.092	0.122
		(0.1117)	(0.1636)	(0.0316)	(0.0080)	(0.0677)	(0.0029)
	NC	0.585	0.672	0.017	0.173	1.305	0.009
		(0.0971)	(0.0601)	(0.0085)	(0.0671)	(0.0378)	(0.0004)
	MCS	0.855	0.708	0.003	1.907	1.419	0.023
		(0.0213)	(0.0156)	(0.0001)	(0.0042)	(0.0046)	(0.0003)

	CBZ	0.780	0.728	0.003	1.941	1.421	0.033
		(0.0235)	(0.0234)	(0.0002)	(0.0037)	(0.0160)	(0.0010)
	CIVI	0.624	0.597	0.002	1.883	1.382	0.014
		(0.0179)	(0.0210)	(0.0000)	(0.0060)	(0.0134)	(0.0002)
	MTW	0.970	0.584	0.002	1.990	1.319	0.185
		(0.0030)	(0.0479)	(0.0001)	(0.0007)	(0.0060)	(0.0054)
Q2	DYN	0.786	0.663	0.002	1.868	1.368	0.337
		(0.0489)	(0.0334)	(0.0000)	(0.0242)	(0.0052)	(0.0046)
	FTAI	0.891	0.461	0.003	1.914	1.470	0.025
		(0.0336)	(0.0573)	(0.0002)	(0.0069)	(0.0209)	(0.0004)
	ELY	0.859	0.601	0.003	1.914	1.353	0.152
		(0.0292)	(0.0417)	(0.0001)	(0.0160)	(0.0102)	(0.0039)
	CUDA	0.190	0.725	0.151	0.033	1.959	0.110
		(0.0600)	(0.0831)	(0.0749)	(0.0117)	(0.0091)	(0.0035)
	HRI	0.978	0.505	0.002	1.948	1.299	0.081
		(0.0073)	(0.0529)	(0.0001)	(0.0042)	(0.0169)	(0.0032)
	TPC	0.647	0.630	0.002	1.977	1.334	0.074
		(0.0889)	(0.0426)	(0.0000)	(0.0019)	(0.0073)	(0.0004)
	AZZ	0.721	0.656	0.002	1.950	1.446	0.034
		(0.0261)	(0.0299)	(0.0001)	(0.0015)	(0.0020)	(0.0001)
	PRTY	0.784	0.332	0.002	1.918	1.347	0.087
		(0.0398)	(0.0478)	(0.0001)	(0.0162)	(0.0150)	(0.0029)
	SEM	0.924	0.806	0.002	1.936	1.350	0.122
		(0.0113)	(0.0269)	(0.0001)	(0.0104)	(0.0046)	(0.0013)
	HTZ	0.675	0.603	0.001	1.949	1.391	0.450
		(0.0365)	(0.1311)	(0.0001)	(0.0162)	(0.0095)	(0.0122)
	ARCH	0.990	0.698	0.002	1.980	1.321	0.101
		(0.0012)	(0.0566)	(0.0000)	(0.0017)	(0.0017)	(0.0019)
	KBR	0.743	0.041	0.002	1.793	1.375	0.239
		(0.0165)	(0.0066)	(0.0001)	(0.0238)	(0.0060)	(0.0014)
	CMC	0.822	0.471	0.002	1.919	1.385	0.280
		(0.0154)	(0.0490)	(0.0000)	(0.0106)	(0.0053)	(0.0033)
	ESL	0.872	0.718	0.002	1.959	1.319	0.053
		(0.0096)	(0.0285)	(0.0001)	(0.0027)	(0.0068)	(0.0009)
	KNX	0.942	0.570	0.003	1.891	1.350	0.226
		(0.0072)	(0.0449)	(0.0003)	(0.0175)	(0.0065)	(0.0061)
Q3	AEO	0.660	0.219	0.002	1.818	1.357	0.399
		(0.0493)	(0.0635)	(0.0001)	(0.0581)	(0.0134)	(0.0056)
	LM	0.979	0.404	0.003	1.871	1.344	0.193
		(0.0069)	(0.1047)	(0.0001)	(0.0375)	(0.0074)	(0.0045)
	ALE	0.647	0.654	0.003	1.963	1.426	0.071
		(0.0396)	(0.0379)	(0.0002)	(0.0049)	(0.0292)	(0.0006)
	OGS	0.861	0.446	0.003	1.636	1.421	0.063
		(0.0164)	(0.0772)	(0.0002)	(0.0563)	(0.0113)	(0.0007)

	ENS	0.954	0.571	0.003	1.535	1.396	0.079
		(0.0080)	(0.0228)	(0.0001)	(0.0539)	(0.0218)	(0.0017)
	RDN	0.954	0.487	0.002	1.750	1.353	0.222
		(0.0127)	(0.0629)	(0.0001)	(0.0762)	(0.0053)	(0.0029)
	VEEV	0.784	0.482	0.002	1.971	1.360	0.243
		(0.0344)	(0.0717)	(0.0002)	(0.0100)	(0.0108)	(0.0030)
	TGNA	0.922	0.493	0.002	1.956	1.359	0.295
		(0.0182)	(0.0957)	(0.0002)	(0.0044)	(0.0108)	(0.0059)
	SON	0.644	0.317	0.003	1.986	1.336	0.132
		(0.0664)	(0.0492)	(0.0001)	(0.0012)	(0.0055)	(0.0015)
	PKI	0.941	0.336	0.003	1.786	1.353	0.143
		(0.0076)	(0.0733)	(0.0001)	(0.0176)	(0.0051)	(0.0012)
	PF	0.871	0.600	0.002	1.954	1.373	0.210
		(0.0532)	(0.0324)	(0.0001)	(0.0066)	(0.0146)	(0.0024)
	JEC	0.939	0.738	0.002	1.913	1.384	0.182
		(0.0079)	(0.0339)	(0.0001)	(0.0197)	(0.0022)	(0.0015)
	NI	0.936	0.332	0.002	1.819	1.378	0.300
		(0.0188)	(0.0883)	(0.0000)	(0.0304)	(0.0058)	(0.0045)
	SPB	0.926	0.606	0.002	1.814	1.355	0.116
		(0.0109)	(0.0327)	(0.0001)	(0.0208)	(0.0202)	(0.0011)
	ROL	0.923	0.595	0.003	1.934	1.353	0.128
		(0.0112)	(0.0760)	(0.0001)	(0.0093)	(0.0056)	(0.0010)
Q4	RAD	0.862	0.189	0.002	1.922	1.474	0.622
		(0.0357)	(0.0505)	(0.0001)	(0.0228)	(0.0105)	(0.0211)
	TMK	0.959	0.451	0.003	1.983	1.344	0.129
		(0.0074)	(0.0345)	(0.0000)	(0.0012)	(0.0058)	(0.0005)
	CBG	0.805	0.641	0.002	1.872	1.385	0.293
		(0.0198)	(0.0277)	(0.0001)	(0.0241)	(0.0150)	(0.0024)
	PFG	0.984	0.696	0.002	1.881	1.375	0.238
		(0.0019)	(0.0268)	(0.0001)	(0.0215)	(0.0055)	(0.0010)
	ABC	0.867	0.818	0.002	1.774	1.390	0.322
		(0.0208)	(0.0354)	(0.0001)	(0.0399)	(0.0066)	(0.0046)
	ES	0.912	0.812	0.003	1.808	1.388	0.283
		(0.0156)	(0.0177)	(0.0002)	(0.0129)	(0.0034)	(0.0033)
	AMP	0.924	0.604	0.002	1.833	1.384	0.262
		(0.0025)	(0.0256)	(0.0000)	(0.0169)	(0.0004)	(0.0003)
	MCO	0.947	0.482	0.002	1.793	1.369	0.219
		(0.0274)	(0.1183)	(0.0001)	(0.0310)	(0.0076)	(0.0040)
	IP	0.892	0.245	0.002	1.867	1.400	0.361
		(0.0177)	(0.0231)	(0.0001)	(0.0150)	(0.0174)	(0.0035)
	HPQ	0.639	0.513	0.001	1.712	1.439	0.590
		(0.0554)	(0.0500)	(0.0001)	(0.0184)	(0.0096)	(0.0063)
	DE	0.924	0.588	0.002	1.742	1.389	0.364
		(0.0051)	(0.0594)	(0.0001)	(0.0381)	(0.0068)	(0.0031)

MMC	0.891	0.010	0.003	1.907	1.428	0.326
	(0.0084)	(0.0017)	(0.0001)	(0.0148)	(0.0091)	(0.0017)
BDX	0.843	0.500	0.002	1.823	1.399	0.320
	(0.0245)	(0.0779)	(0.0001)	(0.0290)	(0.0140)	(0.0060)
PEP	0.912	0.252	0.002	1.747	1.496	0.533
	(0.0275)	(0.0596)	(0.0001)	(0.0156)	(0.0063)	(0.0033)
CVX	0.929	0.216	0.002	1.807	1.544	0.652
	(0.0034)	(0.0729)	(0.0000)	(0.0061)	(0.0040)	(0.0028)

Table A.3: Parameter estimates for the sample of 60 NYSE stocks grouped by market capitalization quartiles. Standard errors are reported in parentheses and are calculated using the bootstrap by resampling trading days. A value of 0.0000 indicates a standard error smaller than 0.00005.

A.3. Robustness

As discussed in Section 4, when constructing our dataset, we insert a number of no-trades equal to the number of seconds without trading activity. As a robustness check, we re-estimate the model for the median stock by market capitalization in each of the four market capitalization quartiles using a no-trade interval of 0.1, 0.2, 0.5, 2 and 5 seconds. We report these estimates in Table A.4. Our estimates are robust to changes in the no-trade interval. The parameters estimates for α, δ, τ and ν are stable across no-trade intervals. The parameters μ and ε, which determine the pattern of trading activity within the day, cannot remain constant as the no-trade interval and therefore the number of no trades in the dataset increase (see, Easley et al. (1997, p. 820) and Cipriani and Guarino (2014, p. 239)): for instance, with a lower no-trade interval, there are more no trades, hence ε must increase. Nevertheless, as first proposed by Easley et al. (1997), we can compute parameter transformations that are economically meaningful and that should remain constant. The first of these transformations is the PIN discussed in Section 5. The second transformation is the proportion of informed traders with correct information on an event day, which we denote by Γ; Γ measures the informativeness of trading activity in a given stock[40]:

$$\Gamma = \frac{\mu}{\mu + \varepsilon(1-\mu)} \left(\frac{1}{2} + \frac{1}{4}\tau \right).$$

[40] In Easley et al. (1997) all informed traders have a correct signal; therefore Γ measures the proportion of informed trading; Cipriani and Guarino (2008a) changed Easley et al. (1997)'s original measure to incorporate the fact that only a proportion of informed traders in their model, like in our model, have correct information.

The third transformation is the probability of at least one trade within 1 second on a no event day, which we denote by Λ:

$$\Lambda = 1 - (1-\varepsilon)^{\frac{1000}{\text{NT interval (in ms)}}}.$$

The last three columns of Table A.4 report the PIN, Γ, and Λ for the four median stocks across 6 no-trade intervals. Overall, the results show that changing the no-trade interval does not change the predicted arrival of trades to the market, the predicted arrival of trades by informed traders with a correct signal and the probability of informed trading (PIN).

stock	NT	α	δ	μ	τ	ν	ε	PIN	Γ	Λ
ECOM	0.1	0.807	0.593	0.0003	1.911	1.430	0.002	0.094	0.111	0.022
(Q1)	0.2	0.717	0.413	0.001	1.898	1.438	0.004	0.083	0.109	0.022
	0.5	0.846	0.483	0.001	1.907	1.404	0.011	0.091	0.104	0.022
	1	0.735	0.524	0.003	1.894	1.432	0.022	0.086	0.111	0.022
	2	0.744	0.399	0.005	1.907	1.435	0.043	0.086	0.110	0.022
	5	0.705	0.455	0.013	1.903	1.443	0.105	0.079	0.106	0.022
PRTY	0.1	0.810	0.389	0.0002	1.920	1.320	0.009	0.018	0.022	0.087
(Q2)	0.2	0.740	0.434	0.0004	1.997	1.357	0.018	0.016	0.022	0.085
	0.5	0.714	0.332	0.001	1.920	1.353	0.043	0.016	0.022	0.085
	1	0.784	0.332	0.002	1.918	1.347	0.087	0.018	0.023	0.087
	2	0.693	0.361	0.004	1.928	1.367	0.166	0.015	0.021	0.087
	5	0.851	0.219	0.009	1.886	1.334	0.353	0.021	0.023	0.083
TGNA	0.1	0.977	0.492	0.0003	1.940	1.322	0.035	0.008	0.008	0.300
(Q3)	0.2	0.954	0.692	0.001	1.966	1.322	0.068	0.008	0.008	0.298
	0.5	0.954	0.725	0.001	1.922	1.326	0.162	0.008	0.008	0.298
	1	0.922	0.493	0.002	1.956	1.359	0.295	0.008	0.008	0.295
	2	0.878	0.674	0.004	1.616	1.396	0.488	0.007	0.007	0.284
	5	0.686	0.657	0.010	1.668	1.485	0.748	0.010	0.013	0.241
MCO	0.1	0.937	0.424	0.0003	1.786	1.359	0.025	0.010	0.010	0.223
(Q4)	0.2	0.956	0.357	0.001	1.840	1.362	0.049	0.010	0.010	0.224
	0.5	0.858	0.554	0.001	1.753	1.366	0.118	0.009	0.010	0.222
	1	0.947	0.482	0.002	1.793	1.369	0.219	0.010	0.010	0.219
	2	0.946	0.737	0.004	1.687	1.377	0.385	0.010	0.010	0.216
	5	0.891	0.072	0.011	1.715	1.476	0.655	0.014	0.015	0.192

Table A.4: Estimation for different no trade intervals (NT reported in seconds) for the median stock by market capitalization within each market capitalization quartile. The last three columns report PIN, Γ and Λ measures as described in Section 5 and in this appendix.

A.4. Model Fit

		buys	sells	trades
Q1	data	343	355	698
	simulation	353	374	727
Q2	data	1,334	1,250	2,584
	simulation	1,388	1,421	2,809
Q3	data	2,020	1,998	4,019
	simulation	2,334	2,333	4,667
Q4	data	4,037	3,949	8,047
	simulation	4,127	4,127	8,254
$\tau < 1$	data	575	596	1,171
	simulation	835	821	1,656
all stocks	data	1,593	1,582	3,175
	simulation	1,816	1,845	3,661

Table A.5: TAQ data and simulated transactions data. Average daily trading activity (buys, sells, total trades) in the NYSE data and in the simulated data described in section 6.4. For the TAQ data, the table reports the median of daily averages across all stocks, for each quartile, and for those stocks with $\tau < 1$. Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$.

A.5. Hasbrouck Decomposition - Calibration of the FTT

In this section we explain how we calibrate λ_H, the size of the increase of the asset value on good event days. The parameters λ_H and λ_L in our model refer to the upward and downward movements of asset values on which informed traders have private information. As we explain in the Section 6, if all price volatility from one day to the next were due to the aggregation of private information by the market price, we would be able to directly estimate λ_H and λ_L.

However, in actual financial markets, not all observed stock price movements from one day to the next are due to informed trading; rather, they may be due to the release of public information. To calibrate λ_H using stock price data, we need to isolate the variability in stock prices that is due to informed trading. To do so, we use the variance decomposition proposed by Hasbrouck (1991) that decomposes the variance of log price changes into a trade-correlated component, interpreted as the component driven by private information, and a trade-uncorrelated component.

We calculate Hasbrouck's variance decomposition using intraday quote updates and all trades that occur during continuous trading. We use a lag length of 500 for the bivariate VAR. For each stock s in our sample, we obtain an estimate of Hasbrouck's R^2 measure, which gives the fraction of the variance of log price changes that can be attributed to private-information based trading. We denote this fraction for stock s by $R_w^2(s)$; if we denote by $\sigma_p(s)$ the standard deviation of daily percentage price changes, then the standard deviation of daily price changes due to private information is $R_w(s)\sigma_p(s)$.

Our model implies that the standard deviation of daily percentage price changes equals $\lambda_H \sqrt{\frac{\alpha\delta}{1-\delta}}$. Hence for stock s we have

$$\lambda_H(s) = R_w(s)\sigma_p(s)\sqrt{\frac{1-\delta(s)}{\alpha(s)\delta(s)}}.$$

Table A.6 reports the square root of Hasbrouck's R_w^2 measure, the standard deviation of log price changes, σ, and λ_H for all 60 stocks in our sample.

	ticker	R_w	σ	λ_H	$\bar{\rho}^N$	$\bar{\rho}^I$	$\bar{\rho}$
Q1	GHM	0.47	0.02	0.01	150	150	102
	VPG	0.30	0.02	0.003	115	115	32
	VRS	0.30	0.04	0.01	211	12	11
	NGVC	0.35	0.04	0.02	207	207	168
	GEN	0.42	0.03	0.01	198	198	112
	TREC	0.30	0.02	0.005	113	113	49
	RUBI	0.45	0.04	0.02	236	236	199
	ECOM	0.39	0.03	0.01	125	125	114
	RRTS	0.41	0.03	0.03	258	23	20
	LL	0.47	0.03	0.02	234	6	6
	NC	0.50	0.09	0.04	815	98	87
	MCS	0.28	0.02	0.003	73	73	30
	CBZ	0.43	0.01	0.004	100	100	37
	CIVI	0.39	0.02	0.01	130	130	88
	MTW	0.44	0.09	0.04	503	503	359
Q2	DYN	0.45	0.04	0.01	251	251	127
	FTAI	0.27	0.01	0.004	37	37	32
	ELY	0.30	0.02	0.005	68	68	45
	CUDA	0.34	0.02	0.01	288	5	5
	HRI	0.44	0.02	0.01	111	111	109

	TPC	0.30	0.02	0.01	112	112	66
	AZZ	0.38	0.02	0.01	122	122	64
	PRTY	0.33	0.02	0.01	124	124	61
	SEM	0.37	0.02	0.004	166	166	40
	HTZ	0.47	0.05	0.02	329	329	216
	ARCH	0.43	0.02	0.01	131	131	56
	KBR	0.26	0.02	0.03	315	315	14
	CMC	0.29	0.02	0.01	85	85	75
	ESL	0.42	0.03	0.01	177	177	70
	KNX	0.39	0.02	0.01	82	82	62
Q3	AEO	0.46	0.03	0.03	271	271	76
	LM	0.31	0.01	0.01	56	56	38
	ALE	0.27	0.01	0.002	42	42	22
	OGS	0.36	0.01	0.004	40	40	32
	ENS	0.46	0.02	0.01	87	87	66
	RDN	0.35	0.02	0.01	59	59	57
	VEEV	0.42	0.02	0.01	93	93	87
	TGNA	0.28	0.03	0.01	99	99	96
	SON	0.28	0.01	0.005	49	49	23
	PKI	0.29	0.01	0.004	42	42	21
	PF	0.30	0.01	0.003	47	47	31
	JEC	0.35	0.01	0.003	78	78	28
	NI	0.37	0.01	0.005	48	48	24
	SPB	0.52	0.02	0.01	102	102	67
	ROL	0.30	0.01	0.003	39	39	27
Q4	RAD	0.60	0.05	0.06	608	608	142
	TMK	0.30	0.01	0.003	27	27	22
	CBG	0.29	0.01	0.003	53	53	30
	PFG	0.28	0.01	0.002	47	47	20
	ABC	0.43	0.02	0.003	152	152	34
	ES	0.39	0.01	0.001	64	64	15
	AMP	0.49	0.01	0.01	77	77	50
	MCO	0.46	0.01	0.005	49	49	46
	IP	0.42	0.01	0.01	86	86	28
	HPQ	0.57	0.01	0.01	99	99	94
	DE	0.46	0.01	0.005	68	68	48
	MMC	0.30	0.01	0.02	241	241	3
	BDX	0.47	0.01	0.01	50	50	50
	PEP	0.40	0.01	0.005	47	47	16
	CVX	0.43	0.01	0.01	78	78	21

Table A.6: Hasbrouck decomposition, calibration of λ_H, and tax thresholds (in bps) for individual stocks grouped by market capitalization quartiles.

A.6. Tax threshold for $\tau < 1$ stocks

	tax	δ_l	δ_h	$\Pr(\delta_\infty^d = \delta_l \mid v_H^d)$	$\Pr(\delta_\infty^d = \delta_h \mid v_L^d)$
VRS	5	0.09	0.91	0.03	0.25
	10	0.14	0.85	0.04	0.44
	20				
RRTS	5	0.04	0.96	0.12	0.01
	10	0.07	0.93	0.23	0.02
	20	0.19	0.80	0.71	0.02
LL	5	0.26	0.73	0.28	0.26
	10				
	20				
NC	5	0.01	0.99	0.004	0.02
	10	0.05	0.94	0.02	0.12
	20	0.04	0.96	0.02	0.09
CUDA	5	0.24	0.75	0.02	0.85
	10				
	20				

Table A.7: Belief thresholds and probabilities of wrong convergence for stocks with estimated $\tau < 1$: thresholds δ_l and δ_h as well as the probabilities of wrong convergence as given in Proposition 3. Empty cells indicate that there is no trading in these stocks for the indicated tax rate.

A.7. PCIN and non-approximated PIN for different rates of the FTT

Table A.8 shows the PIN by quartile (as reported also in the main text - Section 5) together with the Probability of Correct Informed Trade (PCIN),

$$\text{PCIN} = \frac{\alpha\mu}{\alpha\mu + \varepsilon(1-\alpha\mu)}\left(\frac{1}{2} + \frac{1}{4}\tau\right),$$

that is, the probability of an informed trade made by a trader with a correct signal, introduced by Cipriani and Guarino (2014). Since τ is estimated to be very high, PCIN estimates are very close to PIN estimates; the exceptions are those stocks for which τ is less than one, where the PIN is 46% and the PCIN 25%.

PIN and PCIN are approximate measures of informed trading, since they assume that all informed traders and noise traders trade independently of the price, as is the case when they are price inelastic. In Table A.9, we present an alternative measure of informed trading at time 1 that takes into account the impact of the bid-ask spread and, if present, of a tax.

	Q1	Q2	Q3	Q4	$\tau < 1$ stocks	all stocks
PIN	0.090	0.018	0.013	0.006	0.456	0.015
PCIN	0.088	0.018	0.013	0.006	0.247	0.015

Table A.8: Median PIN and PCIN across all stocks, for each quartile, and for those stocks with $\tau < 1$.

Since, for the median parameters for all quartiles, noise traders are more price elastic than informed traders, the probability of informed trading as measured by the PIN increases with the tax rate. For $\tau < 1$ stocks, informed traders are more price elastic than noise traders and the probability of informed trading decreases in the tax.

tax (bps)	Q1	Q2	Q3	Q4	$\tau < 1$ stocks	all stocks
0	0.098	0.018	0.013	0.006	0.398	0.015
5	0.102	0.017	0.014	0.006	0.337	0.016
10	0.107	0.019	0.015	0.006	0.078	0.015
20	0.129	0.023	0.019	0.008	0.000	0.017

Table A.9: Median non-approximated PIN across market capitalization quartiles, for all stocks and for $\tau < 1$ stocks.

B. The Likelihood Function

In our empirical analysis, we estimate the model's parameters through maximum likelihood using financial market data for a period with no transaction tax. Here, we provide a detailed characterization of the model's likelihood function. The derivation of the likelihood function follows a similar logic as in Cipriani and Guarino (2014).

In Section 2.4, we described the equilibrium behavior of market participants for any history of trades. As the bid and ask prices are uniquely determined by the trade sequence and thus do not contain any additional information once we condition on the order flow, we can write the likelihood function in terms of the history of trades only.

Let us denote the history of trades at the end of a trading day by $h^d := h^d_{T_d}$, where T_d is the number of trading times on day d. We denote the likelihood function by

$$\mathcal{L}(\Phi; \{h^d\}_{d=1}^D) = \Pr\left(\{h^d\}_{d=1}^D | \Phi\right)$$

where $\Phi := \{\alpha, \delta, \mu, \tau, \nu, \varepsilon\}$ is the vector of parameters.

Next recall that on day d all market participants know v^{d-1}, and the occurrence of information events is independent across days. Thus, the sequence of trades on day d only depends on the realization of V^d and not on any trading data from days other than d. We can therefore write the likelihood function as the product of the likelihoods of daily trading sequences

$$\mathcal{L}(\Phi; \{h^d\}_{d=1}^D) = \Pr\left(\{h^d\}_{d=1}^D | \Phi\right) = \Pi_{d=1}^D \Pr(h^d | \Phi).$$

Now consider the likelihood of a sequence of trades for a given day. Unlike in the standard market microstructure model of Easley and O'Hara (1987) where only the total number of buys and sells matters for the probability of a given history of trades, in our model, the sequence of trades is important. Informed and price-elastic noise traders update their valuations depending on the trading sequence, and, thus, their probability of trading depends on the observed history of trades up to the time in which they act. Therefore, we compute the likelihood function for the history of trades on day d starting at time 1 and, then, recursively up to time T_d. At trading time t the probability of a given action x_t^d depends on the sequence of previous trades h_t^d, and we have that

$$\Pr\left(h_{t+1}^d | \Phi\right) = \Pr(x_t^d | h_t^d, \Phi) \Pr\left(h_t^d | \Phi\right).$$

To compute $\Pr\left(x_t^d | h_t^d, \Phi\right)$, we express it in terms of the value-contingent trading probabilities

$$\Pr\left(x_t^d | h_t^d, \Phi\right) = \Pr\left(x_t^d | h_t^d, v_H^d, \Phi\right) \Pr\left(v_H^d | h_t^d, \Phi\right) +$$
$$\Pr\left(x_t^d | h_t^d, v_L^d, \Phi\right) \Pr\left(v_L^d | h_t^d, \Phi\right) + \Pr\left(x_t^d | h_t^d, v^{d-1}, \Phi\right) \Pr\left(v^{d-1} | h_t^d, \Phi\right).$$

We now illustrate how to compute the value-contingent probabilities of a trade x_t^d. Consider period $t = 1$ and suppose, for instance, that there was a buy order, $x_1^d = \text{buy}$. The probability of such an order for a given asset value depends on the buy thresholds for informed and price-elastic noise traders, β_1^d and κ_1^d, which are functions of the model's parameters. Having obtained the value-contingent probabilities of a buy order in period 1, we can then update the market makers' beliefs using Bayes' rule. Hence, consider period t and suppose again that x_t^d is a buy order. The equilibrium buy thresholds, β_t^d and κ_t^d, will be functions of the market maker's beliefs given the trading history up to (but not including) period t, as well as the parameters of the model.

Once we have solved for β_t^d and κ_t^d, we can compute the probability of a buy order on a good-event day:

$$\Pr\left(x_t^d = \text{buy} | h_t^d, v_H^d, \Phi\right) =$$
$$\mu\left[1 - F^H\left(\beta_t^d | v_H^d\right)\right] + (1-\mu)\varepsilon\left(1 - \kappa_t^d\right),$$

where $F^H(\cdot | v_H^d)$ is the cumulative distribution function of $f^H(\cdot | v_H^d)$. Recall that a trader active at time t is an informed trader with probability μ and a noise trader with probability $1-\mu$. An informed trader buys if his signal is above the buy threshold β_t^d, which happens with probability $1 - F^H(\beta_t^d | v_H^d)$. A noise trader receives a pseudo signal with probability ε, in which case he buys if his signal is larger than κ_t^d, which happens with probability $1 - \kappa_t^d$ (as these pseudo signals are uniformly distributed). Similarly, on a bad-event day, we have that

$$\Pr\left(x_t^d = \text{buy} | h_t^d, v_L^d, \Phi\right) =$$
$$\mu\left[1 - F^L\left(\beta_t^d | v_L^d\right)\right] + (1-\mu)\varepsilon\left(1 - \kappa_t^d\right).$$

On a no-event day ($V^d = v^{d-1}$), a market order can only come from a noise trader. Therefore,

$$\Pr\left(x_t^d = \text{buy} | h_t^d, v^{d-1}, \Phi\right) = \varepsilon\left(1 - \kappa_t^d\right).$$

By the same logic, we obtain the value-contingent probabilities for a sell order at t by computing σ_t^d and γ_t^d. The probability of a sell order on a good event day, for instance, is

$$\Pr\left(x_t^d = \text{sell}|h_t^d, v_H^d, \Phi\right) =$$
$$\mu F^H\left(\sigma_t^d|v_H^d\right) + (1-\mu)\varepsilon\gamma_t^d.$$

The probability of a no-trade is just the complement to the probabilities of a buy and of a sell. Finally, to compute $\Pr\left(x_t^d|h_t^d, \Phi\right)$, we need the conditional probabilities of V^d given the history until time t; that is, $\Pr\left(V^d = v|h_t^d, \Phi\right)$ for $v \in \{v_L^d, v^{d-1}, v_H^d\}$. These can also be computed recursively by using Bayes's rule.

C. Proofs of Analytical Results

C.1. Proof of Proposition 1

First, observe that since $\nu \geq 1$, a noise trader with shock $N_t^d = 1$ believes the asset value to be v_H^d with probability 1 and a noise trader with shock $N_t^d = 0$ believes the asset value to be v_L^d with probability 1. It follows that if the tax rate ρ is so high that no noise trader trades, no informed trader will trade either; therefore, to prove the first part of the proposition, we will focus on noise traders. At $t = 1$, no noise trader trades if and only if the following two inequalities hold:

$$E(V^d | N_1^d = 1) < a_1^d(1+\rho),$$
$$E(V^d | N_1^d = 0) > b_1^d(1-\rho),$$

that is, a noise trader with the highest signal ($N_1^d = 1$) does not want to buy and a noise trader with the lowest signal ($N_1^d = 0$) does not want to sell given ask and bid prices and the transaction tax. Taking into account that when, in equilibrium, informed traders do not trade, the bid and ask are both equal to the unconditional expected value of the asset, and that $\mathbb{E}(V^d) = v^{d-1}$, noise traders do not trade if and only if

$$v^{d-1}(1+\lambda_H) < v^{d-1}(1+\rho),$$
$$v^{d-1}(1+\lambda_L) > v^{d-1}(1-\rho),$$

which simplifies to

$$\lambda_H < \rho \text{ and } \left(\frac{\delta}{1-\delta}\right)\lambda_H < \rho.$$

These inequalities are jointly satisfied if and only if

$$\rho > \max\left\{1, \frac{\delta}{1-\delta}\right\}\lambda_H = \frac{2\delta}{1-|2\delta-1|}\lambda_H \equiv \bar{\rho}^N.$$

Finally, note that as noise traders do not trade, the market does not open.

We now derive conditions under which informed traders do not buy or sell assuming that noise traders continue to be active. We start with the case in which informed traders do not buy. The informed trader with the highest ex-post valuation is a trader with signal $S_t^d = 1$. The probability they attribute to an increase in the asset value after observing history h_t^d

and this signal is

$$q_h(p) \equiv \Pr\left(V^d = v_H^d | S_t^d = 1, h_t^d\right) = \frac{(1+\tau)p}{(1+\tau)p + (1-\tau)(1-p)},$$

where, to simplify notation, we denote $\Pr\left(V^d = v_H^d | h_t^d\right)$ by p. Informed traders do not buy whenever their highest possible ex-post valuation is below the current ask price including the transaction tax, that is

$$\left[1 + q_h(p)\lambda_H - (1 - q_h(p))\left(\frac{\delta}{1-\delta}\right)\lambda_H\right] v^{d-1} < (1+\rho)\,\text{ask}.$$

If informed traders do not buy, the market maker sets the ask price taking into account that any buy will come from a noise trader. In that case, the market maker's posterior $\Pr\left(V^d = v_H^d | \text{buy}_t^d, h_t^d\right)$ is

$$\frac{(1-\mu)\varepsilon(1-\kappa)a\,p}{(1-\mu)\varepsilon(1-\kappa)\,a + \varepsilon(1-\kappa)(1-a)} = \tilde{a}\,p,$$

where $1-\kappa$ is the probability that a noise trader with a pseudo-signal buys after history h_t^d[41] and $a = \Pr\left(V^d \neq v^{d-1} | h_t^d\right)$. $\tilde{a}$ is the market maker's posterior probability that an event has occurred given that a noise trader is active in t, that is,

$$\tilde{a} = \Pr\left(V^d \neq v^{d-1} | \text{noise trader}, h_t^d\right) = \frac{(1-\mu)a}{1-\mu\,a}.$$

The market maker's valuation of the asset if a noise trader buys after history h_t^d is then given by

$$v^{d-1} + \left[\tilde{a}\,p - \left(\frac{\delta}{1-\delta}\right)\tilde{a}\,(1-p)\right]\lambda_H\,v^{d-1} = v^{d-1} + \tilde{a}\left(\frac{p-\delta}{1-\delta}\right)\lambda_H\,v^{d-1}.$$

It follows that informed traders do not buy whenever their highest possible valuation is below the maker maker's valuation of the asset (plus tax) when only noise traders are active. This

[41] Here, κ designates noise traders' threshold buy pseudo-signal after history h_t^d. Recalling that the pseudo-signal is uniformly distributed on $[0,1]$ irrespective of the asset value, the probability of a noise trader buying is $1-\kappa$ irrespective of the day d event.

condition is met whenever

$$q_h(p) - \delta - (1+\rho)\tilde{a}(p-\delta) < (1-\delta)\frac{\rho}{\lambda_H}. \tag{C.1}$$

Analogous arguments establish the condition under which informed traders do not sell, assuming that noise traders continue to sell. This condition is given by

$$(1-\rho)\tilde{a}(p-\delta) - (q_l(p) - \delta) < (1-\delta)\frac{\rho}{\lambda_H}, \tag{C.2}$$

where $q_l(p)$ is the posterior probability an informed trader with the lowest possible signal, $S_t^d = 0$, assigns to the asset having increased in value, that is

$$q_l(p) = \frac{(1-\tau)p}{(1-\tau)p + (1+\tau)(1-p)}.$$

For $\tau \geq 1$ the proof that no informed trader trades is identical to that for noise traders. Now consider $\tau < 1$. At $t = 1$, we have $p = \Pr(V^d = v_H^d | V^d \neq v^{d-1}) = \delta$. Thus, from (C.1), the condition under which informed traders do not buy at $t = 1$ is given by

$$q_h(p) - \delta < (1-\delta)\frac{\rho}{\lambda_H}.$$

This can be simplified to yield

$$\rho > \left[\frac{2\delta\,\tau}{1-(1-2\delta)\tau}\right]\lambda_H.$$

Likewise, from (C.2) we obtain the condition under which informed traders do not sell at $t = 1$, which is

$$\delta - q_l(p) < (1-\delta)\frac{\rho}{\lambda_H},$$

which is equivalent to

$$\rho > \left[\frac{2\delta\tau}{1+(1-2\delta)\tau}\right]\lambda_H.$$

It follows that informed traders neither buy nor sell whenever

$$\rho > \rho^I \equiv \max\left\{\frac{2\delta\,\min\{\tau,1\}}{1-(1-2\delta)\min\{\tau,1\}}, \frac{2\delta\,\min\{\tau,1\}}{1+(1-2\delta)\min\{\tau,1\}}\right\}\lambda_H$$
$$= \left(\frac{2\delta\,\min\{\tau,1\}}{1-\min\{\tau,1\}|1-2\delta|}\right)\lambda_H.$$

C.2. Proof of Corollary 1

A similar reasoning to that used for the proof of Proposition 1 shows that informed traders are active on both sides of the market for at least some signal realizations if and only if

$$\min\left\{\frac{(1+\tau)\delta}{(1+\tau)\delta+(1-\tau)(1-\delta)}, 1\right\}\lambda_H > \rho,$$
$$\min\left\{\frac{(1-\tau)\delta}{(1-\tau)\delta+(1+\tau)(1-\delta)}, 1\right\}\left(\frac{\delta}{1-\delta}\right)\lambda_H > \rho.$$

The two inequalities are jointly satisfied if and only if

$$\rho < \left(\frac{2\min\{\tau,1\}\,\delta}{1+\min\{\tau,1\}|2\delta-1|}\right)\lambda_H.$$

C.3. Proof of Proposition 2

As a first step to prove Proposition 2, in Lemma 1, we prove that, as the number of trading periods in a day goes to infinity, the market maker learns whether an information event has occurred or not. Intuitively, since $\nu \geq 1$, noise traders always trade (either buy or sell) for some pseudo signal realizations even in the presence of a tax.[42] Therefore, at any time t, the probability of a buy or sell order is different on an event day and a no-event day, which allows the market maker to update their belief α_t^d.

Lemma 1. *Consider a tax rate $\rho < \bar{\rho}$. In equilibrium, the market maker's posterior belief that an event has occurred on day d, $\alpha_t^d =: \Pr(V^d \neq v^{d-1}|h_t^d)$, converges almost surely to $1_{V^d \neq v^{d-1}}$* as $t \to \infty$.

[42] Recall that the assumption $\rho < \overline{\rho}$ means that at time 1 a positive measure of noise traders finds it optimal to buy and a positive measure finds it optimal to sell given their signal realizations. Since $\nu \geq 1$, a positive measure of noise traders does not change their valuation of the asset depending on the history of trades. Therefore, even when the prices change, a positive measure of noise traders continues to buy or to sell.

Proof The process $\{\alpha_t^d\}_{t=1}^{\infty}$ is a bounded martingale under the filtration generated by the successive trades on day d. Hence, by the Martingale Convergence Theorem, α_t^d converges almost surely.

Now consider a given day d and suppose an event has occurred, $V^d \neq v^{d-1}$. First, suppose there is a set of histories with positive measure for which α_t^d converges to a value in $(0,1)$. Then for such trading histories, for any arbitrary value $\eta > 0$, there exists a time T_η such that for all $t > T_\eta$, after observing any given action (buy, sell or no trade), the change in market maker's beliefs about the occurrence of an event is smaller than η. This means that the probabilities of observing any given action (buy, sell, or no trade) conditional on whether an event has occurred or not cannot differ by an amount larger than $\epsilon^S \eta$ where $0 < \epsilon^S < \infty$ can be obtained from Bayes' Rule. Hence, for instance, for a sell order in t it must be that

$$\left|\left(\mu F(\sigma_t^d|V^d) + (1-\mu)\varepsilon\gamma_t^d\right) - \varepsilon\gamma_t^d\right| < \epsilon^S \eta,$$

where σ_t^d is informed traders' threshold sell signal and γ_t^d is noise traders' threshold sell pseudo-signal. We have

$$\left|F(\sigma_t^d|V^d) - \varepsilon\gamma_t^d\right| < \frac{\epsilon^S \eta}{\mu}, \tag{C.3}$$

for both $V^d = v_L^d$ and $V^d = v_H^d$. Similarly, for a buy order, it must be that

$$\left|\left(\mu\left(1 - F(\beta_t^d|V^d)\right) + (1-\mu)\varepsilon\left(1-\kappa_t^d\right)\right) - \varepsilon\left(1-\kappa_t^d\right)\right| < \epsilon^B \eta,$$

that is,

$$\left|\left(1 - F(\beta_t^d|V^d)\right) - \varepsilon(1-\kappa_t^d)\right| < \frac{\epsilon^B \eta}{\mu}, \tag{C.4}$$

for both $V^d = v_L^d$ and $V^d = v_H^d$.

As the probability of selling by a noise trader cannot depend on the asset value, for equality (C.3) to be satisfied, it must be that $\left|F(\sigma_t|V^d = v_L^d) - F(\sigma_t|V^d = v_H^d)\right|$ is arbitrarily close to zero. But this is only possible if informed traders sell for almost all signal realizations or for almost no signal realization.[43] The same considerations apply to the buy order. Since, however, $\varepsilon < 1$, the case in which informed traders sell for almost all signal realizations ($F(\sigma_t^d|V^d) = 1$) cannot satisfy (C.3).[44] Suppose instead that an informed trader does not

[43] For any σ_t in the interior, $\left|F(\sigma_t|V^d = v_L^d) - F(\sigma_t|V^d = v_H^d)\right|$ is bounded away from zero.

[44] Note that, even if $\varepsilon = 1$, the same argument would hold. Indeed, since $\nu \geq 1$, there is always a shock

sell for almost any signal realization ($F(\sigma_t^d|V^d) = 0$); in this case, equality (C.3) could be satisfied if $\gamma_t^d = 0$ (i.e. noise traders do not sell for almost any signal realization). By the same logic, however, for (C.4) to be satisfied as well, noise traders must also not buy for almost any signal realizations ($\kappa_t^d = 1$). If both $\gamma_t^d = 0$ and $\kappa_t^d = 1$, noise traders would never trade (for almost all signal realization). This is, however, impossible, since noise traders find it optimal to buy (or sell) at time 1 for a positive measure of signal realizations ($\rho < \rho^N$) and, therefore, must find it optimal to buy (and to sell) at any time t for at least a positive measure of signal realizations.

To conclude the proof, we must show that α_t^d cannot converge to 0. Suppose it did. Then, for instance, on a good-event day, the probability of a sell order would converge to

$$\mu F(\sigma_t^d|V^d = v_H^d) + (1-\mu)\varepsilon\frac{1}{2},$$

since the market maker would set both bid and ask prices (approximately) equal to the unconditional expected value, v^{d-1}. This probability is different from that of a no-event day, which equals $\varepsilon(1/2)$. The same argument holds for a buy order and a no trade. Hence, as t goes to infinity, the market maker would learn that $\alpha_t^d \neq 0$. The same argument applies to a bad-event day and similar arguments prove the convergence of α_t^d to 0 when $V^d = v^{d-1}$. This concludes the proof.

As an additional step to prove Proposition 2, in the following lemma, we show that when α_t^d is sufficiently high, there exist thresholds for the probability of a good event (conditional on an event having occurred) such that, once they are reached, informed traders do not trade:

Lemma 2. *Consider a tax rate* $\rho < \bar{\rho}$. Let $\delta_t^d =: \Pr(V^d = v_h^d|V^d \neq v^{d-1}, h_t^d)$. *If* $\tau < 1$ *and* $\delta\lambda_H/(1-\delta) < 1$, *in equilibrium, there exist an event probability* $\bar{a} < 1$ *and functions* $0 < \delta_l(\alpha_t^d) < \delta_h(\alpha_t^d) < 1$ *such that for* $\alpha_t^d > \bar{a}$ *informed traders do not trade after history* h_t^d *whenever either* $\delta_t^d < \delta_l(\alpha_t^d)$ *or* $\delta_t^d > \delta_h(\alpha_t^d)$.

Proof Treating the condition under which no informed trader buys as an equality yields a

such that a noise trader's valuation of the asset is v_H^d or v_L^d, that is, the probability of a noise trader buying (or selling) can never be arbitrarily close to 1.

quadratic equation in p,

$$q_h(p) - \delta - (1+\rho)\tilde{a}(p-\delta) = (1-\delta)\frac{\rho}{\lambda_H}, \tag{C.5}$$

where $\tilde{a}$, the market maker's posterior event probability after having observed a noise trader trade, is as defined before

$$\tilde{a}(a) = \frac{(1-\mu)a}{1-\mu a}.$$

If $\rho < \rho^I$, both roots are real and decreasing in the prior event probability a. The original inequality holds whenever p is higher than the larger root, denoted by p_b^1, or lower than the smaller root, denoted by p_b^2. Furthermore, for $\tau < 1$, the larger of the two roots tends to $+\infty$ as a goes to zero and is strictly smaller than 1 when $a = 1$. It follows that there exists a value $0 < a_h < 1$ such that for all $a > a_h$ informed traders stop buying whenever p is higher than $p_b^1(a) < 1$. a_h is given by the solution to $p_b^1(a_h) = 1$, which yields

$$\tilde{a}(a_h) = \left(\frac{1}{1+\rho}\right)\left(\frac{\lambda_H - \rho}{\lambda_H}\right) < 1 \quad \Rightarrow \quad a_h < 1.$$

Similarly, treating the condition under which no informed trader sells as an equality yields a quadratic equation in p,

$$(1-\rho)\tilde{a}(p-\delta) - (q_l(p) - \delta) = (1-\delta)\frac{\rho}{\lambda_H}. \tag{C.6}$$

Both roots are real for $\rho < \rho^I$ and increasing in a. Informed traders do not sell whenever p is higher than the larger root, denoted by p_s^1, or lower than the smaller root, denoted by p_s^2. The smaller of the two roots tends to $-\infty$ as a goes to zero and, if $\lambda_H < (1-\delta)/\delta$, is is strictly above 0 for $a = 1$. Thus, there exists a value a_l such that for all $a > a_l$ no informed trader sells whenever p is lower than $p_s^2(a) > 0$. a_l is given by the solution to $p_s^2(a_l) = 0$, which yields

$$\tilde{a}(a_l) = \left(\frac{1}{1-\rho}\right)\left[1 - \frac{(1-\delta)\rho}{\delta\,\lambda_H}\right] < 1 \quad \Rightarrow \quad a_l < 1.$$

Finally, let

$$\bar{a} = \max\{a_l, a_h\} \;\;,\;\; \delta_h(a) = \max\{p_b^1(a), p_b^2(a), p_s^1(a)p_s^2(a)\}$$

and

$$\delta_l(a) = \min\{p_b^1(a), p_b^2(a), p_s^1(a)p_s^2(a)\}.$$

The above arguments establish that $\bar{a} < 1$ and $0 < \delta_l(a) < \delta_h(a) < 1$ which concludes the proof of Lemma 2.

Intuitively, when after a history of trading the market's belief about the likelihood of a good event is very high or very low, the asymmetry of information between the market maker and informed traders becomes small. In this circumstance, the informational content of a private signal is of little importance with respect to that of the history of trades. There will be a point when the valuations of all informed traders (irrespective of the signal they receive) become so close to the bid and ask prices that the expected gain from trading upon private information becomes smaller than the tax. At this point, all informed traders choose not to trade.

Note that the thresholds for which informed traders stop trading are not constant during the day; they change as a function of the market maker's belief about an information event having occurred.[45] Therefore, it is possible that after informed traders have stopped trading, they resume doing so as the market maker's updates his belief on the likelihood of an information event.

The proof of Proposition 2 follows directly from Lemma 1 and 2. By Lemma 1 we know that on event days α_t^d converges to 1. By Lemma 2, we know that there are two thresholds $\delta_l(1)$ and $\delta_h(1)$ ($0 < \delta_l(1) < \delta_h(1) < 1$) such that informed traders do not trade for $\delta_t^d < \delta_l(1)$ or $\delta_t^d > \delta_h(1)$. Moreover, since δ_t^d is a bounded martingale, it must converge. If it converged to any value in $(\delta_l(1), \delta_h(1))$, informed traders would keep buying or selling for some signal realizations. Since informed traders trade only when their expected gain is greater than the tax paid, the market maker would update δ_t^d by a an amount bounded away from zero, a contradiction. Therefore, δ_t^d converges to a value in $[0, \delta_l(1))$ or $(\delta_h(1), 1]$, such that informed traders stop trading, which proves Proposition 2.

C.4. Proof of Proposition 3

We provide the proof for $V^d = v_H^d$. The proof for $V^d = v_L^d$ is analogous.

Step 1. *The likelihood ratio* $\frac{\Pr\left(V^d = v_L^d | h_t^d\right)}{\Pr\left(V^d = v_H^d | h_t^d\right)}$ *is a martingale conditional on* $V^d = v_H^d$.

[45] The reason is that the market maker's belief on whether the event has occurred affects the market maker's expectation and therefore the bid and ask prices that he sets.

Proof of step 1. Recall that $\delta_t^d =: \Pr\left(V^d = v_H^d | V^d \neq v^{d-1}, h_t^d\right)$. We can then write $\Pr\left(V^d = v_H^d | h_t^d\right) = \alpha_t^d \delta_t^d$ and $\Pr\left(V^d = v_L^d | h_t^d\right) = \alpha_t^d (1 - \delta_t^d)$. We then have:

$$\frac{\Pr\left(V^d = v_L^d | h_t^d\right)}{\Pr\left(V^d = v_H^d | h_t^d\right)} = \frac{1 - \delta_t^d}{\delta_t^d}.$$

We have to show that

$$\mathbb{E}\left(\frac{1 - \delta_{t+1}^d}{\delta_{t+1}^d} \mid h_t^d, V^d = v_H^d\right) = \frac{1 - \delta_t^d}{\delta_t^d}.$$

By Bayes's rule,

$$\mathbb{E}\left(\frac{1 - \delta_{t+1}^d}{\delta_{t+1}^d} \mid h_t^d, v_H^d\right) = \mathbb{E}\left(\frac{\Pr(X_{t+1}^d | h_t^d, v_L^d)(1 - \delta_t^d)}{\Pr(X_{t+1}^d | h_t^d, v_H^d)\delta_t^d} \mid v_H^d\right)$$
$$= \left(\frac{1 - \delta_t^d}{\delta_t^d}\right) \sum_{x \in \{sell, nt, buy\}} \Pr(X_{t+1}^d = x | h_t^d, v_H^d) \left[\frac{\Pr(X_{t+1}^d = x | h_t^d, v_L^d)}{\Pr(X_{t+1}^d = x | h_t^d, v_H^d)}\right] = \frac{1 - \delta_t^d}{\delta_t^d}.$$

Step 2. *The probability that the belief* $\Pr(V^d = v_H^d | h_t^d)$ *remains stuck at a value lower than* $\delta_l(1)$ *on an good event day (*$V^d = v_H^d$*) is bounded above by*

$$\frac{\delta_l(1)[\delta_h(1) - \delta]}{\delta[\delta_h(1) - \delta_l(1)]}.$$

Proof of step 2. Consider the time T at which informed traders stop trading (as defined in Proposition 3) and assume that by that time the market has learned that an event has occurred, i.e. $\alpha_T^d = 1$. Since the likelihood ratio $\frac{\Pr\left(V^d = v_L^d | h_t^d\right)}{\Pr\left(V^d = v_H^d | h_t^d\right)}$ is a martingale conditional on $V^d = v_H^d$, we have that

$$\mathbb{E}\left(\frac{1 - \delta_T^d}{\delta_T^d} \mid V^d = v_H^d\right) = \frac{1 - \delta}{\delta}.$$

As, by Lemma 2, the posterior beliefs δ_t^d have to converge to either $\delta_h(1)$ or $\delta_l(1)$ we have

$$\mathbb{E}\left(\frac{1 - \delta_T^d}{\delta_T^d} \mid v_H^d\right) = \Pr(\delta_T^d = \delta_h(1) | v_H^d)\left(\frac{1 - \delta_h(1)}{\delta_h(1)}\right) + (1 - \Pr(\delta_T^d = \delta_h(1) | v_H^d))\left(\frac{1 - \delta_l(1)}{\delta_l(1)}\right),$$

from which we obtain that

$$\Pr(\delta_T^d = \delta_h(1) | v_H^d) = \frac{\delta_h(1)[\delta - \delta_l(1)]}{\delta[\delta_h(1) - \delta_l(1)]},$$

and

$$\Pr(\delta_T^d = \delta_l(1)|v_H^d) = \frac{\delta_l(1)[\delta_h(1) - \delta]}{\delta[\delta_h(1) - \delta_l(1)]}.$$

Finally, note that since for $\alpha_t^d < 1$, $\delta_h(\alpha_t^d)$ can be higher than $\delta_h(1)$ and $\delta_l(\alpha_t^d)$ can be lower than $\delta_l(1)$, informed traders may stop trading before α_t^d has converged to 1. Hence $\Pr(\delta_T^d = \delta_l(1)|v_H^d)$ is only bounded from above by the expression derived assuming $\alpha_T^d = 1$. Defining $\delta_l \equiv \delta_l(1)$ and $\delta_h \equiv \delta_h(1)$ we obtain the expressions given in Proposition 3 of Section 3.1.

To understand why the formulas in Proposition 3 are only probability bounds, it is useful to focus first on the case of no event uncertainty (that is, $\alpha_t^d = \alpha = 1$ for any t). In this case, the market maker stops updating $\Pr(V^d = v_H^d|h_t^d)$ as soon as the thresholds are reached. Moreover, in this case, the probabilities that $\Pr(V^d = v_H^d|h_t^d)$ converges to the low threshold when $V^d = v_H^d$ or to the high threshold when $V^d = v_L^d$ are equal to $\frac{(1-\delta_h(1))(\delta-\delta_l(1))}{(1-\delta)(\delta_h(1)-\delta_l(1))}$ and $\frac{\delta_l(1)(\delta_h(1)-\delta)}{\delta(\delta_h(1)-\delta_l(1))}$ respectively. In contrast, when there is event uncertainty, the market maker, in principle, may stop learning about whether the event is good or bad before the belief on whether an information event has occurred has converged to 1. This may occur when $\Pr(V^d = v_H^d|h_t^d)$ is above $\delta_h(1)$ or below $\delta_l(1)$. Since the levels at which the belief may be stuck are different, so are the probabilities of these events. That is the reason the probabilities that the price is misaligned with the fundamental value, indicated in Proposition 3, are only bounds.

D. Private Values for Noise Traders

D.1. Equivalence of pseudo signals and multiplicative private values.

In this section, we show that: (i) at each time t, the private valuations generated by the pseudo signals are independent of the fundamental asset value V^d conditional on public information (i.e., on the history of trades up to t), (ii) they can be expressed through a multiplicative shock with a time varying distribution. Specifically, we prove the following propositions:

Proposition. *For any history of trades h_t^d, the distribution of a noise trader's asset valuation v_t^d is independent of the fundamental value of the asset V^d, after conditioning on the history of trades h_t^d.*

Recall that in our model the asset's fundamental value on day d, V^d, is drawn according to

$$V^d = \begin{cases} v_H^d = \lambda_H \, v^{d-1} & \text{with probability } \delta\alpha \\ v^{d-1} & \text{with probability } 1-\alpha \\ v_L^d = \lambda_L \, v^{d-1} & \text{with probability } (1-\delta)\alpha. \end{cases}$$

v^{d-1}, the fundamental on the previous day, is commonly known at the beginning of trading on day d. Denote by δ_t^d the probability that the fundamental value is v_H^d given that an event has occurred and given history h_t^d and by α_t^d the probability that an event has occurred given history h_t^d. As market maker, informed traders, and noise traders start with common priors on day d and observe the same history of trading up to t, these beliefs are commonly known.

Now suppose a noise trader with pseudo signal n_t^d is trading at time t. This noise trader updates their valuation of the asset **as if** it were generated from density functions given in (2), $\tilde{g}^H(n|V^d = v_H^d)$, $\tilde{g}^L(n|V^d = v_L^d)$ and $\tilde{g}(n|V^d = v^{d-1})$ that depend on the realization of V^d. Actually, however, the pseudo signal n_t^d is drawn from the uniform distribution on $[0, 1]$ irrespective the value of V^d. That is, the distribution of the pseudo signal n_t^d does not depend on the realization of V^d.

The noise trader uses Bayes' Rule to calculate an expected value of V^d conditional on n_t^d treating the pseudo signal **as if** it contained information. This pseudo-Bayesian updating yields a valuation of

$$\tilde{\mathbb{E}}\left(V^d | n_t^d, h_t^d\right) = v_{h_t^d}(n_t^d) = [1 - \tilde{p}_{h_t^d}(n_t^d) - \tilde{q}_{h_t^d}(n_t^d)]v^{d-1} + \tilde{q}_{h_t^d}(n_t^d)v_L^d + \tilde{p}_{h_t^d}(n_t^d)v_H^d,$$

where $\tilde{p}_{h_t^d}(n_t^d) = \tilde{\Pr}(V^d = v_H^d | n_t^d, h_t^d)$ and $\tilde{q}_{h_t^d}(n_t^d) = \tilde{\Pr}(V^d = v_L^d | n_t^d, h_t^d)$ are the pseudo-posterior beliefs of the noise trader for the asset value being high, respectively low, given pseudo signal n_t^d and are given by

$$\tilde{p}_{h_t^d}(n) = \begin{cases} 0 & \text{if } n \leq (\nu-1)/\nu \\ \frac{\tilde{g}^H(n)\alpha_t^d\delta_t^d}{\tilde{g}^H(n)\alpha_t^d\delta_t^d + \tilde{g}^L(n)\alpha_t^d(1-\delta_t^d) + \nu/(2-\nu)(1-\alpha_t^d)} & \text{if } (\nu-1)/\nu < n < 1/\nu \\ 1 & \text{if } n \geq 1/\nu \end{cases}$$

and

$$\tilde{q}_{h_t^d}(n) = \begin{cases} 1 & \text{if } n \leq (\nu-1)/\nu \\ \frac{\tilde{g}^L(n)\alpha_t^d(1-\delta_t^d)}{\tilde{g}^H(n)\alpha_t^d\delta_t^d + \tilde{g}^L(n)\alpha_t^d(1-\delta_t^d) + \nu/(2-\nu)(1-\alpha_t^d)} & \text{if } (\nu-1)/\nu < n < 1/\nu \\ 0 & \text{if } n \geq 1/\nu. \end{cases}$$

The noise trader's valuation $v_t^d \equiv \tilde{\mathbb{E}}\left(V^d | n_t^d, h_t^d\right)$ does not contain any information that is not already contained in the commonly known history of trades h_t^d. The pseudo signal n_t^d itself is drawn from the uniform distribution $U[0,1]$ irrespective of V^d and the mapping from n_t^d to v_t^d only depends on the functions $\tilde{g}^H, \tilde{g}^L, \tilde{g}$ and the beliefs δ_t^d and α_t^d that are commonly known after history h_t^d. Hence, conditional on h_t^d, the distribution of v_t^d does not depend on the realization of V^d which concludes the proof.

As the pseudo signal, n_t^d, is uniformly distributed on $[0,1]$ and, given h_t^d, there is a monotone one-to-one mapping from n_t^d to the noise trader's valuation, v_t^d, it is straightforward to characterize the CDF of v_t^d given h_t^d. Denote by $n_{h_t^d}$ the inverse mapping from v_t^d to n_t^d, that is $n_t^d = n_{h_t^d}(v_t^d)$ whenever $v_t^d = v_{h_t^d}(n_t^d)$ for all $(\nu-1)/\nu \leq n_t^d \leq 1/\nu$. Then the CDF of a noise trader's valuation v_t^d given public history h_t^d is

$$F_{h_t^d}(v) = \begin{cases} 0 & \text{for } v < v_L^d, \\ n_{h_t^d}(v) & \text{for } v_L^d \leq v < v_H^d, \\ 1 & \text{for } v \geq v_L^d. \end{cases}$$

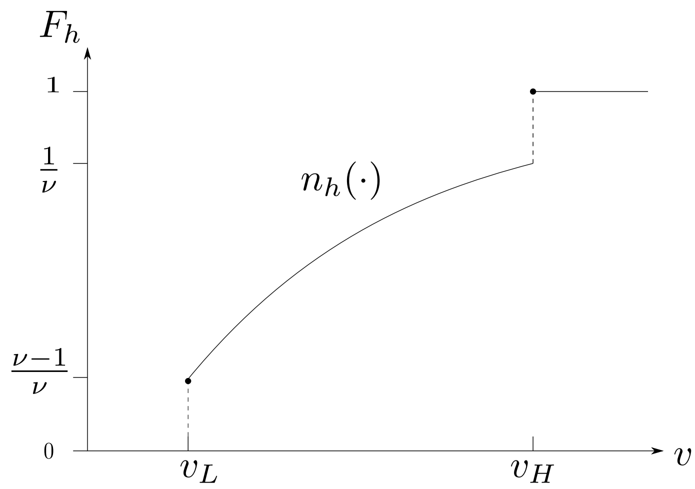

Figure D.1: CDF of noise trader's valuation v given history h.

Proposition. *For any history of trades h_t^d and realization of the pseudo signal n_t^d, a noise trader's valuation of the asset can through a multiplicative private value shock ρ_t^d with history-dependent distribution function $G_{h_t^d}$ such that*

$$\tilde{\mathbb{E}}\left(V^d|h_t^d, n_t^d\right) = \rho_t^d \, \mathbb{E}(V^d|h_t^d),$$

where $\tilde{\mathbb{E}}$ is the expectation of a noise trader obtained via pseudo-Bayesian updating given a realization n_t^d of the pseudo signal.

Recall the notation $v_t^d = \tilde{\mathbb{E}}\left(V^d|h_t^d, n_t^d\right)$ and $p_t^d = \mathbb{E}(V^d|h_t^d)$. As for all h_t^d, $p_t^d > 0$ we can express the private value shock as

$$\rho_t^d = \frac{v_t^d}{p_t^d}.$$

Given h_t^d, p_t^d is fixed and we have already characterized the distribution of v_t^d given h_t^d, namely $F_{h_t^d}$. If follows that the CDF of the multiplicative private value shock ρ_t^d given history h_t^d is

$$G_{h_t^d}(\rho) = \begin{cases} 0 & \text{for } \rho < v_L^d/p_t^d, \\ n_{h_t^d}(\rho \, p_t^d) & \text{for } v_L^d/p_t^d \leq \rho < v_H^d/p_t^d, \\ 1 & \text{for } \rho \geq v_L^d/p_t^d. \end{cases}$$

D.2. Microfoundation of Pseudo Signals

In this section we provide a microfoundation of the noise traders' private values given by the pseudo signals; in particular, we prove that hedging motives generate noise traders' valuations with the same distribution as those generated by the pseudo-signals.

Let us denote the private value for a noise trader receiving a pseudo signal n_t^d by $B^d \equiv E(V^d|h_t^d, n_t^d)$ and the distribution of these private values by $f_{B^d}(B^d)$.

Proposition 1 *Consider noise traders with preferences represented by a HARA utility function over wealth w with risk aversion parameter r distributed on $[0,1]$: $u(w,r) = w(r-2)/(2w(r-1)-r)$. Assume they have an endowment of the asset of either V^d or $1-V^d$ with equal probability. There exists a distribution of r such that the distribution of private values from the asset for these noise traders is equal to $f_{B^d}(B^d)$.*

Proposition 2 *Consider noise traders with preferences represented by a CRRA utility function over wealth w with risk aversion parameter equal to 2, that is, $u(w) = -(1/w)$. Assume they have an endowment of M units of cash and either V^d or $1-V^d$ of the asset with equal probability . There exists a distribution of M on the interval $[v^L, \infty]$ such that the distribution of private values from the asset for these noise traders is equal to $f_{B^d}(B^d)$.*

Proof of Propositions

First we establish some results useful to prove both propositions.

Preliminary result

Let us assume that at time t, $\Pr(V^d = v^H|h_t^d, n_t^d) = p$ and $\Pr(V^d = v^L|h_t^d, n_t^d) = q$. Let us consider the transformation

$$V = (1-\delta)\left(\frac{V^d - v^{d-1}}{\lambda_H v^{d-1}}\right),$$

and note that V is distributed on $\{-\delta, 0, 1-\delta\}$. Its expectation is $(1-\delta)p - \delta q$. The noise trader's private value in our model is defined as

$$B^d = \mathbb{E}(V^d|h_t^d, n_t^d) = \sum_i V_i^d \Pr(V_i^d|h_t^d, n_t^d).$$

Let us define

$$B = \mathbb{E}(V|h_t^d, n_t^d) = \sum_i V_i \ \Pr(V_i|h_t^d, n_t^d).$$

Note that $B = ((1-\delta)/(\lambda_H v^{d-1}))B^d - v^{d-1}$. Since B^d is a one-to-one, differentiable, function $g(B)$ of B, the distribution f_{B^d} of B^d is given by

$$f_{B^d}(B^d) = f_B(g^{-1}(B^d)) \left| \frac{dB}{dB^d} \right|.$$

We, therefore, prove our result for the random variable V.

Pseudo signals

First, we compute a noise trader's private value, given their beliefs p and q and their pseudo signal n:

$$\Pr(V = 1-\delta|n) = (((1+\nu(2n-1))p)/((1+\nu(2n-1))p + (1-p-q) + (1-\nu(2n-1))q))$$
$$\Pr(V = 0|n) = (((1-p-q))/((1+\nu(2n-1))p + (1-p-q) + (1-\nu(2n-1))q))$$
$$\Pr(V = -\delta|n) = (((1-\nu(2n-1))q)/((1+\nu(2n-1))p + (1-p-q) + (1-\nu(2n-1))q))$$

Therefore,

$$B = (((1+\nu(2n-1))p)/((1+\nu(2n-1))p + (1-p-q) + (1-\nu(2n-1))q))(1-\delta) -$$
$$(((1-\nu(2n-1))q)/((1+\nu(2n-1))p + (1-p-q) + (1-\nu(2n-1))q))\delta =$$
$$((p - p\delta - p\nu - q\delta + 2pn\nu + p\delta\nu - q\delta\nu - 2pn\delta\nu + 2qn\delta\nu)/(q\nu - p\nu + 2pn\nu - 2qn\nu + 1).$$

B is a strictly increasing function of n. Its inverse is

$$n(B) = (((B - p + p\delta + p\nu + q\delta - Bp\nu + Bq\nu - p\delta\nu + q\delta\nu))/(2p\nu - 2Bp\nu + 2Bq\nu - 2p\delta\nu + 2q\delta\nu)).$$

and its derivative with respect to B is

$$\frac{dn}{dB} = ((p(1-\delta)(1-p) + q(p + \delta(1-q)))/(2\nu(p - p\delta + q\delta - Bp + Bq)^2))$$

Finally, the distribution of B is

$$f_B(B) = f_n(n(B))\left|\frac{dn}{dB}\right| = \left|\frac{dn}{dB}\right|$$
$$= ((p(1-\delta)(1-p) + q(p+\delta(1-q)))/(2\nu(p - p\delta + q\delta - Bp + Bq)^2))$$

on the support $[-\delta, 1-\delta]$, since pseudo signals are distributed uniformly on $[0,1]$.

Proof of Proposition 1

Consider a noise trader who has an endowment of 1 unit of the asset. By selling the asset they can completely hedge their risk. At time t, the noise trader has expected utility

$$Eu(w,r) = u(1-\delta, r)p + u(-\delta, r)q =$$
$$(((1-\delta)(r-2))/(2(1-\delta)(r-1)-r))p + (((-\delta)(r-2))/(2(-\delta)(r-1)-r))q.$$

If the noise traders sell the asset at a bid price B, their expected utility is

$$Eu(w,r) = u(B,r) = ((B(r-2))/(2B(r-1)-r)).$$

To obtain an expression for the noise trader's private value, we find the lowest price at which he is willing to sell:

$$(((1-\delta)(r-2))/(2(1-\delta)(r-1)-r))p + (((-\delta)(r-2))/(2(-\delta)(r-1)-r))q$$
$$= ((B(r-2))/(2B(r-1)-r)).$$

This implies that

$$B = ((pr^2 - 2pr^2\delta^2 - 2qr^2\delta^2 - 2pr\delta - 2qr\delta + 2pr\delta^2 + pr^2\delta + 2qr\delta^2 + qr^2$$
$$\delta)/(2r - 4\delta + 4r^2\delta^2 + 4p\delta + 4q\delta + 4r\delta + 4\delta^2 + 2pr^2 - 4p\delta^2 - 4q\delta^2$$
$$- 8r\delta^2 - 2pr - r^2 - 4pr^2\delta^2 - 4qr^2\delta^2 - 6pr\delta - 6qr\delta + 8pr\delta^2 + 2pr^2\delta + 8qr\delta^2 + 2qr^2\delta)).$$

Let us denote this function by $B = h(r)$. It is equal to 0 for $r = 0$ (extreme risk aversion) and to the asset's expected value, $(1-\delta)p - \delta q$, for $r = 1$ (risk neutrality). Therefore, B is strictly increasing on the support $[-\delta, (1-\delta)p - \delta q]$ with density

$$f_B(B) = f_r(h^{-1}(B)) \left| \frac{dr}{dB} \right|.$$

This distribution is equal to

$$f_B(B) = \left| \frac{dn}{dB} \right|,$$

when the distribution of r is

$$f_r(h^{-1}(B)) = \left| \frac{ds}{dB} \right| / \left| \frac{dr}{dB} \right|$$

on $[0, 1]$. Analogous steps for a noise trader with endowment of $1 - V$ (i.e., who is short of one unit of the asset) prove the existence of an $f_r(h^{-1}(B))$ on $[0, 1]$ that has the same distribution $f_B(B)$ on the support $[(1-\delta)p - \delta q, 1 - \delta]$.

Proof of Proposition 2

Let us study a noise trader's asset valuation when they have has an endowment of M of cash (or certainty equivalent of a portfolio uncorrelated with the asset) and an endowment of 1 unit of the asset. Without hedging, their utility is

$$\mathbb{E}\, u(M) = -\left(\frac{1}{M+1+\delta} \right) p - \left(\frac{1}{M} \right) (1-p-q) - \left(\frac{1}{M-\delta} \right) q.$$

If the noise trader hedges by selling the asset at a bid price B, their expected utility is

$$\mathbb{E}\, u(M) = -\left(\frac{1}{M+B} \right).$$

They are indifferent between hedging and not when

$$\left(\frac{1}{M+1+\delta} \right) p + \left(\frac{1}{M} \right) (1-p-q) + \left(\frac{1}{M-\delta} \right) q = \frac{1}{M+B},$$

that is, when

$$B = -((-M^2p + Mp\delta + Mq\delta + Mp\delta^2 - M^2p\delta + Mq\delta^2 + M^2q\delta)/(M - \delta + p\delta + q\delta - \delta^2 + p\delta^2 + q\delta^2 - Mp + M^2 - Mp\delta + Mq\delta)).$$

Standard arguments show that B is a strictly increasing function $B(M)$ of M. Moreover, for $M = \delta$, $B = -\delta$ and for M that goes to infinity, B converges to $(1-\delta)p - \delta q$ (the asset's expected value). Therefore, B is strictly increasing on the support $[-\delta, (1-\delta)p - \delta q]$ with density

$$f_B(B) = f_M(k^{-1}(B)) \left| \frac{dM}{dB} \right| .$$

This distribution is equal to

$$f_B(B) = \left| \frac{dn}{dB} \right| ,$$

when the distribution of M is

$$f_M(k^{-1}(B)) = \left| \frac{dn}{dB} \right| / \left| \frac{dM}{dB} \right|$$

on $[\delta, \infty]$.

Analogous steps for a noise trader with endowment of $M + 1 - V$ (i.e., they are short on the asset by one unit) prove the existence of an $f_M(k^{-1}(B))$ on $[\delta, \infty]$ that has the same distribution $f_B(B)$ on the support $[(1-\delta)p - \delta q, 1 - \delta]$.

D.3. Standard Deviation of Private Values

	stddev	range
Q1	119	442
Q2	69	258
Q3	39	145
Q4	40	136
$\tau < 1$	274	953
all stocks	62	228

Table D.1: Standard deviation and range of the percentage difference between noise traders' private asset valuations and the price, $\tilde{\mathbb{E}}(V^d|h_t^d, n_t^d) - p_t^d$, in basis points. Simulations are computed at the median parameters across all stocks, for each quartile, and for those stocks with $\tau < 1$. For comparison, in Hendershott and Menkveld (2014, p.417), the ranges for private values by NYSE market capitalization quintile are 614, 364, 292, 266, 348 and 370 bps (median from smallest to largest and across all stocks); the corresponding standard deviations are 177, 105, 84, 77, 100 and 107. The unconditional standard deviation of private values in Hollifield et al. (2006, p.2784) is 21% (i.e., 2100 bps), which is the standard deviation of the mixture of two standard normal distributions.