

7496818-who-trades-against-the-maker-adverse-selection-with-counterparty-identity-on-an-

SUMMARY: This study examines the market maker's performance on the Derive platform, which matches orders off-chain and settles them on-chain. The analysis focuses on 603,940 trades in BTC, ETH, and HYPE options between January 2024 and September 2026. The study reveals that the market maker loses, on average, 13.05 USDC per contract 30 minutes after a trade, with adverse selection accounting for a significant portion of this loss. The research finds that the maker's loss is concentrated among a few taker wallets, with 10 wallets accounting for 90.5% of the maker's aggregate loss. The study also identifies the importance of counterparty identity in understanding market maker performance, as the maker's loss is not influenced by trade size or sweeps. The findings suggest that the literature has underestimated the importance of adverse selection and counterparty identity in options trading.

METHOD: The study uses a large dataset of 603,940 trades in BTC, ETH, and HYPE options on the Derive platform. The dataset includes information on the wallets and subaccounts of both parties, the transaction hash, and the realized profit of each side. The research measures the market maker's performance by calculating the mean markout 30 minutes after a trade, and splits the results by counterparty class. The study also includes fixed effects for instrument-by-day to control for instrument-specific effects.

KEY FINDINGS:

- The market maker loses, on average, 13.05 USDC per contract 30 minutes after a trade.
- Adverse selection accounts for a significant portion of the market maker's loss.
- The loss is concentrated among a few taker wallets, with 10 wallets accounting for 90.5% of the maker's aggregate loss.
- Trade size and sweeps do not explain the market maker's loss.
- The counterparty identity is crucial in understanding market maker performance.

IMPLICATION FOR TRADING: The findings suggest that market makers should be aware of adverse selection and the importance of counterparty identity in options trading. The study highlights the need for more research on adverse selection and the impact of counterparty identity on market maker performance. Practitioners can use this information to make more informed decisions about their trading strategies and to identify potential counterparty risks.

Who trades against the maker? Adverse selection with counterparty identity on an on-chain options order book

Gregor Albiez^a

^aUniversity of Applied Sciences and Arts Northwestern Switzerland (FHNW), Switzerland

ARTICLE INFO

Keywords:

market microstructure
options
adverse selection
decentralised exchange
market making
pre-registration
JEL: G13, G14, G12, D47

ABSTRACT

A market maker on an options book loses to counterparties who know something and earns from those who do not, but on an anonymous venue the two can only be told apart by proxy. Derive settles on chain, so its public tape names both wallets. On 603 940 fills in BTC, ETH and HYPE options between January 2024 and September 2026, the maker's mean markout thirty minutes after the fill is +13.05 USDC per contract: a half spread of +15.70 less adverse selection of -2.65. Fees, rebates and delta hedging leave a net edge of +8.93. The loss is extremely concentrated, with ten taker wallets carrying 90.5 per cent of it, and those ten are almost the entirety of the two professional counterparty classes. Trade size and sweeps, the standard proxies for toxic flow, explain none of it once instrument-by-day fixed effects are included. Four hypotheses were registered before any markout was computed; three are rejected.

1. Introduction

A market maker on an options book loses money to counterparties who know something and earns it from those who do not. Telling the two apart is the central problem of the literature that follows Glosten and Milgrom (1985) and Copeland and Galai (1983). Because the counterparty is anonymous, toxicity has to be inferred from what is observable about the trade itself, above all its size and its aggressiveness (Easley and O'Hara, 1987). That inference is indirect, and it is the weakest link in any estimate of what an options maker actually earns.

Derive is an options venue whose book is matched off chain and settled on chain. Every fill lands in a public tape that names both wallets, the subaccount, the request-for-quote identifier where there is one, the fee, the rebate and the realised profit of each side. The counterparty is not a statistical construct but an address with a history. This paper uses that to ask directly what the literature has had to approach by proxy: against whom does an options maker lose? In equity options the question has been answered through the price impact of net demand (Bollen and Whaley, 2004; Gârleanu, Pedersen and Poteshman, 2009) and through the returns that follow order flow (Muravyev, 2016), always with the counterparty unobserved (Easley, O'Hara and Srinivas, 1998).

The answer is that the maker's average result is a statement about the counterparty rather than about the option. Taking the passive side of all 603 940 fills in the sample, the mean markout after thirty minutes is +13.05 USDC per contract, while the median fill returns 0.98. Splitting that average by counterparty class moves it from +24.73 against ordinary takers to -25.41 against the seventeen wallets that dominate maker volume. Nothing about the option changes across that range; only the address on the other side does.

The loss is carried by very few of them. Ten taker wallets account for 90.5 per cent of the maker's aggregate loss, and a single address for 39.0 per cent. The same ten wallets supply 98.8 per cent of the fills of the dominant-maker class and 92.4 per cent of those of the market-maker programme, so the concentration result and the counterparty result are one finding seen twice. What does not follow is the usual explanation. Neither trade size above the ninetieth percentile nor sweeps across several price levels carry a coefficient distinguishable from zero once instrument-by-day fixed effects absorb the composition of the flow. On this venue the standard observables do not find the informed trader; the identity does.

Three contributions follow. The first is a measurement. Section 3 decomposes the maker's result into a half spread, adverse selection, the fee, the rebate and the cost of delta hedging, marked against the volatility surface that the venue itself writes on chain rather than against a mark of my own construction. The second is the counterparty fact above, together with its negative half: the observables that the literature uses in place of identity do not reproduce it. The third is a map. Section 6 reports the net edge per unit of notional across moneyness and tenor for three underlyings, which is the form in which a quoting decision is actually taken.

A fourth contribution is procedural. Four hypotheses, their rejection rules and the inference were committed to the repository before a single markout was computed (Burlig, 2018). Three of the four are rejected. They are reported as such, because a pre-registered rejection is evidence about the market and not a failure of the study: it is the reason the reader can believe the one result that survives.

2. The venue and the tape

Crypto markets are fragmented across venues with very different designs (Makarov and Schoar, 2020). Derive

✉ gregor.albiez@students.fhnw.ch (G. Albiez)
ORCID(s): 0009-0006-6036-4932 (G. Albiez)

matches orders off chain and settles them on an Optimism-stack chain of its own. Positions are held in subaccounts under a portfolio margin system, and the mark that margin is computed against is not a private number. A feed contract receives, for every expiry, the parameters of a stochastic volatility inspired surface (Gatheral and Jacquier, 2014), and anyone can read the mark implied volatility of any strike at any block by evaluating that surface. During the sample the median curve was 32 seconds old at the moment of a fill and the ninetieth percentile was 58 seconds.

The public trade tape is unusual: it is not anonymous. Each fill appears twice, once from each side, and carries the wallet and subaccount of both parties, the transaction hash, the fee and rebate booked to each side, and the realised profit. Where a trade came from a request for quote, the tape carries the identifier that ties its legs together; 20.3 per cent of fills arrive this way.

The sample runs from 11 January 2024 to 17 September 2026 and covers every option fill on BTC, ETH and HYPE: 603 940 fills across 42 510 instruments, 11 573 taker wallets and 7 033 maker wallets. ETH carries 405 250 fills, BTC 154 676 and HYPE 44 014. Liquidations do not appear: they are settled outside the trade tape, so the liquidation class defined in the pre-registration is empty throughout.

Almost everything known about adverse selection in decentralised markets concerns automated market makers, where the liquidity provider does not quote but posts a curve, and where the loss to arbitrageurs is a mechanical consequence of that curve rather than a decision (Milionis, Moallemi, Roughgarden and Zhang, 2022; Cartea, Drissi and Monga, 2024; Lehar and Parlour, 2025). A central limit order book with a maker who chooses a price is a different object, closer to the venues studied by Hasbrouck and Sofianos (1993) and Menkveld (2013) than to a constant-product pool, and the evidence on decentralised market quality has so far been drawn from the pool side (Barbon and Ranaldo, 2026; Aramonte, Huang and Schrimpf, 2021). Prices on this venue are quoted by people, and what follows measures what those quotes earned.

What the venue does not publish is the history of the book itself. Quotes that were never hit, the depth standing at any moment, and cancellations are absent, so every statement in this paper is about trades that happened and none is about the cost of quoting. Section 7 returns to this.

3. What a markout measures

For every fill I take the maker's side, $s = -1$ when the taker bought and $s = +1$ when the taker sold, and compare the traded price P with the venue's mark M at a series of horizons. The markout at horizon τ is $MO_\tau = s(M(t + \tau) - P)$. At $\tau = 0$ this is the half spread, the distance between the fill and the mark at the moment of the fill, and the difference $MO_\tau - MO_0$ is adverse selection: what the mark did afterwards. Figure 1 shows the construction on a single fill. The two pieces sum to the markout by construction, so

the half spread must not be added a second time when the maker's result is totalled.

Three marks are available and they answer different questions. The primary one is the on-chain surface, evaluated at the last curve pushed at or before the horizon, because it is the number the venue itself uses for margin and settlement. The second is the mark attached to the next fill in the same instrument, which is a genuine transaction but arrives at a median distance of 15 026 seconds and is therefore no control at a thirty-minute horizon. The third is the settlement payoff, which is exact but only available once. Their agreement is examined in Figure 9.

Subtracting the maker fee, adding the maker rebate and subtracting a modelled cost of holding the resulting delta to the horizon turns the markout into a net edge. Figure 2 shows the whole identity with an interval on every component. The hedge is the only component that is not observed, and the figure marks it as such; its assumed cost is varied in Section 5.

Four units are used and none of them is redundant. USDC per contract is the quantity the pre-registration fixed. Vol points are the only unit that pools across underlyings whose contracts differ in size by three orders of magnitude. The share of the premium is what a maker books. Basis points of notional is what a quoting decision compares against a perpetual spread. Figure 3 shows why reading the result in only one of them is misleading.

Counterparties are assigned to one class each by a cascade in which the first match wins: liquidation, vault, request for quote, dominant maker, market-maker programme, large wallet, and otherwise ordinary taker. A dominant maker is a wallet that supplied a large share of maker volume in the same month; the programme membership comes from the venue's own published epochs. Figure 5 shows the cascade with the fills it assigns.

4. Pre-registration and inference

Four hypotheses were committed, with their rejection rules, on 17 September 2026 before any markout existed, and two dated addenda resolved questions of definition that arose while the data were being built. H1 holds that toxic flow is concentrated in few wallets *and* identifiable from trade size and sweeps; it is rejected unless the concentration exceeds one half and both coefficients are negative and significant. H2 holds that vault flow is toxic to the maker; it is rejected if the mean markout against vaults is significantly positive. H3 holds that the listing of HYPE options on a centralised venue on 23 June 2026 raised the toxicity of HYPE flow on Derive; it is rejected unless the difference-in-differences estimate is positive, significant, and extreme against a placebo distribution. H4 holds that a net edge survives in a majority of moneyness-by-tenor cells or in short-dated at-the-money options; it is rejected otherwise.

All inference clusters on the taker wallet. Coefficients are tested with a wild cluster bootstrap using restricted residuals and Rademacher weights (Cameron, Gelbach and Miller,

Who trades against the maker?

ETH-20250110-3600-C maker bought 1.00 contracts




Figure 1: What a markout measures. Panels a and b stop at the registered 30 minute horizon, panel c carries every horizon. The shaded band is the interquartile range of the 100 nearest fills in the same cell, so the example can be read against its neighbours rather than on its own.




Figure 2: From the half spread to the net edge. Error bars in panel a are 95 percent cluster bootstrap intervals over taker wallets. The hedge bar is dotted because it is the only component that is modelled rather than observed. Panel b is the distribution of the fee as a share of the option premium, for both sides of the trade.

2008; MacKinnon and Webb, 2017), with $B = 9999$ draws and a fixed seed (Roodman, Nielsen, MacKinnon and Webb, 2019); means and cell averages carry percentile intervals from a cluster pairs bootstrap. Clustering on the taker wallet is the level at which the sampling here is plausibly independent (Abadie, Athey, Imbens and Wooldridge, 2023). Fixed effects are instrument by day and are applied as a within transformation. Every interval reported in this paper was drawn the same number of times as the headline it belongs to.

5. Results

5.1. The average is not a trading outcome

The mean markout is thirteen times the median, and the gap is not a curiosity but the first thing a reader has

to absorb. Weighting each fill by the contracts it moved rather than counting it once takes the average from +13.05 to +0.264 USDC per contract; trimming five per cent from each tail takes it to +6.36. Figure 3 shows why. The three underlyings contribute almost the same aggregate dollar amount to the maker, 3.67, 3.68 and 3.20 million USDC, while their average index prices differ by a factor of 1599. A per-contract average is therefore in large part a statement about contract size. Expressed as a share of the premium actually paid, the distribution a maker books runs from -1.6 per cent at the lower quartile through +1.1 per cent at the median to +21.1 per cent at the upper quartile. Every number that follows is reported in the unit that makes it comparable, and the pre-registered unit is kept alongside.

Who trades against the maker?




Figure 3: How to read the headline number. The shaded core of panel a is linear and everything outside it is logarithmic with equal area per decade. Panel b sets the aggregate dollar gain against the average index price, which is what makes a per-contract average a statement about contract size.




Figure 4: How long adverse selection lasts. The subsample is restricted to fills that have every horizon, so a falling line cannot be a shrinking sample. Bands are 95 percent cluster bootstrap intervals of the median, not dispersion.

5.2. Adverse selection is small, and it does not grow

Restricting the sample to the 511 237 fills that have every horizon, so that a falling line cannot be a shrinking sample, the maker's median result is almost flat from one minute to one day. Figure 4 shows it in three units. In vol points the median moves from 2.09 to 1.89 across the five horizons; in USDC per contract the mean moves from 15.96 to 12.97. Whatever the counterparty learns, it is priced within the first minute and does not continue to accrue against the maker afterwards. Splitting the flow makes the point sharper: the professional classes sit below zero at every horizon and the remaining flow above it, and the two curves are already separated at one minute. This is the first sign that the maker's problem is not information arriving after the trade.

5.3. The counterparty, not the option

Figure 5 reports the six classes with the number of wallets behind each. The ordering is stark, from +24.73 USDC per contract against ordinary takers and +15.21 against

request-for-quote flow down to -20.67 against the market-maker programme and -25.41 against dominant makers. The decomposition explains where the difference comes from, and it is not where the literature would look. Against dominant makers the maker's half spread is already -16.60 before anything happens: the fill is struck on the wrong side of the mark. Adverse selection adds a further -8.81. Against the market-maker programme the split is -5.17 and -15.50. Against ordinary takers the half spread is +25.60 and adverse selection only -0.87. The professional counterparties are not principally trading on information the maker lacks; they are getting a better entry price. That is a different claim, and it points at quoting policy rather than at information.

Two cautions belong here. The two classes that carry the negative numbers rest on seventeen and thirty-three wallets, and their wild cluster bootstrap p -values against zero are 0.100 and 0.101, above the five per cent the pre-registration demands. The percentile intervals exclude zero and the p -values do not, so the class result is reported as a description and not as a test.




Figure 5: Who takes back which part of the spread. G is the number of taker wallets behind a class and p is the wild cluster bootstrap p-value against zero; shaded rows carry fewer than 40 wallets. Medians are shown per component and are not additive, so the mean markout is marked separately.

5.4. Concentration, and the failure of H1

Ten wallets carry 90.5 per cent of the maker's aggregate loss, with an interval from 69.8 to 94.3 per cent that is wide because a single address carries 39.0 per cent of it on its own. Only 1 388 of the 11 573 taker wallets are loss-making for the maker at all. Figure 6 shows the curve with its interval and, beside it, the reason H1 fails. Raw differences look convincing: fills above the ninetieth size percentile and sweeps across levels both show a markedly worse result for the maker. Under instrument-by-day fixed effects the size coefficient is -0.77 with $t = -0.74$ and the sweep coefficient -0.36 with $t = -0.10$. The raw differences were composition: large and aggressive orders are where the professional classes trade, and once the comparison is made within the same instrument on the same day, neither observable carries information about the outcome. Concentration is real; the standard proxies for it are not. H1 is rejected because it asserted both.

5.5. Where an edge survives, and the failure of H4

Of the 97 moneyless-by-tenor cells with at least two hundred fills, 48 carry a net edge whose interval excludes zero, three are significantly negative and all three are low-delta ETH, and 46 are inconclusive. Figure 7 shows them. That the most common verdict is *inconclusive* is itself the finding, and it is not visible in a coloured surface. The pre-registered threshold is a bare majority, and the result sits just under it at 49.5 per cent; at the assumed perpetual half spread of zero it would be 57.7 per cent and at three basis points 38.1 per cent. A hypothesis whose verdict turns on an unobserved hedging cost between zero and one basis point is not one the data can settle. H4 also fails on its second branch: short dated at-the-money BTC is positive at +23.95 USDC




Figure 6: Concentration, and why size and sweeps explain nothing. The horizontal axis of panel a is logarithmic because the first ten wallets carry most of the loss. Panel b sets the raw difference against the coefficient under instrument by day fixed effects, which is where the hypothesis fails.

per contract, so the rejection rule fires either way. Read per unit of notional the map is orderly, rising with the absolute delta and with the tenor in all three underlyings, which is




Figure 7: Where the edge survives. The upper row is the median net edge per notional, the unit a quoting decision uses, shaded within each panel only. The lower row is the registered quantity in USDC and its verdict per cell, with an empty cross where a cell has fewer than 200 fills.

the form Section 6 uses, and which is where the illiquidity premium documented for equity options would lead one to look (Christoffersen, Goyenko, Jacobs and Karoui, 2018).

5.6. One sample or two, and the failure of H3

The sample is not homogeneous and Figure 8 makes that checkable rather than assumed. HYPE enters only in November 2025 and the market-maker programme class cannot be assigned before November 2024, so every pooled average in this paper spans a market that changed composition underneath it. Against that background the event study has no power. The difference-in-differences estimate for the listing of HYPE options on a centralised venue is -1.05 vol points with $t = -1.23$, and 46 per cent of a hundred placebo windows produce a more extreme estimate. The sign is wrong for the hypothesis and the magnitude is indistinguishable from noise drawn before the event; H3 is rejected for want of discrimination, not because the effect was shown to be absent.

6. What this means for a maker

A quoting decision is not taken in USDC per contract. It is taken per unit of risk, against the alternative of quoting the same notional somewhere else, and in that unit the map in Figure 7 is orderly in a way the pre-registered unit hides. The median net edge per notional rises monotonically with the absolute delta and with the tenor in all three underlyings. In BTC it runs from about nine basis points in short dated

at-the-money options to twenty-nine in the thirty to ninety day bucket and higher still in the wings. In ETH the same cells carry roughly one and three basis points. In HYPE the entire surface sits between zero and one basis point, which after any realistic cost of hedging is nothing at all. The first consequence is therefore negative and useful: two of the three underlyings on this venue do not pay a maker to be there across most of the surface.

The second consequence concerns what a maker should spend effort on. The result of Section 5 is that size and aggressiveness carry no information about the outcome once the instrument and the day are held fixed, while the identity of the counterparty carries almost all of it. A maker who prices flow by its observable shape is modelling the wrong variable. A maker who maintains a list of a dozen addresses and widens against them is using the variable that works, and on a venue that settles on chain that list is cheap to build and needs no inference at all. The concentration makes the policy tractable: ten wallets, not ten thousand.

The third is a warning about the map itself. It describes fills that happened at prices that cleared. It does not describe what a new quote would have earned, because a quote that is never hit leaves no trace in the data. A cell with a large positive edge and few fills may be a cell where the maker was rarely willing to quote, not one where quoting was profitable. Sizing a strategy from this map without the book history would mistake selection for opportunity.




Figure 8: The sample over 33 months, and the event study. Panels a and b share a calendar axis; hollow markers mark months with fewer than 1000 fills. Panel c is the distribution of the 100 placebo estimates with the estimated effect marked.

7. What this cannot show

Seven limits bound the reading of the results, and each is stated once. First, the mark is the venue's own surface, so the maker is measured against the number the venue uses rather than against an independent valuation; Figure 9 shows the markout is flat across eight classes of curve age, from 11.30 to 17.01 USDC per contract, so the lag of the curve does not drive the result. Second, the alternative mark taken from the next fill in the same instrument arrives at a median distance of 15 026 seconds and is therefore not a control at a thirty-minute horizon; where the comparison is fair, with a next fill inside three hundred seconds, the two paths correlate at 0.896. Third, one trader may hold several addresses, so every concentration figure in this paper is a lower bound. Fourth, the vault class rests on eight wallets, which is why H2 is reported as a hint that vault flow is not toxic and not as a demonstration. Fifth, liquidations settle outside the trade tape, so the paper says nothing about the flow that a maker probably fears most. Sixth, the volatility unit is undefined for 10 745 fills whose traded price sits outside the arbitrage bounds, and those fills are excluded from every vol-point statement. Seventh, and most importantly, the venue publishes no history of the book, so nothing here measures the cost of quoting, only the outcome of trading.

8. Conclusion

On a venue where the counterparty is an address rather than an inference, the maker's result is mostly a statement about who was on the other side. Adverse selection after the fill is small, stable across horizons and a minor part of the decomposition; the difference between a good and a bad counterparty shows up at the moment of the fill, in the half spread, not in what the mark does afterwards. The loss is carried by ten wallets, and the observables that the literature uses in place of identity do not find them.

Three of four pre-registered hypotheses are rejected, and the rejections are informative rather than empty. Toxicity is concentrated but not identifiable from trade characteristics; the venue's net edge sits so close to the registered threshold that an unobserved hedging cost decides it; and an event study on a thirty-three month panel whose composition changes underneath it has no power. What remains is a measurement, a fact about identity, and a map that says where on this surface a maker is paid at all.

Data, code and pre-registration

Every input to this paper is public. The trade tape and reference data come from the venue's public API, the volatility surfaces from the events of the feed contract on Derive Chain. The code that builds the sample, computes the markouts, runs the inference and draws every figure is at

Who trades against the maker?




Figure 9: Is the on-chain mark good enough? Panel a puts the age of the curve and the distance of the next fill on one axis, which is why the next fill is no control at 30 minutes. A flat line in panel b is the reassurance: the markout does not move with the age of the curve.

<https://github.com/gelatotrader/Derive-Option-Surface>, together with a test suite that covers it. The four hypotheses, their rejection rules and the inference were committed in docs/paper1/PRAEREGISTRIRUNG.md on 17 September 2026 at 15:47 UTC+2, in commit 3fd9caa, before any markout had been computed; two dated addenda resolved definitional questions that arose while the sample was being built and are in the history of the same file. A reader who wants to check the order of events can read that history rather than take this paragraph on trust.

Competing interest

I am building a market-making bot for the venue studied here and have been in contact with its team about obtaining order-level data for a follow-up paper. No data, funding or review was received from the venue for this paper, and every number in it comes from public sources. Readers should weigh the practitioner conclusions in Section 6 in that light.

Use of generative tools

Large language models were used for code, for drafting and for checking the bibliography against publisher records. All results were produced by the code in the repository, and every reference was verified against the publisher, RePEc or the DOI. The author is responsible for the content.

References

- Abadie, A., Athey, S., Imbens, G.W., Wooldridge, J.M., 2023. When should you adjust standard errors for clustering? *Quarterly Journal of Economics* 138, 1–35. doi:10.1093/qje/qjac038.
- Aramonte, S., Huang, W., Schrimpf, A., 2021. DeFi risks and the decentralisation illusion. *BIS Quarterly Review* (December), 21–36.
- Barbon, A., Rinaldo, A., 2026. On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. *Management Science*

doi:10.1287/mnsc.2024.07703. articles in advance; no volume, issue or pages assigned.

- Bollen, N.P.B., Whaley, R.E., 2004. Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59, 711–753. doi:10.1111/j.1540-6261.2004.00647.x.
- Burlig, F., 2018. Improving transparency in observational social science research: A pre-analysis plan approach. *Economics Letters* 168, 56–60.
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* 90, 414–427. doi:10.1162/rest.90.3.414.
- Cartea, Á., Drissi, F., Monga, M., 2024. Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. *SIAM Journal on Financial Mathematics* 15, 931–959.
- Christoffersen, P., Goyenko, R., Jacobs, K., Karoui, M., 2018. Illiquidity premia in the equity options market. *Review of Financial Studies* 31, 811–851. doi:10.1093/rfs/hhx113.
- Copeland, T.E., Galai, D., 1983. Information effects on the bid-ask spread. *Journal of Finance* 38, 1457–1469. doi:10.1111/j.1540-6261.1983.tb03834.x.
- Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. *Journal of Financial Economics* 19, 69–90. doi:10.1016/0304-405X(87)90029-8.
- Easley, D., O'Hara, M., Srinivas, P.S., 1998. Option volume and stock prices: Evidence on where informed traders trade. *Journal of Finance* 53, 431–465. doi:10.1111/0022-1082.194060.
- Gârleanu, N., Pedersen, L.H., Poteshman, A.M., 2009. Demand-based option pricing. *Review of Financial Studies* 22, 4259–4299. doi:10.1093/rfs/hhp005.
- Gatheral, J., Jacquier, A., 2014. Arbitrage-free svi volatility surfaces. *Quantitative Finance* 14, 59–71. doi:10.1080/14697688.2013.819986.
- 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. doi:10.1016/0304-405X(85)90044-3.
- Hasbrouck, J., Sofianos, G., 1993. The trades of market makers: An empirical analysis of nyse specialists. *Journal of Finance* 48, 1565–1593. doi:10.1111/j.1540-6261.1993.tb05121.x.
- Lehar, A., Parlour, C.A., 2025. Decentralized exchange: The Uniswap automated market maker. *Journal of Finance* 80, 321–374. doi:10.1111/jofi.13405.
- MacKinnon, J.G., Webb, M.D., 2017. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics* 32, 233–254. doi:10.1002/jae.2508.

- Makarov, I., Schoar, A., 2020. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics* 135, 293–319. doi:10.1016/j.jfineco.2019.07.001.
- Menkveld, A.J., 2013. High frequency trading and the new market makers. *Journal of Financial Markets* 16, 712–740.
- Milionis, J., Moallemi, C.C., Roughgarden, T., Zhang, A.L., 2022. Automated Market Making and Loss-Versus-Rebalancing. Technical Report. arXiv. arXiv:2208.06046. stand der Recherche 2026-09-18: keine Zeitschriftenpublikation der Originalarbeit gefunden; zirkuliert als arXiv-/SSRN-Preprint (arXiv:2208.06046, Erstversion 08/2022). Ein davon abgeleitetes, eigenstaendiges Folgepaper wurde 2025 bei AFT (Advances in Financial Technologies) veroeffentlicht, ist aber nicht dieselbe Arbeit.
- Muravyev, D., 2016. Order flow and expected option returns. *Journal of Finance* 71, 673–708. doi:10.1111/jofi.12380.
- Roodman, D., Nielsen, M.Ø., MacKinnon, J.G., Webb, M.D., 2019. Fast and wild: Bootstrap inference in stata using boottest. *Stata Journal* 19, 4–60. doi:10.1177/1536867X19830877.