

The Role of Transaction Costs for Financial Volatility

Evidence from the Paris Bourse

Harald Hau (2006)

Journal of the European Economic Association

1. Motivation: Keynes vs. Friedman debate
2. Problems in prior empirical literature
3. Natural experiment: Paris Bourse tick size discontinuity
4. Measuring transaction costs and volatility
5. **Mathematical derivation: std. deviation bias**
6. Key results
7. Policy implications

Motivation

The Tobin Tax Debate

Keynes (1936) / Tobin (1978):

- Transaction costs curb short-term speculation
- Higher costs $\Rightarrow$ fewer noise traders $\Rightarrow$ **lower volatility**
- “Throw sand in the wheels of finance”

Friedman (1953):

- Transaction costs discourage stabilizing arbitrage
- Higher costs $\Rightarrow$ less efficient prices $\Rightarrow$ **higher volatility**
- Costs reduce market liquidity and informational efficiency

This paper: provides *causal* evidence using a natural experiment.

Why Prior Evidence Is Inconclusive

- **Cross-country studies:** different market microstructure, regulation, investor composition $\Rightarrow$ omitted variable bias
- **Event studies** (e.g. Swedish tax 1984): tax changes coincide with other reforms $\Rightarrow$ confounding
- **Endogeneity:** volatile stocks naturally have wider spreads
- **Core problem:** cannot separate cause from effect

$\Rightarrow$ Need **exogenous** variation in transaction costs.

Natural Experiment

The Paris Bourse Tick Size Rule

The Paris Bourse imposed a **discontinuous minimum tick size**:

	Price < FF 500	Price $\geq$ FF 500
Minimum tick	FF 0.1	FF 1.0
Relative tick (at FF 500)	0.02%	0.20%

- At the FF 500 threshold: **10 $\times$ jump** in relative tick size
- Stocks just above vs. just below the cutoff are otherwise similar
- Creates **quasi-random assignment** of transaction costs

Measuring Transaction Costs

Effective Spread

The **effective spread** captures the actual cost of a round-trip trade:

$$\text{ESPREAD}_t = 200 \times \frac{|P_t^T - P_t^M|}{P_t^M}$$

where:

- P_t^T = transaction price at t
- P_t^M = midpoint of bid-ask quote at t
- Factor 200 = doubling (half-spread $\rightarrow$ full) $\times$ 100 (percentage)

For regressions, use log transform: $\text{LESPREAD}_t = \ln(\text{ESPREAD}_t)$

Table 2: Spread Distribution

TABLE 2. Log effective spreads for trades and executed orders by tick regime.

Panel A: Effective spreads for trades and subtrades						
Percentiles	Small ticks			Large ticks		
	Centiles	1% C.I.		Centiles	1% C.I.	
1%	-3.8669	-3.8669	-3.8669	-1.7893	-1.7893	-1.7893
5%	-3.7567	-3.7567	-3.7564	-1.7758	-1.7758	-1.7758
10%	-3.1551	-3.1555	-3.1547	-1.7639	-1.7639	-1.7639
25%	-2.4753	-2.4761	-2.4744	-1.7272	-1.7272	-1.7272
50%	-1.8488	-1.8499	-1.8483	-1.6668	-1.6668	-1.6668
75%	-1.3005	-1.3009	-1.2997	-1.0219	-1.0225	-1.0207
90%	-0.7602	-0.7608	-0.7595	-0.5869	-0.5869	-0.5869
95%	-0.4340	-0.4350	-0.4340	-0.2983	-0.3001	-0.2983
99%	0.1522	0.1520	0.1534	0.2886	0.2886	0.2894
Observations	2,540,764			2,155,658		
Mean	-1.923			-1.371		
Std. dev.	0.944			0.528		
Skewdness	-0.157			1.465		
Kurtosis	2.791			4.702		

Key observations:

- Median LSPREAD: -1.849 (small) vs -1.667 (large tick)
- $\Rightarrow$ median effective spread $\sim 20\%$ higher for large-tick stocks: $e^{0.182} \approx 1.20$
- Large ticks: skewness 1.47 (strongly right-skewed) vs -0.16 for small ticks
- Censoring effect: tick constraint truncates left tail

Figure 2: Spread vs. Price Level




FIGURE 2. The log effective spread is plotted for a random sample of 20,000 trades on stocks in the price range from FF 400 to FF 600. At FF 500, the minimal tick size increases from FF 0.1 to FF 1. A small amount of noise is added to each observation to render it visually distinguishable.

What to see:

- Each dot = one trade (random sample of 20,000)
- Below FF 500: many bands at 2, 4, 6... bp
- Above FF 500: bands below 20 bp **vanish** — tick constraint binds
- Sharp discontinuity right at FF 500, not gradual

Figures 3–4: Spread Distributions




FIGURE 3. The kernel density estimate of the log effective spread is presented for 2,540,764 trades in the small tick regime and 2,155,658 trades in the large tick regime.

Key features:

- Small ticks (solid): broad distribution with peaks at 2, 4, 6... bp
- Large ticks (dashed): huge **censoring spike** at 20 bp (= minimum possible spread)
- Near-zero density below 17 bp for large ticks
- $\Rightarrow$ tick constraint is **binding** — forces wider spreads than market would set

Volatility Measurement

Standard Deviation Bias — Setup

Observed price = true price + rounding error: $P'_t = P_t^*(1 + e_t)$, where $e_t \sim U[-\frac{1}{2}\text{PTS}, \frac{1}{2}\text{PTS}]$.

True return: $R_t = (P_{t+1}^* - P_t^*)/P_t^*$. Observed return:

$$R'_t = \frac{P_{t+1}^*(1 + e_{t+1}) - P_t^*(1 + e_t)}{P_t^*(1 + e_t)}$$

Since $P_{t+1}^*/P_t^* = 1 + R_t$:

$$R'_t = \frac{(1 + R_t)(1 + e_{t+1}) - (1 + e_t)}{1 + e_t}$$

Expanding: $(1 + R_t)(1 + e_{t+1}) - (1 + e_t) = R_t + e_{t+1} - e_t + R_te_{t+1}$.

Since e_t is small and $R_t \cdot e_{t+1} \approx 0$, and $1/(1 + e_t) \approx 1$:

$$R'_t \approx R_t + e_{t+1} - e_t$$

Standard Deviation Bias — Variance

Taking variance of both sides, with e_t independent of R_t and e_{t+1} :

$$\text{Var}(R'_t) = \text{Var}(R_t) + \text{Var}(e_{t+1}) + \text{Var}(e_t) = \text{Var}(R_t) + 2\sigma_e^2$$

Since $e_t \sim U[-\frac{1}{2}\text{PTS}, \frac{1}{2}\text{PTS}]$:

$$\sigma_e^2 = \frac{\text{PTS}^2}{12} \quad \Rightarrow \quad \boxed{\text{Bias} = 2 \times \frac{\text{PTS}^2}{12} = \frac{\text{PTS}^2}{6}}$$

$\Rightarrow$ Standard deviation is **severely biased** for large-tick stocks $\Rightarrow$ need a robust alternative.

Measurement Bias — Standard Deviation

From the variance bias $\text{Var}(\tilde{R}_t) \approx \text{Var}(R_t) + \frac{1}{6}\text{PTS}^2$, define:

$$\text{Measurement Bias} = \frac{SD(\tilde{R}_t) - SD(R_t)}{SD(R_t)}$$

Substituting:

$$= -1 + \sqrt{\frac{\text{Var}(\tilde{R}_t)}{\text{Var}(R_t)}} \approx -1 + \sqrt{1 + \frac{\frac{1}{6}(\text{PTS})^2}{\text{Var}(R_t)}}$$


The bias is always **positive** and increasing in PTS, decreasing in true volatility.

At median daily $\sigma = 1.2\%$:

- Small tick (PTS = 0.02%): $-1 + \sqrt{1 + \frac{0.0002^2/6}{0.012^2}} \approx 2\%$
- Large tick (PTS = 0.20%): $-1 + \sqrt{1 + \frac{0.002^2/6}{0.012^2}} \approx 23\%$

⇒ Standard deviation is **unusable** as a volatility measure for large-tick stocks.

Figure 1: Monte Carlo Confirmation



Monte Carlo setup:

- 1 million random draws from uniform initial distribution
- Normally distributed return innovations
- Compared true vs. measured volatility under small tick (PTS = 0.1/500) and large tick (PTS = 1/500)

Results (at median hourly $\sigma = 1.2\%$):

- Small tick: std. dev. bias $\sim 2\%$
- Large tick: std. dev. bias $\sim 23\%$

⇒ confirms analytical derivation

The Range Estimator — Solution

Hau uses the **intraday range** as a volatility measure:

$$\text{RANGE}_t = 200 \times \frac{\max_{\{s \in I_t\}} P_s^M - \min_{\{s \in I_t\}} P_s^M}{\max_{\{s \in I_t\}} P_s^M + \min_{\{s \in I_t\}} P_s^M}$$

where I_t is the set of all transactions within hourly trading period t , and P_s^M is the midquote price at transaction s .

Why is this better?

- At median volatility: range bias $< 5\%$ vs. std. dev. bias up to 23%
- Can be further corrected by controlling for PTS in regressions

Log transformation: $\text{LRANGE}_t = \ln(0.1 + \text{RANGE}_t)$

Results

Table 3: Per-Stock Volatility

TABLE 3. Summary statistics on hourly log percentage range by stock.

Stock Name	Small ticks					Large ticks				
	Obs.	Mean	St. dev.	Skew.	Kurt.	Obs.	Mean	St. dev.	Skew.	Kurt.
Sanofi	609	-0.679	0.627	0.245	2.945	1,460	-0.226	0.601	-0.091	3.134
Total	1,033	-0.617	0.545	0.072	2.598	582	0.012	0.513	-0.126	2.926
Elf Aquitaine	646	-0.765	0.518	0.135	2.703	630	-0.170	0.526	-0.068	2.895
Bouygues	753	-0.756	0.649	0.028	2.805	3,785	-0.582	0.659	-0.330	3.403
Lyonnais des Eaux	2,886	-0.761	0.595	0.150	2.812	1,441	-0.540	0.618	-0.243	3.568
Lafarge	945	-0.263	0.613	0.083	2.555	627	0.010	0.547	-0.258	3.385
AXA	615	-0.569	0.576	0.812	4.580	269	0.046	0.565	-0.230	3.687
CFF	52	-0.771	0.683	0.277	2.442	336	-0.437	0.749	-0.199	3.208
BIC	1,230	-0.247	0.634	0.086	2.733	96	-0.021	0.601	0.345	2.573
Banque (CIE)	629	-0.588	0.655	-0.145	2.649	2,661	-0.521	0.700	-0.338	3.192
Général des Eaux	1,091	-0.545	0.614	0.113	2.895	2,016	-0.571	0.570	-0.180	3.385
Spie Batignolles	404	-0.281	0.606	0.364	2.860	58	-0.035	0.616	-0.055	3.071
Paribas	1,165	-0.304	0.603	0.136	2.777	461	-0.125	0.552	-0.185	4.262
CCF	869	-0.262	0.667	0.151	2.729	288	-0.098	0.590	-0.121	3.277
Alcatel Alstom	3,201	-0.626	0.603	0.335	3.143	486	-0.197	0.741	0.474	3.594
Havas	2,291	-0.630	0.628	0.090	2.813	275	-0.232	0.558	-0.253	4.488
Viéo	741	-0.071	0.671	0.023	2.959	457	0.043	0.546	0.048	3.051
Société Générale	110	-0.392	0.521	-0.093	3.107	3,255	-0.583	0.594	0.051	3.777
Crédit LCL France	2,758	-0.816	0.629	0.168	2.852	1,256	-0.366	0.607	-0.267	3.556
BNP	396	-0.039	0.602	0.028	2.576	363	-0.029	0.568	0.119	3.989
SGS Thomson	648	-0.130	0.668	-0.129	2.580	115	-0.071	0.595	-0.832	5.320
France Télécom	409	0.191	0.562	-0.134	3.020	—	—	—	—	—
Accor	—	—	—	—	—	799	-0.611	0.607	-0.236	3.523
Peugeot	—	—	—	—	—	1,181	-0.429	0.654	-0.196	3.180
Saint-Gobain	—	—	—	—	—	711	-0.432	0.535	-0.210	3.824
Canal Plus	—	—	—	—	—	124	-0.422	0.617	-0.199	2.679
All Stocks	23,481	-0.542	0.661	0.186	2.881	23,732	-0.423	0.651	-0.190	3.456

Notes: Distribution statistics of the hourly log percentage range (LRANGE) are calculated for all stocks in the French CAC40 index which trade for at least 20 hours in the price range (in French francs, FF) from FF 400 to FF 600 over the four-year period from January 1995 to December 1998. The log percentage range over hourly trading periods t , from 10:00 a.m. to 5:00 p.m. is defined as

$$LRANGE_t = \ln \left[0.1 + 200 \times \frac{\max_{(p \in I)} P_p^M - \min_{(p \in I)} P_p^M}{\max_{(p \in I)} P_p^M + \min_{(p \in I)} P_p^M} \right],$$

Key observations:

- **19 out of 21** stocks trading in both regimes show higher LRRANGE under large ticks
- All stocks: mean LRRANGE -0.542 (small) vs -0.423 (large tick)
- Example — Sanofi: mean LRRANGE jumps from -0.679 to -0.226 (+70% range increase)
- Pattern is nearly universal, not driven by outliers

Volatility: Percentile Comparison

The LRANGE distribution for large-tick stocks **stochastically dominates** the small-tick distribution:

Percentile	Small tick	Large tick	Difference
25th	0.94	1.17	+24%
50th	1.27	1.56	+23%
75th	1.68	1.95	+16%

⇒ Higher transaction costs ⇒ higher volatility **at every quantile**.

Exception: 1st percentile is lower for large ticks (more zero-range hours due to coarser price grid).

Figure 5: CDF Comparison




FIGURE 5. The cumulative distribution of the hourly log percentage range for 23,481 volatility measures in the small tick regime and 23,732 volatility measures in the large tick regime.

Interpretation:

- Large-tick CDF (dashed) lies **to the right** of small-tick CDF (solid) $\Rightarrow$ first-order stochastic dominance
- Difference largest at median ($\sim 20\%$ range increase)
- Effect visible across entire distribution, not just the mean
- Exception only at < 2 nd percentile (zero-range hours)

Regression Analysis

From OLS to FGLS — Why It Matters

OLS (Ordinary Least Squares): minimizes $\sum \varepsilon_i^2$. Assumes all errors have the **same variance** and are **uncorrelated**.

Problem: financial data violates both assumptions:

- **Heteroskedasticity:** some stocks are inherently more volatile $\Rightarrow$ their residuals have larger variance
- **Autocorrelation:** volatile hours cluster together $\Rightarrow$ residuals within a stock are correlated

GLS (Generalized Least Squares): weights observations by the *inverse* of their error variance. Gives more weight to precise observations, less to noisy ones. Requires knowing the exact error structure Ω .

FGLS (Feasible GLS): in practice, Ω is unknown. FGLS *estimates* the error structure from the data (first run OLS, estimate $\hat{\Omega}$ from residuals, then run GLS with $\hat{\Omega}$). Used by Hau to allow stock-specific autocorrelation patterns.

Regression Setup

Panel regression (many stocks $\times$ many hours) with fixed effects:

$$\text{LRANGE}_{i,t} = \underbrace{\alpha_i}_{\text{stock FE}} + \underbrace{\gamma_t}_{\text{hour FE}} + \beta \cdot \text{TICK}_{i,t} + \delta \cdot X_{i,t} + \varepsilon_{i,t}$$

What each part does:

- α_i (stock fixed effects) — absorb permanent differences between stocks (sector, size, liquidity). Like giving each stock its own baseline.
- γ_t (hour fixed effects) — absorb market-wide shocks (news, macro events). Compares stocks *within the same hour*.
- $\text{TICK}_{i,t} = 1$ if stock i has large tick on day t ; this is the key variable
- $X_{i,t}$ — controls: lagged LRANGE, index volatility, proportional tick size

$\Rightarrow \beta$ captures the effect of large tick **after removing** stock characteristics and market conditions.

Table 5: Regression Results

TABLE 5. Volatility regressions.

Dependent Variable	LRANGE (= log hourly percentage range of the stock midprice)		
	Specification I	Specification II	Specification III
TICK DUMMY (Large ticks = 1)	0.2535*** (0.0091)	0.2258*** (0.0074)	0.1241*** (0.0059)
LRANGE(INDEX)	—	0.5680*** (0.0061)	0.5479*** (0.0061)
LRANGE(INDEX)(-1)	—	0.0928*** (0.0060)	-0.0441*** (0.0068)
LRANGE(INDEX)(-2)	—	0.0565*** (0.0060)	-0.0277*** (0.0068)
LRANGE(INDEX)(-3)	—	0.0344*** (0.0061)	-0.0281*** (0.0067)
LRANGE(-1)	—	—	0.2070*** (0.0047)
LRANGE(-2)	—	—	0.0805*** (0.0048)
LRANGE(-3)	—	—	0.0499*** (0.0048)
LRANGE(-4)	—	—	0.0346*** (0.0045)
LRANGE(-5)	—	—	0.0233*** (0.0045)
LRANGE(-6)	—	—	0.0276*** (0.0045)
LRANGE(-7)	—	—	0.0358*** (0.0043)
HOUR DUMMIES	Yes	Yes	Yes
STOCK DUMMIES	Yes	Yes	Yes
Observations	47, 213	45, 950	44, 731
Stocks	26	26	26
Wald χ^2 ($n = 32, 36, 43$)	12940***	27436***	41708***
Log Likelihood	-36680	-31167	-29799
Permanent effect			
TICK DUMMY			0.2293

Three specifications:

- I Basic: $TICK = 0.2535^{***}$
- II + index volatility controls: $TICK = 0.2258^{***}$
(barely changes $\Rightarrow$ not driven by market conditions)
- III + 7 lagged LRANGE: $TICK = 0.1241^{***}$
(short-run effect; permanent effect requires rescaling)

Index volatility coefficient: $0.568^{***} \Rightarrow$ strong co-movement, but doesn't explain TICK effect.

Short-Run vs. Permanent Effect

In Spec. III, lagged LRange captures **volatility clustering** (volatile hours follow volatile hours). The TICK coefficient (0.1241) is only the *one-period shock*.

In steady state, today's shock feeds into tomorrow via lags:

$$\text{LRange}^* = \alpha + \beta \cdot \text{TICK} + (\rho_1 + \rho_2 + \dots + \rho_7) \cdot \text{LRange}^*$$

$$\text{Solving: permanent effect} = \frac{\beta}{1 - \sum \rho_k} = \frac{0.1241}{1 - 0.4587} = \mathbf{0.2293}$$

	TICK coefficient	Implied range increase
Spec. I (static)	0.2535	~35%
Spec. II (+ index vol.)	0.2258	~31%
Spec. III (permanent)	0.2293	~31%

⇒ All three converge: large tick increases volatility by ~30%.

Policy Implications

Implications for the Tobin Tax Debate

- Results support **Friedman's view**: higher transaction costs increase volatility
- The effect is **causal** (natural experiment, not correlation)
- Mechanism: costs discourage **stabilizing arbitrage** more than speculation
- A Tobin-style financial transaction tax would likely **increase** rather than decrease price volatility

Caveat: the experiment captures a specific type of cost (tick-induced spread). A flat-rate tax might have different distributional effects across trader types.