Competition and Manipulation in Derivative Contract Markets*

Anthony Lee Zhang†

December 2018

Please click HERE for the latest version of the paper

Abstract

This paper develops metrics and methods for quantifying the manipulability of cash-settled derivative contract markets. Many derivative contracts, such as futures, options, and swaps, are settled based on price benchmarks, which are calculated based on trade prices of underlying assets. Derivative contract volume is often much larger than the volume of underlying trade used to construct price benchmarks, so these markets may be vulnerable to manipulation: contract holders may trade the underlying asset in order to move benchmarks and influence contract payoffs. I show that derivative contract markets can be much larger than underlying markets without creating large incentives for manipulation, as long as underlying markets are sufficiently competitive. I show how to estimate manipulation-induced benchmark distortions using commonly observed market data. I propose a simple manipulation index which can be used as a diagnostic metric to detect potentially manipulable contract markets, similar to the Herfindahl-Hirschman index (HHI) in antitrust. I apply my results to study contract market competitiveness using the CFTC Commitments of Traders reports, to measure the manipulability of the LBMA gold price benchmark, and to propose a less manipulable design for the CBOE Volatility Index (VIX).

Keywords: derivative contracts, manipulation, regulation

JEL classifications: D43, D44, D47, G18, K22, L40, L50

*The online appendix to the paper is available here. I appreciate comments from Mohammad Akbarpour, Yu An, Sam Antill, Anirudha Balasubramanian, Alex Bloedel, Jeremy Bulow, Shengmao Cao, Gabe Carroll, Juan Camilo Castillo, Scarlet Chen, Wanning Chen, Yiwei Chen, Yuxin (Joy) Chen, Cody Cook, Peter DeMarzo, Rob Donnelly, Tony Qiaofeng Fan, Winston Feng, Robin Han, Benjamin Hebert, Hugh J. Hoford, Emily Jiang, David Kang, Zi Yang Kang, Fuhito Kojima, Nadia Kotova, Eddie Lazear, Wenhao Li, Yucheng Liang, Jacklyn Liu, Jialing Lu, Carol Hengheng Lu, Danqi Luo, Hanno Lustig, Suraj Malladi, Giorgio Martini, Negar Matoorian, Ellen Muir, Evan Munro, Mike Ostrovsky, Paul Oyer, Agathe Pernoud, Peter Reiss, Sharon Shiao, Ryan Shyu, Andy Skrzypacz, Paulo Somaini, Takuo Sugaya, Cameron Taylor, Lulu Wang, Bob Wilson, Milena Wittwer, Alex Wu, David Yang, Jessica Yu, Ali Yurukoglu, Becky Zhang, Jeffrey Zhang, and Mingxi Zhu. I am very grateful to my advisors, Lanier Benkard and Paul Milgrom, as well as my committee members, Darrell Duffie, Glen Weyl and Brad Larsen, for their continuous guidance and support. I especially thank Darrell Duffie, who originally inspired this work.

†Stanford Graduate School of Business, 655 Knight Way, Stanford, CA 94305; anthonyz@stanford.edu.

1 Introduction

This paper studies manipulation in cash-settled derivative contract markets. In these markets, agents trade contracts with payments determined by price benchmarks, which are set based on trade prices of certain underlying assets. Price benchmarks and associated derivative contracts exist for a variety of physical commodities, such as oil, gas, gold, and cattle, as well as financial assets, such as stocks, interest rates and foreign currencies. Some benchmarks aggregate over prices of multiple underlying assets, such as the S&P 500 and the CBOE Volatility Index (VIX). Practically, derivative contracts are simply cash bets on price benchmarks; economically, derivative contracts are liquid and standardized instruments for trading exposure to prices of underlying assets.

An important stylized fact about these markets is that the volume of derivative contracts settled based on a given price benchmark is often much larger than the volume of underlying asset trades used to calculate the benchmark. As a result, agents holding large contract positions may have incentives to manipulate price benchmarks, trading relatively small volumes of underlying assets to move price benchmarks and influence payoffs on their much larger contract positions. Derivative contract markets have been manipulated or suspected of manipulation as long as they have existed; fines for manipulation of oil, gas, foreign exchange, and interest rate benchmarks, among others, have totalled billions of dollars in the past two decades alone. Contract market manipulation harms welfare by making price benchmarks noisier measures of economic fundamentals, causing derivative contracts to be less effective tools for hedging against fundamental uncertainty; manipulation also distorts allocative efficiency in underlying markets, as manipulators create shortages or surpluses of underlying assets in order to move benchmark prices.

Contract market manipulation is illegal in the US and many other jurisdictions, but it is notoriously difficult to prosecute. US authorities currently regulate manipulation using a primarily behavioral approach. Trading with the intent to move prices and influence contract payoffs is a crime *per se*, and regulators often bring cases against manipulators based on evidence from email or telephone communications between traders. This approach is difficult to apply in settings where regulators cannot easily access traders' communications, and it is insensitive to the magnitude of traders' economic incentives to manipulate contract markets.

Antitrust regulation, in contrast, is primarily structural rather than behavioral. Regulators do not punish firms for pricing above marginal cost; regulators instead aim to

keep industries sufficiently competitive that market competition disciplines firms' pricing behavior, without the need for case-by-case regulatory intervention. Operationally, regulators use concentration metrics such as the Herfindahl-Hirschman index (HHI), and methods such as demand system and markup estimation, to monitor the extent of market power-induced distortions in oligopolistic industries; these metrics and methods guide the application of policy interventions, such as blocking mergers and forcing divestitures, for structurally controlling industry competitiveness.

Contract market regulators could in principle adopt a similar approach, controlling the size and structure of contract markets to limit market participants' incentives to manipulate. Regulators have many policy tools, such as contract position limits and margin requirements, for controlling contract market structure. But regulators do not have metrics and methods, like HHIs and demand systems, for measuring contract market manipulability. Contract markets are often much larger than underlying markets, but we do not know how to tell when a contract market is too large. In order to structurally regulate manipulation in derivative contract markets, we need a theory-based way to measure contract market manipulability.

This paper develops metrics and methods to quantify contract market manipulability. I assume that agents with concave utility for an underlying asset hold contract positions tied to a price benchmark. As a reduced-form model, I assume that the benchmark is set in a uniform-price auction for the underlying asset. If an agent holds a large positive contract position, she has incentives to manipulate the benchmark, buying the underlying asset to increase the benchmark price and thus her contract payoffs. However, buying large quantities of the underlying asset is costly to the agent, since it decreases her marginal utility for the underlying asset. I show that the size of an agent's manipulation incentives depend on how large the slope of residual supply facing the agent is, relative to the slope of the agent's own demand for the underlying asset; that is, agents' manipulation incentives depend on how quickly the agent's trades of the underlying asset move benchmark prices, compared to how quickly they change the agent's own marginal utility for the underlying asset.

I define an agent's *capacity share* as the slope of the agent's demand curve divided by the slope of aggregate market demand. I precisely characterize agents' equilibrium manipulation incentives: in equilibrium, an agent with capacity share s_i trades approximately s_i additional units of the underlying asset per unit contract that she holds. Thus, manipulation incentives are small, and contract markets can be much larger than underlying

markets, as long as underlying markets are competitive and all agents' capacity shares are small. Intuitively, the slope of residual supply facing a given agent is approximately the sum over all other agents' demand slopes; in competitive markets, residual supply is thus much steeper than any given agent's demand curve. Trades that the agent makes move her marginal utility for the underlying asset much faster than they move benchmark prices, so manipulation is costly and equilibrium manipulation incentives are small. The relationship between demand slopes and manipulation incentives is quantitatively precise, implying that we can measure agents' equilibrium manipulation incentives in any setting where we can estimate agents' slopes of demand for the underlying asset.

My results can be used to measure the extent to which manipulation distorts price benchmarks: I show that manipulation-induced bias and variance in price benchmarks can be calculated based on agents' contract positions and agents' slopes of demand, both of which can be estimated or approximated using commonly observed contract market data. I also propose the *manipulation index*, a quantity which can be calculated using only aggregate data on total contract volume and underlying trade volume, along with the maximum capacity share among market participants. I show that, under some assumptions, the square of the manipulation index is an approximate upper bound for the share of total benchmark variance which is attributable to manipulation. The manipulation index is thus a simple diagnostic metric that regulators can calculate using widely available data, to identify contract markets which may be vulnerable to manipulation.

My results have implications for the design of aggregate benchmarks, which are calculated based on prices from multiple submarkets. I show that optimal submarket weights for minimizing manipulability are proportional to submarket liquidity and competitiveness; they are not necessarily proportional to submarket volume. Regulators can observe indicators of liquidity and competitiveness in many contract markets in practice, and my theory suggests that benchmark administrators can use these indicators to adjust submarket weights during benchmark calculation. Volume-weighted average price benchmarks, which are commonly used in practice, are not necessarily optimal for minimizing benchmark manipulability.

Using my theoretical results, I study three empirical applications. First, I analyze the Commodity Futures Trading Commission (CFTC) Commitments of Traders reports, which document the positions of large traders in contract markets. I show that competitiveness varies widely across contract markets: I estimate that the capacity share of the largest agent ranges from approximately $\frac{1}{10}$ to $\frac{1}{3}$ across markets. Thus, in some markets, contract volume

can be approximately 10 times greater than underlying volume before manipulation distorts price benchmarks substantially; in other markets, manipulation may be a concern even if contract volume is not much larger than underlying volume. I also show that agents' demand slopes and contract positions appear to be correlated in many contract markets; my theory predicts that benchmark prices in these markets are likely to be biased.

Second, I study the London Bullion Market Association (LBMA) gold price. Gold dealers in London participate in auctions twice each day to set benchmark prices for gold. Using auction bidding data, I estimate auction participants' slopes of demand, allowing me to predict manipulation-induced benchmark variance for any conjectured contract position size. The popular COMEX gold futures contract is currently settled by physical delivery of gold; I ask whether the LBMA gold price benchmark could be used to financially settle COMEX gold futures contracts, without leading to large distortions in benchmark prices. Counterfactually, even if the volume of contracts held by auction participants were 5-10 times larger than the total volume of gold traded in the auction, manipulation would only increase benchmark variance by around 4-15% of the variance in auction-to-auction gold price innovations.

Third, I propose a redesign for the Chicago Board Options Exchange (CBOE) Volatility Index (VIX). The VIX is an index of market volatility, calculated as a weighted sum of prices of S&P500 options. Recently, Griffin and Shams (2018) have provided evidence suggesting that the VIX is being manipulated, as trade volume spikes in far-out-of-money S&P500 options, which are otherwise rarely traded, on VIX settlement dates. I show that the CBOE's current methodology for VIX calculation is very vulnerable to manipulation: my model predicts that, under the CBOE's formula for calculating the VIX, manipulation-induced benchmark variance is approximately 22 times larger than true changes in volatility reflected in the benchmark. I propose a Bayesian volatility estimator which is less manipulable than the VIX; my estimator lowers manipulation variance by a factor of 140, while retaining 68.8% of the VIX's sensitivity to true changes in volatility.

The contribution of this paper is thus to develop methods and metrics for measuring manipulation incentives and manipulation-induced distortions in derivative contract markets, and to illustrate their use in a number of empirical applications. The results of this paper can be used to guide the application of structural policy tools for regulating manipulation in derivative contract markets.

1.1 Related literature

To my knowledge, this is the first paper to quantify the manipulability of derivative contract markets. There are a number of related strands of literature. There is a small but growing literature theoretically and empirically analyzing market and benchmark manipulation. Abrantes-Metz et al. (2012) and Gandhi et al. (2015) analyze LIBOR manipulation. Griffin and Shams (2018) analyzes VIX manipulation. A number of theoretical papers analyze the question of optimal benchmark design, such as Duffie, Dworczak and Zhu (2017), Duffie and Dworczak (2018), Coulter, Shapiro and Zimmerman (2018) and Baldauf, Frei and Mollner (2018). Duffie and Dworczak (2018) and Baldauf, Frei and Mollner (2018) propose using volume-weighted average price schemes, whereas Coulter, Shapiro and Zimmerman (2018) proposes an incentivized announcement scheme which bears some resemblence to an auction. In contrast to these papers, I do not attempt to find an optimal mechanism for benchmark determination in this paper; instead, I adopt a reduced-form model of benchmark setting in order to quantify agents' manipulation incentives. In terms of modelling, this paper builds on the literature on linear-quadratic double auctions (Vayanos, 1999; Rostek and Weretka, 2015; Du and Zhu, 2017; Duffie and Zhu, 2017).

A number of earlier papers qualitatively study price manipulation by agents who hold cash-settled futures contracts. Kumar and Seppi (1992) analyzes the equilibrium of a Kyle (1985) common-valued competitive market-maker model with agents holding futures contracts linked to spot prices of an asset. Dutt and Harris (2005) discusses how to set futures position limits for agents to prevent manipulation from dominating prices. There are also a number of papers about prediction market manipulation. Hanson and Oprea (2008) show that the presence of manipulators can improve price discovery by increasing agents' incentives for costly information acquisition. Other models of prediction market manipulation include Oprea et al. (2008); Chen et al. (2010); Huang and Shoham (2014).

1.2 Outline

The remainder of the paper proceeds as follows. Section 2 discusses institutional background on contract markets and manipulation. Section 3 introduces the model, and derives the main theoretical results. Section 4 analyzes optimal submarket weights for constructing aggregated benchmarks. Section 5 studies the CFTC's Commitments of Traders reports, section 6 studies the LBMA gold price, and section 7 studies the CBOE VIX. Sec-

tion 8 discusses implications of my findings, and section 9 concludes. Proofs, derivations and other supplementary material are presented in the appendix, and supplementary theoretical results are presented in an online appendix.

2 Institutional background

2.1 Price benchmarks and derivative contract markets

In a *derivative contract market*, agents trade a variety of contracts – futures, options, swaps, and others – linked to some underlying asset. This paper focuses on cash-settled derivative contracts, which entitle contract holders to monetary payments linked to *price benchmarks*, which are constructed based on trade prices of the underlying asset. An example of a cash-settled derivative contract is the CME Feeder Cattle futures contract. The CME Feeder Cattle Index is a benchmark which measures the average sale price per pound of cattle in the US over the past seven days, based on cattle sale prices published by the USDA.[1] An agent who holds a long position in CME Feeder Cattle futures, at contract expiration, is entitled to a cash payment totalling 50,000 times the price benchmark, and a short contract holder owes the same amount; effectively, one contract represents 50,000 pounds of exposure to cattle prices.

Practically, derivative contracts are simply bets on price benchmarks. Economically, derivative contracts allow market participants to trade exposure to price risk in underlying assets using liquid and standardized financial instruments. Suppose that a firm wishes to insure against the possibility that cattle prices may rise in the future. In an Arrow-Debreu economy, the firm would purchase contracts with large payoffs in all future states of the world in which cattle prices are high. Such a contract clearly cannot be used in practice; states of the world are too numerous, and are not sufficiently well-defined, to be referenced in enforcible contracts. If a reliable cattle price benchmark exists, a much simpler solution is to write a contract with payoff determined by the level of the benchmark at some future date; if the benchmark accurately reflects the market price of cattle at any given point in time, the contract allows the firm to perfectly hedge its cattle price risk.

The risk-transfer function that derivative contracts play is indispensable in many settings. As Baer and Saxon (1949, ch. 2) describe, many commodity dealers and

[1] See Understanding the CME Feeder Cattle Index

warehouses operate on very low profit margins; in the absence of contract markets, these agents would be exposed to large swings in the price of stored commodities between purchase and sale, and their profit margins would have to be much higher on average in order for them to stay solvent.

Commodity futures contracts, settled by physical delivery of the underlying asset, have existed in the US since the 1800's; cash-settled derivative contracts based on price benchmarks are a more recent invention. The first cash-settled contract was the Eurodollar futures contract, introduced on the Chicago Mercantile Exchange in 1981.[2] Cash-settled contracts have since proliferated widely. Futures and options contracts for livestock[3] and many dairy products[4] are financially settled based on USDA-published average transaction prices for the underlying commodity. Many derivative contracts for energy products, such as oil,[5] gas,[6] and electricity[7] are settled using price benchmarks calculated based on government or industry sources. Interest rate benchmarks such as LIBOR and SOFR are used to determine payments on many floating-rate loans and related products, such as mortgages, student loans, bonds, and interest rate swaps.[8] FX benchmarks such as the WM/Reuters rates are widely used to value portfolios and evaluate execution quality,[9] and a variety of exchange-traded FX futures and options are also financially settled based on these and other benchmarks.[10] There are also a variety of contracts written based on indices derived from aggregated prices of many underlying securities: aggregated indices and associated derivative contracts exist for equities,[11] commodities,[12] FX,[13] volatility,[14] and many others.

Derivative contract activity for any given commodity tends to be concentrated in a relatively small number of markets.[15] As of December 2016, over half of wheat futures

[2] Market Begins Trading In Eurodollar Futures

[3] CME Lean Hog Futures Contract Specs and Feeder Cattle Futures Contract Specs.

[4] CME Class III Milk Futures Contract Specs and Cash-Settled Cheese Futures Contract Specs.

[5] ICE Brent Crude Futures and Platts Dubai Swap.

[6] ICE Fixed Price Swap - Henry Hub - Tailgate, Louisiana

[7] ICE ERCOT North 345KV Real-Time Peak Daily Fixed Price Future

[8] Financial Stability Board (2018)

[9] Financial Stability Board (2014)

[10] CME Korean Won futures are settled based on a benchmark price reported by Seoul Money Brokerage Service Limited, which represents activity in the Korean Won spot market, and the CME Chinese Renminbi futures are settled based on a similar benchmark reported by the Treasury Markets Association, Hong Kong.

[11] FTSE 100 Index

[12] Bloomberg Commodity Index

[13] U.S. Dollar Index

[14] VIX

[15] This is likely because of increasing returns to liquidity. Agents who wish to buy or sell wheat or oil

open interest recorded by the CFTC was based on the Chicago SRW contract.[16] Over 3/4 of all open interest in crude oil futures in December 2016 was in contracts based on WTI and Brent crude, and over half of gas futures open interest was in contracts based on gas traded at Henry Hub, Louisiana.[17] In financial markets, LIBOR and SOFR are used as the main USD interest rate benchmarks for determining floating interest rates on a large variety of floating-rate loans and other derivatives.

Derivative contracts also tend to be highly leveraged, so investors can purchase large volumes of exposure to underlying assets at low upfront capital cost. An investor buying ICE Brent Crude futures can purchase exposure to 1,000 barrels of oil, worth approximately $70,000 at the time of writing, by depositing at ICE a margin payment of approximately $3,500.[18] Hence, each dollar an investor spends on oil futures buys exposure to $20 of underlying oil. Other derivative contracts, such as swaps and options, similarly allow market participants to trade large volumes of exposure to underlying asset prices with low capital costs.

Due to their concentration and leverage, the volume of derivative contracts settled using a given benchmark is often much larger than the volume of trade in underlying assets used to construct the benchmark. In a few examples, the ratio of contract volume to underlying volume ranges between 3 and 50. The Platts *Inside FERC* Houston Ship Channel natural gas price benchmark is based on weekly trade volume totalling around 1.4 million MMBtus of natural gas;[19] open interest in the ICE HSC basis future, which is financially settled based on the Platts benchmark, is more than 75 million MMBtus for many delivery months.[20] The ICE Brent Crude index is based on around 75 million barrels

for commercial use may have specific demands for the variety of wheat purchased, as well as its delivery location and time, leading to a large number of differentiated spot markets for the underlying asset. In contrast, hedgers and speculators who wish to trade exposure to global wheat or oil prices, who do not care about the specific variety or delivery location of the underlying commodity, will tend to trade in whichever contract market is largest and most liquid, leading to fewer and more concentrated contract markets.

[16] My calculations based on the CFTC Commitments of Traders report for December 27, 2016. In a CME report titled "Understanding Wheat Futures Convergence", Fred Seamon writes that "Often, because of its liquidity, the CBOT Wheat futures contract is used as a benchmark for world wheat prices".

[17] My calculations based on the CFTC Commitments of Traders report for December 27, 2016.

[18] Brent Crude Futures margin rates.

[19] See "Liquidity in Noth American Monthly Gas Markets" on the Platts website. In 2008, daily trade volume during "bid week," the five business days during which index prices are set, was between 50 to 280 thousand MMBtu per day; multiplying by 5, we get an upper bound of 1.4 million. Volume in earlier years is much larger; in 2008-2010, many delivery months have total trade volume over 10 million MMBtu.

[20] ICE Report Center, End of Day reports for HSC basis futures, as of October 24th, 2018. Open interest is above 30,000 contracts for many delivery months, and the contract multiplier is 2,500 MMBtus.

of physical oil traded in the North Sea each month;[21] open interest in Brent Crude futures is over 200 million barrels of oil for some delivery months.[22] The CME Feeder Cattle Index is based on 10-20 million pounds of cattle traded each week;[23] open interest in CME feeder cattle futures for many delivery months is over 500 million pounds of cattle.[24] The Secured Overnight Financing Rate (SOFR), designed to replace LIBOR as an interest rate benchmark, is based on average daily volumes of approximately $800 billion of overnight treasury-backed repo loans;[25] as of 2014, the total notional volume of contracts linked to LIBOR was estimated to be greater than $160 trillion.[26] Since contract markets are much larger than underlying markets, agents holding large contract positions may attempt to manipulate contract markets, trading relatively small volumes of the underlying asset to move price benchmarks and generate large profits in contract markets.

2.2 Manipulation

The basic strategy of contract market manipulation is as follows. Suppose that a gas trader is able to acquire a long position in Houston Ship Channel (HSC) basis futures, totalling 10 million MMBtu of natural gas. If gas is currently trading at $3 USD/MMBtu at HSC, the gas trader can then purchase a relatively large amount of physical gas at HSC – for example, 2 million MMBtus – thus raising the Platts *Inside FERC* price at HSC to, for example, $5 USD/MMBtu. The trader makes a loss in the underlying market, as she purchases physical gas at an elevated price, and accumulates a large amount of physical gas that she has no commercial use for; however, by raising the gas benchmark price, she increases her total contract payoff by $20 million. Symmetrically, the trader could accumulate a large short contract position and then sell large quantities of gas in order to lower benchmark prices. The result of both kinds of manipulation is that the trader makes losses from buying or selling large undesired quantities of the underlying asset at unfavorable prices, but profits from moving price benchmarks and thus increasing her

[21] ICE Brent FAQ.

[22] ICE Report Center, End of Day reports for Brent Crude futures, as of October 24th, 2018.

[23] See "Feeder Cattle Daily Index Data" at the CME website.

[24] As of 2018-09-29, open interest in the Nov 18 and Jan 19 contracts were over 20,000 and 15,000 respectively, and the contract multiplier is 50,000 pounds.

[25] NY Fed's Secured Overnight Financing Rate Data.

[26] Financial Stability Board (2018). Note that the $160 trillion number measures the total volume of outstanding contracts across expiration dates; the volume of interest rate derivatives expiring on any given settlement date will be substantially smaller, so I do not include this in calculating the range of ratios of contract volume to underlying volume.

derivative contract payoffs.

Contract market manipulation is socially harmful for two reasons. First, the HSC benchmark price becomes a noisier signal of true supply and demand conditions at HSC for natural gas, so derivative contracts based on the benchmark become noisier hedges against shifts in true supply and demand for natural gas. Second, allocations of gas are distorted, as manipulative trades made to move benchmark prices may create shortages or surpluses of gas at HSC.

A number of examples illustrate the basic structure of contract market manipulation. In 1888, Benjamin Hutchinson cornered the wheat market by building up a large long position in wheat contracts, then chartering a fleet of ships to move wheat out of Chicago; Hutchinson successfully raised the price of wheat from below $1 per bushel to over $2.[27] In 1980, Nelson Bunker Hunt and his brother Herbert attempted to corner the world silver market, taking on large long positions in silver futures, then purchasing enough physical silver to drive silver prices from $6 per ounce in 1979 to around $50 in 1980. The attempt ultimately failed, the Hunt brothers were forced to close out their positions, and silver prices dropped from $39.50 to $10.80 in a single day.[28] In 2005, Energy Transfer Partners, L.P. and three subsidiaries entered into basis swap contracts with payoffs linked to the *Inside FERC* Houston Ship Channel (HSC) natural gas index, then sold large quantites of physical gas at HSC to drive prices down; in 2008, the Commodity Futures Trading Commision (CFTC) fined Energy Transfer Partners $10 million for its actions.[29]

Manipulation has haunted contract markets since their inception, but it is notoriously hard to define and prosecute. In the US, manipulation and attempted manipulation of contract markets is illegal under the Commodity Exchange Act of 1936.[30] However, the act does not define manipulation, leaving interpretation entirely to the courts. The CFTC operationally defines manipulation as trading to intentionally and successfully create "artificial prices" which do not reflect "legitimate forces of supply and demand".[31] But

[27] Markham (2014, pg. 21)

[28] Markham (2014, pg. 180)

[29] CFTC Press Release 5471-08

[30] 17 CFR Part 180, Rule 180.2, "Prohibition on price manipulation", states: "It shall be unlawful for any person, directly or indirectly, to manipulate or attempt to manipulate the price of any swap, or of any commodity in interstate commerce, or for future delivery on or subject to the rules of any registered entity."

[31] 17 CFR Part 180: "the Commission reiterates that, in applying final Rule 180.2, it will be guided by the traditional four-part test for manipulation that has developed in case law arising under 6(c) and 9(a)(2): (1) That the accused had the ability to influence market prices; (2) that the accused specifically intended to create or effect a price or price trend that does not reflect legitimate forces of supply and demand; 128 (3) that artificial prices existed; and (4) that the accused caused the artificial prices."

this definition is still difficult to apply in legal proceedings: prosecutors must prove both that manipulation was intentional and it successfully created "artificial prices", which is similarly difficult to define precisely. The CFTC achieved limited success combating manipulation using the language of the Commodity Exchange Act, losing many of the manipulation cases it brought to court in its early days.[32]

Since the turn of the millenium, the legal standard for proving manipulation has been decreased substantially. Congress granted broad anti-manipulative authority in derivative contract markets to the Federal Energy Regulatory Commission (FERC) in 2005 and the Federal Trade Commission (FTC) in 2007, and the CFTC's authority was revised under the Dodd-Frank Act of 2010. While legislators left the language of the original Commodity Exchange Act in place, they granted regulators additional authority to prosecute market participants for employing "manipulative or deceptive devices and contrivances" in trading. The new language incorporates a broader, though similarly vaguely defined, body of behavior, and additionally does not require prosecutors to prove that market participants successfully moved benchmark prices. Armed with new legal tools for combating manipulation, regulators have policed manipulation aggressively in the recent two decades. The CFTC and the FERC have fined energy traders millions of dollars for manipulating oil and gas futures markets.[33] Fines for financial market manipulation are orders of magnitude larger: banks have been fined over $10 billion for FX manipulation,[34] over $8 billion USD for manipulation of LIBOR and other interest rate benchmarks,[35] and over $500 million for manipulation of the ISDAFIX interest rate swap benchmark.[36]

In most recent manipulation cases, authorities have adopted a primarily behavioral approach to regulation. Manipulation charges are brought largely based on "smoking gun" evidence of manipulative intent: regulators search through traders' email and telephone communications to find claims demonstrating that traders acted with the intent to move benchmark prices.[37] But it is difficult to behaviorally regulate manipulation

[32]Markham (2014, pg. 204)

[33]See, for example, CFTC Press Release 6041-11, 128 FERC 61,269, and

[34]Levine (2015)

[35]Ridley and Freifeld (2015)

[36]See Leising (2017). A number of other manipulation cases are discussed throughout the paper, and in appendix C.4 in particular. Manipulation is also illegal in many jurisdictions other than the US, although there are differences in how it is defined and investigated; Annex 1 of the International Organization of Securieties Commissions (IOSCO) report on Investigating and Prosecuting Market Manipulation surveys approaches taken in different jurisdictions.

[37]Levine (2014) quotes a number of trader chat messages used in FX manipulation lawsuits; some other

when legislators, regulators, and academics have not reached a consensus definition of manipulation precise enough to use in court. The role that traders play in contract markets is precisely to manage and optimize the price impact of their trades; it is impractical to prohibit traders' acknowledgement of price impact *per se*, but in many cases it is difficult to disentangle legitimate versus illegitimate forms of price impact.[38]

Putting aside the difficulty of defining manipulation, the behavioral approach to regulation is also insensitive to the economics of contract market manipulation. In cases where where regulators are able to obtain sufficient evidence to prosecute, massive punitive fines are levied against manipulators, which may not be proportionate to agents' *ex ante* incentives to manipulate. On the other hand, many economically harmful cases of manipulation may be unprosecutable in settings where regulators cannot easily access traders' commmunications.

An approach more grounded in economic theory would be to regulate markets structurally, limiting the size of contract markets so that market participants do not have large economic incentives to manipulate. But we do not have a structural economic understanding of manipulation in derivative contract markets. We do not know which markets are most vulnerable to manipulation, or how large manipulation-induced distortions will be, despite the fact that manipulation repeatedly occurs across contract markets in practice. Since contract markets are often much larger than underlying markets, limiting contract volume to the size of underlying markets would effectively regulate most markets out of existence. But we do not know how much larger contract markets can be than underlying markets without producing excessively large incentives for manipulation.

Regulators have a variety of policy tools to control the size and structure of contract markets; regulators impose speculative position limits on the size of agents' contract positions,[39] and impose rules for how benchmarks are constructed using trade prices from

examples of evidence for manipulation drawn from traders' communications can be found in the CFTC's lawsuits brought against Parnon Energy, Inc. and others for crude oil manipulation, Energy Transfer Partners, L.P. and others for natural gas manipulation, and Barclays for ISDAFIX manipulation.

[38] As an example of the difficulties involved in defining manipulation, the CFTC dismissed a 1977 cotton futures manipulation case against the Hohenberg Brothers Company, writing: "Even though respondents' activities may have involved a "profit motive," absent a finding of manipulative intent, trading with the purpose of obtaining the best price for one's cotton does not constitute, in itself, a violation of the Commodity Exchange Act." The legal literature has proposed a number of definitions for manipulation, but has not reached consensus. For example, Perdue (1987) proposes defining manipulation as "conduct that would be uneconomic and irrational, absent an effect on market price." Fischel and Ross (1991) argue that manipulation can only be defined with respect to the intent of the trader, but also argue that the law should not prohibit manipulation.

[39] For example, see the CFTC's website on Speculative Limits

underlying markets.[40] In order to apply these tools to effectively combat manipulation, regulators need theory-based metrics and methods to measure manipulation incentives and manipulation-induced distortions in derivative contract markets.

3 Model

In the baseline model of the paper, I assume that a price benchmark for a single underlying good is determined in a uniform-price double auction. Agents participating in benchmark setting may be heterogeneous, and all agents can both hold exogeneously determined derivative contract positions and trade the underlying asset. Agents holding contracts have incentives to trade the underlying asset to move benchmarks and thus contract payoffs; the primary force limiting agents' ability to manipulate is that agents have decreasing marginal utility for the underlying asset, so it is costly for agents to trade volumes of underlying assets to move benchmark prices.

Formally, benchmarks are determined in a game between $n \geqslant 3$ agents indexed by i, with four stages:

1. Agents draw privately observed demand shocks y_{di} and contract positions y_{ci}.
2. Agents simultaneously submit bid curves $z_{Di}(p; y_{di}, y_{ci})$ in a uniform-price double auction.
3. The auction clears at price p_b; agent i receives $z_{Di}(p_b; y_{di}, y_{ci})$ units of the underlying asset and pays $p_b z_{Di}(p_b; y_{di}, y_{ci})$.
4. The benchmark price p_b is used to settle contracts: an agent holding contract position y_{ci} receives monetary payment $p_b y_{ci}$.

I discuss agents' utility functions, derivative contract positions, and the auction model for benchmark setting in subsections 3.1, 3.2 and 3.3 respectively. The following subsections solve for agents' best responses and equilibrium behavior, derive agents' equilibrium manipulation incentives, and show how to measure manipulation-induced benchmark distortions.

[40] For example, see the Financial Conduct Authority website on EU Benchmarks Regulation, and the IOSCO report on Principles for Financial Benchmarks.

3.1 Utilities and demand shocks

I assume that agent i's utility for trading z units of the underlying asset, and paying net monetary transfer t, is:

$$U_i(z,t) = \pi z + \frac{y_{di}z}{\kappa_i} - \frac{z^2}{2\kappa_i} - t \tag{1}$$

The $\frac{z^2}{2\kappa_i}$ term implies that agents have linearly declining marginal utility for the underlying asset. Declining marginal utility can be thought of as driven by, for example, agents' capital costs for purchasing the underlying asset, or their physical infrastructure costs for storing the underlying asset once it is purchased.[41] κ_i represents the slope of agent i's demand for the underlying asset; κ_i will be large for agents who can trade large amounts of the underlying asset without large changes in their marginal utility for the asset. For example, a large oil pipeline operator firm will have a relatively high slope of demand, because it has large financial resources, so that it has relatively low costs of capital, and because it has stocks of oil to sell and tanks to store oil that it purchases, so that it has relatively low infrastructure costs for trading oil. Retail investors will have lower demand slopes, since they have limited financial resources, so their capital costs are high, and since they do not have stocks of oil to sell or storage space for oil that they purchase.

The term y_{di} is a *demand shock*, which vertically shifts i's marginal utility for the underlying asset. y_{di} is normalized by κ_i in expression (1), so that an agent who faces a unit increase in y_{di} optimally purchases an additional unit of the asset at fixed prices. I assume that y_{di} is privately observed by i and independently distributed across agents. For analytical convenience, I assume that the means of agents' demand shocks sum to 0:

$$\sum_{i=1}^{n} E[y_{di}] = 0 \tag{2}$$

This is not a substantive assumption, as we can always define the constant term π such that (2) holds.

I assume that agents have private values, so a given agent's utility for the underlying asset is not affected by other agents' private information. Many of the underlying markets for benchmark setting, such as FX, interest rates, oil and gas, are commodity-like;

[41] My model is static, and I do not explicitly consider agents' decisions to trade the asset over time, but the quadratic utility term could also be thought of as representing curvature of the agent's value function in z in a richer dynamic model.

information about the value of the underlying asset is likely to be close to symmetric, so trading is likely to be driven mostly by preference shocks rather than asymmetric information about the value of the asset. I analyze an extension of the model in which agents have interdependent values in section 3 of the online appendix. The assumption that agents' utilities are quadratic in z is also strong; this is needed because it is difficult to solve for equilibrium in multi-unit auctions with general utility functions. However, in section 1 of the online appendix, I show that my results in subsection 3.4, characterizing agents' manipulation incentives given the slope of residual auction supply, approximately apply in settings with general nonlinear utility functions and residual supply curves.

3.2 Derivative contract positions

In practice, market participants trade complex derivative contracts with payoffs that may be nonlinear functions of benchmark prices; for simplicity, I assume that agents' derivative contract payoffs are linear in benchmark prices. Formally, I model derivative contracts held by agent i as a random, exogeneously determined *contract position* y_{ci}, where i receives monetary payment $p_b y_{ci}$ if the benchmark price is p_b. y_{ci} does not need to be positive, as agents can be positively or negatively exposed to prices of the underlying asset.

From expression (1) above, the utility of agent i with demand shock y_{di} and contract position y_{ci}, if she buys z units of the underlying asset and the benchmark price is p_b, is:

$$U_i(z, p_b;\ y_{di}, y_{ci}) = \pi z + \frac{y_{di} z}{\kappa_i} - \frac{z^2}{2\kappa_i} - p_b z + p_b y_{ci}$$

where the agent pays $p_b z$ to purchase z units of the underlying asset, and receives $p_b y_{ci}$ in contract payoffs.

I assume that y_{ci} is privately observed by i and independently drawn across agents. I additionally assume that i's contract position is independent of her demand shock y_{di}.[42] I do not require contract positions to sum to 0 across the n agents in my model. Since derivative contracts are zero-sum bets on prices, contract positions must sum to 0 across agents in the broader economy; we can think of other contracts being held by unmodelled agents who do not participate in benchmark setting, perhaps because they have very high

[42]This assumption is not essential for solving the model, but it allows us to decompose benchmark price variance cleanly into components attributable to demand shocks and contract positions.

costs of trading the underlying asset.

I assume that contract positions are exogeneous. In practice, market participants may hold derivative contracts for many reasons: to hedge existing risk, to speculate on prices of the underlying asset, or even to profit from attempted manipulation. However, agents holding a given contract position have the same incentives to manipulate, regardless of the original reason for entering into contract positions. Moreover, regulators can directly observe the size of agents' contract positions at settlement. For the purpose of analyzing agents' manipulation incentives given the observed size of contract positions, it is not necessary to take a stance on why agents originally choose to hold contract positions.[43]

Contract positions at settlement are zero-sum bets on prices; they contain no information about demand and supply for the underlying asset at settlement, and thus should not influence benchmark prices. Put another way, if an agent cannot move benchmark prices by trading the underlying asset, any contract position she holds is equivalent to a fixed monetary transfer, and does not affect her trading decisions. However, agents who have price impact have incentives to trade the underlying asset to that moves benchmark prices and contract payoffs. In the context of my model, I define manipulation as the effect that contract positions have on agents' trading decisions in the underlying asset.[44]

While I do not explicitly model the welfare effects of manipulation in the main text of the paper, we can think of manipulation as leading to two main sources of welfare loss. First, manipulation causes price benchmarks to be noisier signals for supply and demand conditions in underlying markets. If we think of the common component of agents' values, π, as being uncertain *ex ante*, risk-averse agents may originally enter into contracts to hedge against exogeneous exposures to π. For example, airlines have exogeneous positive exposure to oil prices, and may enter into short contracts to hedge their exogeneous

[43] A natural question is why agents would not build up extremely large contract positions just before settlement, then manipulate the underlying market to increase contract payoffs. There are a number of reasons why this is difficult in practice. First, manipulation is illegal in the US, and entering large contract positions close to settlement dates and then trading aggressively in underlying markets would be strong evidence of manipulative intent. Second, while I treat contract positions as exogeneous in this paper, in practice there is a decreasing market demand curve for contracts in practice, so agents entering into large positive contract positions face progressively higher prices. Third, every contract has a counterparty, and if a large agent blatantly attempted to purchase large contract positions for manipulation, at some point she would be unable to find contract counterparties at reasonable prices.

[44] In some markets, agents holding expiring contracts will "rollover" contract positions, taking on an equivalently sized position in the underlying asset just as contracts expire. This is conceptually distinct from manipulation: agents who do not have price impact may still wish to rollover contract positions, whereas I define manipulation in my model as trading the underlying asset to move benchmark prices and increase contract payoffs, which is only possible when agents have nonzero price impact.

price risk; oil drillers have exogeneous long exposures, and may hold short contracts to hedge. As I discuss in subsection 3.6 below, manipulation at contract settlement then adds variance to the benchmark price p_b, making derivative contracts less effective tools for hedging exposure to π and exposing contract holders to additional risk.

Second, manipulation distorts allocative efficiency in underlying markets. In order to move benchmark prices, manipulators buy or sell large quantities of underlying assets, which can creates shortages or surpluses of the underlying asset at the time of contract settlement. Manipulation can distort allocative efficiency even if it does not add variance to price benchmarks: if a large trader systematically holds long oil contract positions at contract settlement, she will buy oil at each contract settlement event, causing benchmark prices to be biased upwards. Benchmark bias need not affect the utility of hedgers, as *ex ante* derivative contract prices may reflect foreseen manipulation; however, the manipulator must still buy large quantities of oil at the point of benchmark determination, taking oil away from commercial users and decreasing allocative efficiency.

I construct a simple model which qualitatively demonstrates both of these sources of welfare losses in section 2 of the online appendix. For the remainder of the main text, I assume that it is undesirable for agents' contract positions to influence trading decisions and thus benchmark prices, without imposing a formal model of social welfare.

3.3 Benchmark determination

As a reduced-form model, I assume that benchmarks are set in a uniform-price double auction. After observing demand shocks y_{di} and contract positions y_{ci}, agents simultaneously submit affine and strictly decreasing bid curves $z_{Di}(p; y_{di}, y_{ci})$, representing the net amount of the underlying asset they are willing to purchase at price p. The auction clearing price is the unique price that equates supply and demand:

$$p_b = \left\{ p : \sum_{i=1}^{n} z_{Di}(p; y_{di}, y_{ci}) = 0 \right\}$$

Each agent then pays $p_b\, z_{Di}(p; y_{di}, y_{ci})$, and purchases net quantity $z_{Di}(p; y_{di}, y_{ci})$ of the underlying asset. The auction clearing price p_b, which is also the benchmark price, is then used to settle agents' contracts, so agent i receives a transfer of $p_b y_{ci}$ in contract payments.

No single model can perfectly match all of the many benchmark-setting mechanisms used in practice, but auctions are arguably a reasonable approximation to many mechanisms. Some benchmarks, such as the CBOE VIX and the LBMA gold price, are in fact set using uniform-price auctions for the underlying assets. Benchmarks for some assets are based on trades of a relatively homogeneous commodity over a relatively short period of time; examples are benchmarks for oil and gas, the WM/Reuters FX benchmark and the and the ISDAFIX interest rate swap benchmark. For these benchmarks, it may be reasonable to assume that all trades within the benchmark-setting period happen at a uniform price, in which case the uniform-price auction is potentially a reasonable model of benchmark setting.

Auctions are a less appropriate model for benchmarks based on underlying markets with large distortions other than market power, and the analysis of this paper may not be suitable for these settings. Some assets are traded in decentralized markets with large search frictions; the assumption that all trades happen at the same price is not appropriate. In other markets, such as the interbank loan market which the LIBOR interest rate benchmark is based on, trades are difficult to verify, so false reporting is an important concern. I further discuss the extent to which my model applies to different benchmark-setting mechanisms in subsection 8.2.

3.4 Best responses

From the perspective of agent i, auction equilibrium defines a *residual supply curve*, $z_{RSi}(p)$, specifying the number of units i is able to trade at price p. The residual supply curve facing agent i is the negative of the sum of all other agents' bid curves:

$$z_{RSi}(p) = -\sum_{j \neq i} z_{Dj}\left(p; y_{jd}, y_{jf}\right) \tag{3}$$

In equilibrium, given the assumption that agents' utility functions are linear-quadratic, residual supply facing a given agent i will be affine with a fixed slope:

$$z_{RSi}(p, \eta_i) = d_i(p - \pi) + \eta_i \tag{4}$$

where d_i is the slope of residual supply, and η_i is a random term which incorporates uncertainty in other agents' submitted bid curves, induced by uncertainty in their demand

shocks and contract positions. When bidding in uniform-price auctions, agents cannot separately choose the quantity they purchase for every possible realization of residual supply; they can only submit bid curves $z_{Di}(p; y_{di}, y_{ci})$, purchasing the quantity and price pair at which demand and residual supply are equal. However, this constraint does not restrict agents' behavior in the linear-quadratic auction model, as it is possible to submit an *ex-post optimal* bid curve, which traces through the set of optimal price-quantity pairs for all possible realizations of residual supply. Formally, if agent i submits bid curve $z_{Di}(p; y_{di}, y_{ci})$ and residual supply is $z_{RSi}(p)$, define the price which equates residual supply to i's bid as:

$$p\left(z_{Di}(p; y_{di}, y_{ci}), z_{RSi}(p, \eta)\right) \equiv \{p : \ z_{Di}(p; y_{di}, y_{ci}) = z_{RSi}(p, \eta)\} \tag{5}$$

Definition 1. Bid curve $z_{Di}(p; y_{di}, y_{ci})$ is *ex-post optimal* with respect to the residual supply function (4) if $z_{Di}(p; y_{di}, y_{ci})$ implements the optimal price for every realization of η_i; that is,

$$p\left(z_{Di}(p; y_{di}, y_{ci}), z_{RSi}(p, \eta_i)\right) = \arg\max_p U_i\left(z_{RSi}(p, \eta_i), p; y_{di}, y_{ci}\right) \ \forall \eta_i$$

When agents with quadratic utility functions face affine residual supply functions with additive uncertainty, the unique ex-post optimal bid curves are also linear, and depend only on the slope of residual supply d_i, not the distribution of η_i.[45]

Proposition 1. *If agent* i *has demand shock* y_{di} *and contract position* y_{ci}*, and the slope of residual supply is* d_i*, then agent* i*'s unique ex-post optimal bid curve is:*

$$z_{Di}(p; y_{di}, y_{ci}) = \frac{d_i}{\kappa_i + d_i} y_{di} + \frac{\kappa_i}{\kappa_i + d_i} y_{ci} - \frac{\kappa_i d_i}{\kappa_i + d_i}(p - \pi) \tag{6}$$

According to proposition 1, bid curves are increasing in both i's demand shock, y_{di}, and i's contract position, y_{ci}. Intuitively, when i's demand shock y_{di} is high, i's marginal utility for the underlying asset is high, so i increases her bid curve to purchase more of the underlying asset. When i's contract position y_{ci} is positive and large, agent i want benchmark prices to be higher to increase contract payoffs, so i also increases her bid to purchase more of the underlying asset. However, the coefficients on y_{di} and y_{ci} differ.

A natural measure of i's manipulation incentives is the coefficient $\frac{\kappa_i}{\kappa_i + d_i}$ on y_{ci}; I

[45]This is well known in the literature on multi-unit double auctions with quadratic utility (Klemperer and Meyer, 1989; Vayanos, 1999; Rostek and Weretka, 2015; Du and Zhu, 2012).

will call this i's *manipulation coefficient*. This coefficient measures the pass-through of i's contract positions into i's bids: if i's contract position increases by one unit, she increases her bid curve by $\frac{\kappa_i}{\kappa_i+d_i}$, so she is willing to buy $\frac{\kappa_i}{\kappa_i+d_i}$ more units of the underlying asset at any given price.

The manipulation coefficient depends on the relative sizes of i's demand slope κ_i and the slope of residual supply d_i. Intuitively, the benefits from manipulation are that trading the underlying asset moves benchmark prices and thus contract payoffs; the costs are that trading the underlying asset changes agents' marginal utility for the underlying asset. When d_i is large relative to κ_i, i's trades move i's marginal utility for the asset much faster than they move benchmark prices, so the costs of manipulation are large relative to the benefits. Thus, i has low incentives to manipulate, the ratio $\frac{\kappa_i}{\kappa_i+d_i}$ is small, and i's contract positions have small effects on i's trading decisions. In the limit as d_i approaches infinity, the manipulation coefficient decreases to 0: if i's trades do not affect benchmark prices, i has no incentive to manipulate, so i's contract positions do not affect i's bids.[46]

Proposition 1 thus implies that manipulation coefficients are small, so contract positions have small effects on agents' trading decisions, when underlying markets are competitive and residual supply slopes are large relative to agents' demand slopes. I proceed to analyze how residual supply slopes and manipulation coefficients are determined in auction equilibrium.

3.5 Manipulation incentives in equilibrium

Following the literature, I define a *linear ex-post equilibrium* as a collection of bid curves for all agents, such that each agent is best-responding to the residual supply functions induced by other agents' bid curves.

Definition 2. Bid curves $z_{Di}(p; y_{di}, y_{ci})$ constitute a *linear ex-post equilibrium* if:

1. As in (6), agents' bid curves are ex-post best responses given the slope of residual supply d:

$$z_{Di}(p; y_{di}, y_{ci}) = \frac{d_i}{\kappa_i + d_i} y_{di} + \frac{\kappa_i}{\kappa_i + d_i} y_{ci} - \frac{\kappa_i d_i}{\kappa_i + d_i}(p - \pi) \tag{7}$$

[46] In contrast, as d_i approaches infinity, the coefficient on demand shocks y_{di} in bids approaches 1; in a perfectly competitive market, agents bid their true marginal utility for the underlying asset, unaffected by any contract positions they may hold.

2. The slope of residual supply facing agent i is the sum over all other agents' bid slopes:

$$d_i \equiv z'_{RSi}(p) = -\sum_{j \neq i} z'_{Dj}(p; y_{dj}, y_{cj}) = \sum_{j \neq i} b_i \tag{8}$$

where,

$$b_i \equiv \frac{\partial z_{Di}}{\partial p} = \frac{\kappa_i d_i}{\kappa_i + d_i}$$

Equations (7) and (8) are two sets of equations linking agents' bid slopes b_i to the slopes of residual supply d_i facing each agent. Du and Zhu (2012), in a model without derivative contracts, show that there is a unique linear ex-post equilibrium for any collection of demand slopes $\kappa_1 \dots \kappa_N$, characterized by a unique collection of bid slopes b_i and residual supply slopes d_i that satisfy definition 2. Their proof immediately extends to my model, leading to the following proposition.

Proposition 2. *(Du and Zhu, 2012) There is a unique linear ex-post equilibrium in the asymmetric model, in which i submits the bid curve:*

$$z_{Di}(p; y_{ci}, y_{di}) = y_{di}\frac{b_i}{\kappa_i} + y_{ci}\frac{b_i}{\sum_{j \neq i} b_j} - (p - \pi) b_i \tag{9}$$

and the unique equilibrium price is

$$p_b - \pi = \frac{1}{\sum_{i=1}^{n} b_i}\left[\sum_{i=1}^{n}\left[\frac{b_i y_{di}}{\kappa_i} + \frac{b_i y_{ci}}{\sum_{j \neq i} b_j}\right]\right] \tag{10}$$

Where,

$$b_i = \frac{B + 2\kappa_i - \sqrt{B^2 + 4\kappa_i^2}}{2} \tag{11}$$

and $B = \sum_{i=1}^{n} b_i$ is the unique positive solution to the equation

$$B = \sum_{i=1}^{n} \frac{2\kappa_i + B - \sqrt{B^2 + 4\kappa_i^2}}{2} \tag{12}$$

I use the exact equilibrium expressions of proposition 2 in my empirical analysis of the LBMA gold price and the CBOE VIX, in sections 6 and 7 respectively. Equilibrium expressions can be solved for analytically in the symmetric case, as I show in appendix A.7. However, the exact equilibria are difficult to interpret in general, since the quantities b_i

and B cannot generally be solved for analytically. For ease of interpretation, I demonstrate the core theoretical intuitions of the model using an analytically simpler approximation to agents' equilibrium bid curves.

Define agent i's *capacity share* s_i as i's demand slope κ_i divided by the sum of all agents' demand slopes, that is:

$$s_i \equiv \frac{\kappa_i}{\sum_{i=1}^{n} \kappa_i} \tag{13}$$

The denominator of expression (13) is the sum of all market participants' demand slopes, $\sum_{i=1}^{n} \kappa_i$; this is the market's slope of aggregate demand, which represents the market's capacity to supply or absorb quantities of the underlying asset at off-equilibrium prices. Agent i's capacity share s_i thus represents i's share of the slope of aggregate demand, or how much i is responsible for the market's aggregate capacity to supply or absorb the asset at off-equilibrium prices. Capacity shares, like standard market shares, are a measure of market competitiveness: capacity shares are higher for agents with high demand slopes, and if there are many market participants with similar demand slopes, the capacity share of any given participant will be low.[47]

Define the largest capacity share across agents, s_{max}, as:

$$s_{max} \equiv \max_i s_i$$

The following proposition describes approximate expressions for agents' equilibrium bid curves; the approximation is valid when the underlying market is fairly competitive and s_{max} is small.[48]

Proposition 3. *If s_{max} is small, in the unique ex-post equilibrium, agent i bids approximately:*

$$z_{Di}(p; y_{di}, y_{ci}) \approx y_{di} + s_i y_{ci} - \kappa_i (p - \pi) \tag{14}$$

To interpret proposition 3, first suppose that agents are symmetric, so that there are n agents with identical slopes of demand, $\kappa_i = \kappa$. In equilibrium, the slope of residual supply d_i facing any given agent i is approximately the sum over all $n-1$ other agents' slopes of demands; that is, $d_i \approx (n-1)\kappa$. Thus, the manipulation coefficient $\frac{\kappa}{\kappa + d_i}$ of any

[47] Capacity shares can in principle differ significantly from standard market shares based on trade volume. Certain market participants may not regularly trade large volumes of the underlying asset, but may be willing to supply or absorb large quantities of the asset at off-equilibrium prices. In such settings, trade volume could be quite concentrated, but market participants' capacity shares may still be low.

[48] Appendix A.3 proves proposition 3 and bounds the error of the approximation.

given agent is approximately $\frac{1}{n}$: one additional contract held by agent i causes her to trade approximately $\frac{1}{n}$ additional units of the underlying asset. Manipulation coefficients are thus small when the underlying market is competitive and n is large. A simple intuition for this result is that, in order to raise benchmark prices by a given amount, agent i must purchase enough of the underlying asset to raise all other agents' marginal utility for the asset by the same amount; in competitive markets, this is very costly, because it substantially decreases i's own marginal utility for the underlying asset.

If agents are asymmetric, the slope of residual supply facing agent i in equilibrium is approximately the sum over all other agents' slopes of demand, $d_i = \sum_{j \neq i} \kappa_i$, implying that i's manipulation coefficient $\frac{\kappa_i}{\kappa_i + d_i}$ is approximately equal to i's capacity share, s_i. Qualitatively, this result implies that large contract positions have small effects on agents' equilibrium trading decisions, as long as markets are competitive and all agents' capacity shares are small. This result is surprisingly quantitatively precise: manipulation coefficients are equal to agents' capacity shares, and do not depend on any other parameters of the model. Thus, in any setting where we can estimate market participants' demand slopes, we can construct capacity shares and measure agents' equilibrium manipulation incentives.

Manipulation coefficients depend only on agents' capacity shares; manipulation coefficients are not affected by factors which shift the absolute level of agents' demand slopes κ_i. Intuitively, any factor which incerases all agents' demand slopes proportionally has two exactly offsetting effects on manipulation incentives. If agent i's demand slope κ_i is larger, i has a lower cost of purchasing the underlying asset to to push benchmark prices upwards. However, if all other agents' demand slopes are also larger, the slope of residual supply facing i becomes steeper, so i must buy more of the underlying asset to move benchmark prices a given amount. Thus, the level of agents' demand slopes does not affect manipulation incentives; factors which affect all agents' slopes of demand, such as the storability of the underlying asset or the costs for transporting the asset to or from other markets, only affect manipulability insofar as they also affect agents' capacity shares.

3.6 Manipulation-induced benchmark distortions

In this subsection, I consider the distortions induced by manipulation to benchmark prices. Throughout this subsection and subsection 3.7 below, for notational simplicity, I treat the

approximation of proposition 3 as exact. Setting the sum of agents' bid curves to 0 and solving for p_b, we have:

$$p_b - \pi = \frac{\sum_{i=1}^{n} y_{di} + s_i y_{ci}}{\sum_{i=1}^{n} \kappa_i} \tag{15}$$

Expression (15) shows how contract positions and demand shocks affect benchmark prices in equilibrium. Taking expectations over y_{di}, y_{ci}, and using the assumption made in subsection 3.1 that $\sum_{i=1}^{n} E[y_{di}] = 0$, we have:

$$E[p_b - \pi] = \frac{\sum_{i=1}^{n} s_i E[y_{ci}]}{\sum_{i=1}^{n} \kappa_i} \tag{16}$$

Since agents' manipulation incentives are approximately equal to their capacity shares, the net effect of manipulation on benchmark prices depends on the s_i-weighted sum of expected contract positions across agents. Thus, if contract positions are correlated with agents' capacity shares, benchmark prices will be biased towards the contract positions of agents with large capacity shares, even if expected contract positions sum to 0 across agents.

Intuitively, suppose that oil pipeline operators tend to hold short positions in oil contracts, so $y_{ci} < 0$, and retail traders tend to hold long positions, so $y_{ci} > 0$. Since oil pipeline operators have higher demand slopes than retail traders, pipeline operators will sell more oil into the market, per unit contract that they hold, than retail traders are willing to buy, so manipulation will cause benchmark prices to be biased downwards.[49]

If market participants enter into contract settlement with *ex ante* uncertain contract positions, manipulation will also add variance to benchmark prices, making contracts less effective for hedging against uncertainty in economic fundamentals. From expression (15), the variance of benchmark prices is:[50]

$$Var[p_b - \pi] = \frac{\sum_{i=1}^{n} Var(y_{di}) + s_i^2 Var(y_{ci})}{\left(\sum_{i=1}^{n} \kappa_i\right)^2} \tag{17}$$

[49] As I discuss briefly in Subsection 3.2 above, benchmark bias may not necessarily decrease the effectiveness of contracts for hedging against fundamental uncertainty; if bias at settlement is foreseen by all market participants, the ex-ante contract price will change to reflect expected bias at settlement. Manipulation will still distort allocative efficiency, as manipulators need to trade the underlying asset to move prices, but systematic bias in benchmarks need not harm the effectiveness of contracts for hedging.

[50] If demand shocks and contract positions are not independent, expression (17) will contain an additional covariance term, $\sum_{i=1}^{n} s_i \sqrt{Var(y_{ci}) Var(y_{di})}$. I assume demand shocks and contract positions are independent throughout the paper for analytical simplicity, and also because the data I use in sections 6 and 7 do not allow me to identify these covariances.

The contribution of manipulation to benchmark variance is thus:

$$\frac{\sum_{i=1}^{n} s_i^2 Var(y_{ci})}{(\sum_{i=1}^{n} \kappa_i)^2} \tag{18}$$

The components of manipulation-induced benchmark bias and variance, expressions (16) and (18) respectively, can be measured in many contract markets. Contract positions y_{ci} are often tracked closely by regulators, so the expectations and variances of agents' contract positions can be estimated. If regulators are considering policy interventions that change the size of agents' contract positions, regulators can then plug counterfactual values for $E[y_{ci}]$ and $Var(y_{ci})$ into expressions (16) and (18), to predict the effects of these interventions on manipulation-induced benchmark bias and variance. The empirical exercises of sections 6 and 7 will calibrate contract position variances using data on, respectively, gold futures position limits, and the size of VIX futures open interest at contract expiration.

Besides contract position moments, expressions (16) and (18) contain agents' capacity shares s_i and the slope of aggregate demand, $\sum_{i=1}^{n} \kappa_i$. Both are functions of individual agents' demand slopes, κ_i; if the econometrician observes rich enough data to measure individual agents' slopes of demand, she can estimate expressions (16) and (18) given the means and variances of contract positions. In section 6, I measure agents' demand slopes for gold using auction bidding data.

If the econometrician cannot directly estimate agents' demand slopes, capacity shares s_i and the slope of aggregate demand $\sum_{i=1}^{n} \kappa_i$ can be estimated separately. There are many standard methods, such as instrumental variables regressions or price impact regressions, for estimating the slope of aggregate market demand. In the empirical application to VIX in section 7, I estimate the slope of aggregate demand for S&P 500 options using limit order book data. Capacity shares s_i are harder to estimate; however, if the econometrician observes some proxy variable x_i, and is willing to assume that x_i is proportional to agents' demand slope κ_i, capacity shares can be estimated using x_i in place of κ_i:

$$s_i \approx \frac{x_i}{\sum_{i=1}^{n} x_i} \tag{19}$$

For example, x_i could be equal to contract position size, total trade volume, or some other metric external to the model, such as total portfolio size or market cap of agents. Since s_i is a concentration metric, it lies between 0 and 1 by construction; even if the assumption

that x_i is proportional to κ_i is violated, the approximation error in s_i from expression (19) may be quantitatively small. Intuitively, if an econometrician or regulator is confident that a market has roughly n similarly sized participants, participants' capacity shares are unlikely to be orders of magnitude away from $\frac{1}{n}$.[51]

Expressions (16) and (18) thus allow the econometrician to predict the size of manipulation-induced bias and variance in benchmark prices, given data on agents' demand slopes and contract positions. What constitutes an acceptable level of benchmark bias and variance will vary depending on the specific market in question. One criterion regulators could use, which I apply in the empirical exercises of sections 6 and 7, is that manipulation-induced distortions should be small relative to the variance of innovations in benchmark prices between contract settlement events.

3.7 The manipulation index

Another possible measure of manipulation-induced benchmark distortions is the fraction of total benchmark variance which is attributable to manipulation. From expression (17), the ratio of manipulation-induced variance to total benchmark variance is:

$$\frac{\sum_{i=1}^{n} s_i^2 Var\left(y_{ci}\right)}{\sum_{i=1}^{n} Var\left(y_{di}\right) + \sum_{i=1}^{n} s_i^2 Var\left(y_{ci}\right)} \tag{20}$$

Demand shock variance can be thought of as variance which is intrinsic to the benchmark, caused by idiosyncratic variance in participants' desire to buy and sell the underlying asset at the time of benchmark settlement; benchmark variance from demand shocks would exist even if the benchmark were not used for contract settlement. Thus, if expression (20) is small and manipulation variance is much smaller than demand shock-induced variance, benchmark variance is relatively close to what it would be in a world with no manipulation.

The ratio in expression (20) cannot be estimated directly, as it depends on the unobserved variance of demand shocks, σ_{di}^2. However, since agents with larger demand shocks trade more of the underlying asset, the size of demand shocks should be reflected in the

[51] One issue with using proxy variables to estimate capacity shares is that there may be agents who are not usually active in a given underlying market, but are willing to buy or sell large quantities of the underlying asset if prices move significantly. If the econometrician estimates capacity shares s_i using proxy variables, such as trade or contract volume, which consider only participants who are usually active in underlying markets, the econometrician will tend to overestimate capacity shares and thus manipulation incentives.

total volume of trade involved in benchmark calculation. Define v_i as i's unsigned volume of trade in the underlying asset in the benchmark-setting auction, that is,

$$v_i = |z_{Di}(p_b; y_{ci}, y_{di})|$$

Under two additional assumptions about the distributions of demand shocks y_{di} and contract positions y_{ci}, we will be able to approximate expression (20) using aggregate data on contract volume and underlying trade volume used in benchmark calculation.

Assumption 1. *Demand shocks y_{di} and contract positions y_{ci} are normally distributed:*

$$y_{di} \sim N\left(0, \sigma_{di}^2\right), \; y_{ci} \sim N\left(0, \sigma_{ci}^2\right)$$

The assumption that demand shocks and contract positions are normal is convenient, as it allows us to analytically characterize expected trade volume and the expected volume of outstanding contracts. The assumption that both y_{di} and y_{ci} have mean 0 is more substantive. Technically, assumption 1 is needed so that we can infer the variance of contract positions and demand shocks from the total contract and underlying trade volume. Assuming that contract positions are mean 0 rules out the possibility that market participants hold systematically long or short contract positions, and assuming that demand shocks are mean 0 implies that market participants are equally likely to be buyers or sellers of the underlying asset. Both assumptions are most appropriate in settings where market participants act like dealers, who may buy or sell in underlying markets depending on their current inventory conditions. This is potentially an accurate description of some underlying markets for financial assets, such as equities and foreign exchange markets; it is potentially less appropriate for markets for physical commodities, such as oil, wheat, or livestock, where markets are more vertically structured, with producers selling to downstream resource consumers.

Assumption 2. *The variance of demand shocks and the variance of contract positions are proportional across agents:*

$$\sigma_{di}^2 \propto \sigma_{ci}^2$$

Assumption 2 can be interpreted as saying that agents are primarily differentiated by size: larger agents both hold larger contract positions and receive proportionally larger demand shocks. Technically, assumption 2 is needed because the variance induced by contract positions and demand shocks depends on the skew of their distributions

across agents; if contract volume is substantially more concentrated than demand shocks, manipulation-induced variance will tend to be higher than demand shock-induced variance.

Under assumptions 1 and 2, together with the approximation of proposition 3, we can prove the following proposition.

Proposition 4. *Suppose assumptions 1 and 2 hold, and* s_{max} *is small. Define* $C = \sum_{i=1}^{n} |y_{ci}|$*, and* $V = \sum_{i=1}^{n} v_i$*. An approximate upper bound for (20) is:*

$$\frac{\sum_{i=1}^{n} s_i^2 Var(y_{ci})}{\sum_{i=1}^{n} Var(y_{di}) + \sum_{i=1}^{n} s_i^2 Var(y_{ci})} \leqslant \left(\frac{E[C]\, s_{max}}{E[V]}\right)^2 \tag{21}$$

We can construct a simple index of contract market manipulability by taking the square root of expression (21).

Definition 3. Suppose the regulator can estimate expected total contract volume $E[C]$, expected total underlying trade volume $E[V]$ used to construct the benchmark, and the capacity share of the largest agent, s_{max}. Define the *manipulation index* as:

$$\frac{E[C]\, s_{max}}{E[V]} \tag{22}$$

Given proposition 4, the square of the manipulation index is an approximate upper bound for the ratio of manipulation variance to total benchmark variance.

A simple nontechnical interpretation for the manipulation index is that it is an upper bound for the fraction of trade in the underlying asset which is attributable to manipulation. Proposition 3 says that each unit contract held by an agent with capacity share s_i generates s_i units of trade in the underlying asset; $E[C]\, s_{max}$ is thus an approximate upper bound for the expectation of total manipulation-driven trade, so the manipulation index is the ratio of this upper bound to total trade volume in the underlying asset. If the assumptions underlying proposition 4 hold, and the manipulation index is relatively low, then manipulation-induced variance is a relatively small fraction of the total variance of p_b around π, so the benchmark is functioning about as well as it would in a world where no contracts were settled based on the benchmark.

The manipulation index relies on a number of assumptions, which are not appropriate for all contract markets. Proposition 4 builds on the approximation formula of proposition 3, which is appropriate when the underlying market is competitive and s_{max} is relatively

small; proposition 4 also requires assumptions 1 and 2, which are most appropriate when market participants behave like dealers who are differentiated primarily by one-dimensional size. The manipulation index is also a relatively crude upper bound for the share of total benchmark variance induced by manipulation, as it essentially assumes that all contract holders have manipulation coefficients equal to s_{max}, the share of the largest agent.[52] Moreover, the manipulation index can be large without necessarily implying that contract markets are not functioning well. In some markets, the volume of trade in underlying assets used for benchmark settlement may generally be low, but there may be many agents who are willing to trade large volumes of the underlying asset at off-equilibrium prices; in these markets, manipulation-induced benchmark distortions may be small in absolute terms, even if manipulation-induced variance is large relative to demand shock-induced variance.

The manipulation index can potentially play a role similar to the Herfindahl-Hirschman Index (HHI) in antitrust regulation. The HHI is only a formally correct measure of market power-induced distortions under fairly stringent theoretical assumptions. However, it is a simple and intuitive metric which can be calculated using very limited data, so regulators use the HHI as a first-pass metric to identify industries which are concentrated enough that distortions from market power may be large. If an industry is identified as potentially problematic, better data can be collected and more precise methods, such as demand system estimation, can be applied to more accurately model market outcomes and guide policy interventions.[53] Similarly, the manipulation index is an imperfect metric in many ways, but it is easy to interpret and can be calculated using very limited data. Regulators can similarly use the manipulation index as a diagnostic metric, to identify

[52] If large agents by demand slopes in fact hold relatively small contract positions, contract market manipulability could be much lower than the manipulation index suggests. This also implies that, when calculating the manipulability index, total contract volume C only needs to count the volume of contracts held by underlying market participants, as contracts held by agents who do not trade the underlying asset have no influence on benchmarks. This can dramatically change the size of total contract volume C. For example, the volume of outstanding interest rate derivative contracts is approximately 3 orders of magnitude larger than the volume of trade used to determine the SOFR interest rate benchmark; however, most holders of interest rate derivatives do not participate in the treasury-backed overnight repo markets which are used to calculate SOFR, and it is possible that the manipulation index is relatively small when considering only contract positions held by agents who participate in treasury-backed repo markets.

[53] Tirole (1988, pp. 221-223) shows how the HHI can be derived as a measure of industry profitability under Cournot competition with constant marginal costs. The Department of Justice Horizontal Merger Guidelines classifies markets as unconcentrated, moderately concentrated, and highly concentrated based on the HHI. Mergers which would increase industry HHI by more than a certain amount warrant further scrutiny. Thus, the Horizontal Merger Guidelines essentially suggests using the HHI as a diagnostic metric, to be supplemented with more detailed industry-specific analysis.

contract markets which may be vulnerable to manipulation; if a market is identified as potentially problematic, better data can be collected and more precise methods, such as those described in subsection 3.6, can be used to estimate manipulation-induced benchmark distortions more accurately.

4 Aggregated benchmarks

Many price benchmarks aggregate over prices from multiple submarkets. For example, commodities such as oil may be differentiated into multiple heterogeneous types or quality levels, and may be traded in different locations; it may be desirable to construct benchmarks which aggregate across prices of different varieties of oil traded at different locations. In this section, I analyze the optimal design of aggregated benchmarks. The model in this subsection is intentionally stylized, relying on the equilibrium approximation of proposition 3, in order to demonstrate the main intuitions of aggregated benchmarks. I discuss the restrictions I impose in subsection 4.3, and I use exact equilibrium expressions and generalize the model in various directions in the applications of sections 6 and 7.

4.1 Model

There are $m \in \{1 \dots M\}$ submarkets, each with $n_m \geqslant 3$ traders, indexed by i. Trader i in submarket m has slope of demand κ_{mi}, and demand shocks $y_{mdi} \sim N\left(0, \sigma_{dm}^2\right)$. Traders $i \in \{2 \dots n_m\}$ are regular traders, who have no contract positions. The utility of a regular trader in submarket m for trading z_m at price p_m is:

$$U_{reg}\left(p_m, z_m\right) = \pi z_m + \frac{y_{mdi} z_m}{\kappa_{mi}} - \frac{z_m^2}{2\kappa_{mi}} - p_m z_m$$

There is a single manipulator, who we call agent $i = 1$ in each submarket. In addition to demand shocks, she has a contract position $y_c \sim N\left(0, \sigma_c^2\right)$. The contract position payoff depends on an aggregated benchmark, P, which is an affine function of submarket prices, as I describe below. The manipulator has slope of demand κ_{m1} in each submarket m; her utility is the sum of utilities across all submarkets, plus her contract payoffs. That is, if the manipulator trades z_m in each submarket at p_m, and the aggregated benchmark is P, her

payoff is:

$$U_{manip}\left(p_1 \dots p_M, z_1 \dots z_M\right) = \left[\sum_{m=1}^{M} \pi z_m + \frac{y_{md1} z_m}{\kappa_{m1}} - \frac{z_m^2}{2\kappa_{m1}} - p_m z_m\right] + y_c P$$

Thus, the manipulator optimally bids across all submarkets in order to maximize the sum of her physical utilities and contract payoffs.

Unlike section 3, in this section I will assume that $\pi \sim N\left(\mu_\pi, \sigma_\pi^2\right)$. By assuming a prior for π, I take a Bayesian view of benchmarks; benchmark prices will be averages of observed prices and the posterior mean. This implies that risk-minimizing benchmarks will be biased towards the historical mean.[54]

I consider a class of benchmarks of the form:[55]

$$P_b = \mu_\pi + \sum_{m=1}^{M} \phi_m \left(p_m - \mu_\pi\right) \tag{23}$$

As I show below, the mean of prices in each submarket m in equilibrium is equal to π; prices differ from π because of demand shocks specific to each submarket, and the manipulator acting across submarkets. The goal of the benchmark administrator is to choose weights ϕ_m to minimize the squared difference between P_b and π.

4.2 Equilibrium and price variance

If the manipulator trades quantities z_m in each submarket, and the auction clearing prices are p_m, the manipulator's utility is:

$$\left[\sum_{m=1}^{M} \pi z_m + \frac{y_{md1} z_m}{\kappa_{m1}} - \frac{z_m^2}{2\kappa_{m1}} - p_m z_m\right] + y_c \left[C + \sum_{m=1}^{M} \phi_m p_m\right]$$

Or,

$$\sum_{m=1}^{M} \left[\pi z_m + \frac{y_{md1} z_m}{\kappa_{m1}} - \frac{z_m^2}{2\kappa_{m1}} - p_m z_m + p_m \phi_m y_c\right] + C y_c$$

[54] In settings where it is not appropriate to assume a prior for π, all results can be applied taking the limit as $\sigma_\pi^2 \to \infty$, so the prior is fully diffuse.

[55] Appendix B.2 shows that any general affine benchmark $P = C + \sum_{m=1}^{M} \phi_m p_m$ is dominated by a benchmark that takes the form (23).

The outcomes z_m and p_m each only influence a single component of the sum. Thus, the manipulator's optimal bidding behavior is separable across submarkets; she bids as if she holds $\phi_m y_c$ contracts in submarket m. Thus, agents bid as in the unique equilibrium of proposition 2 in section 3, except that the manipulator bids as if her contract position in each market is $\phi_m y_c$.

Proposition 5. *Fixing ϕ_m, there is a unique equilibrium of the aggregated benchmark game, in which regular traders in submarket m bid approximately:*

$$z_{Dmi}(p_m; y_{dmi}) = y_{dmi} - \kappa_{mi}(p_m - \pi) \tag{24}$$

The manipulator bids, in each submarket m:

$$z_{Dm1}(p_m; y_c) = y_{dmi} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m} \kappa_{m1}} \phi_m y_c - \kappa_{m1}(p_m - \pi) \tag{25}$$

The equilibrium price in submarket m is:

$$p_m - \pi = \frac{1}{\sum_{i=1}^{n_m} \kappa_{mi}} \left[\sum_{i=1}^{n_m} y_{dmi} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m} \kappa_{m1}} \phi_m y_c \right] \tag{26}$$

For notational simplicity, we will define the aggregate demand shock $\bar{y}_{dm}$ in submarket m as:

$$\bar{y}_{dm} = \sum_{i=1}^{n_m} y_{dmi}$$

$\bar{y}_{dm}$ is distributed as $N\left(0, \bar{\sigma}^2_{dm}\right)$, where:

$$\bar{\sigma}^2_{dm} = \sum_{i=1}^{n} \sigma^2_{dm}$$

Then, we can write expression (26) as:

$$p_m - \pi = \frac{1}{\sum_{i=1}^{n_m} \kappa_{mi}} \left[\bar{y}_{dm} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m} \kappa_{m1}} \phi_m y_c \right] \tag{27}$$

Appendix B.2 shows that expression (27) implies:

$$\mathrm{E}\left[(P_b-\pi)^2\right] = \underbrace{\left(1-\sum_{m=1}^{M}\phi_m\right)^2\sigma_\pi^2}_{\text{Bias}} + \underbrace{\sum_{m=1}^{M}\phi_m^2\frac{\bar{\sigma}_d^2}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}}_{\text{Demand Shocks}} + \underbrace{\left[\sum_{m=1}^{M}\phi_m^2\frac{\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\right]^2\sigma_c^2}_{\text{Manipulation}} \quad (28)$$

The "Bias" term reflects the fact that, if weights ϕ_m do not sum to 1, the benchmark is biased towards the prior.[56] If submarket prices are sufficiently noisy signals of underlying values π, estimators minimizing the distance between the benchmark P_b and the true state of the world π optimally shrink the estimate towards the prior mean of π. The "Demand Shocks" term aggregates variance from demand shocks over submarkets; since demand shocks are independent across submarkets, benchmark variance can be decreased by taking a weighted average over submarket prices. The "Manipulation" term reflects the manipulator's correlated attempts across markets to move prices to influence contract payoffs. Manipulators preferentially trade in markets where n_m and κ_m are low, so the cost of moving p_m is low, and where weights ϕ_m is high, so that movements in p_m affect the aggregate benchmark P significantly. Hence, manipulation variance is decreased when the market administrator puts higher weight on submarkets which are more competitive, and in which agents have higher slopes of demand κ_m, both of which increase the cost of price impact for any individual agent.

The optimal benchmark weights ϕ_m minimize the sum of variance from both sources. Formally, the benchmark administrator solves:

$$\min_{\phi_m}\left(1-\sum_{m=1}^{M}\phi_m\right)^2\sigma_\pi^2 + \sum_{m=1}^{M}\phi_m^2\frac{\bar{\sigma}_d^2}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2} + \left[\sum_{m=1}^{M}\phi_m^2\frac{\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\right]^2\sigma_c^2 \quad (29)$$

The objective function is strictly convex in ϕ_m, so the optimal weight vector is unique, and since $\frac{dVar(P)}{d\phi_m}$ has the same sign as ϕ_m, the problem is guaranteed to have an interior solution. Hence, the FOC is necessary and sufficient for a set of weights ϕ_m to constitute

[56] When the prior is fully diffuse, so that $\sigma_\pi^2 \to \infty$ the bias is infinite whenever $\sum_{m=1}^{M}\phi_m \neq 1$, so we will have always have $\sum_{m=1}^{M}\phi_m = 1$; the benchmark P will be a weighted average of submarket prices, and will be an unbiased estimator of π. The benchmark administrator's problem then becomes to minimize the variance of P around π, by choosing weights ϕ_m which sum to 1. Hence the "frequentist" case can be derived as a limit of the Bayesian case with a fully diffuse prior.

an optimal benchmark.

Proposition 6. *The solution to the aggregated benchmark problem is the unique set of weights* ϕ_m *that satisfies, for all* m*:*

$$-\left(1-\sum_{m=1}^{M}\phi_m\right)\sigma_\pi^2+\frac{\phi_m\bar{\sigma}_d^2}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}+\left[\sum_{m=1}^{M}\phi_m^2\frac{\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\right]\frac{2\phi_m\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\sigma_c^2 \;=\; 0 \quad (30)$$

Rearranging expression (30), we get:

$$\phi_m=\frac{\left(1-\sum_m\phi_m\right)\sigma_\pi^2}{\frac{\bar{\sigma}_d^2}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}+\left[\sum_m\phi_m^2\frac{\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\right]\frac{2\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\sigma_c^2}$$

Or,

$$\phi_m\propto\frac{1}{\frac{\bar{\sigma}_d^2}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}+\left[\sum_m\phi_m^2\frac{\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\right]\frac{2\kappa_{m1}}{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}\sigma_c^2} \quad (31)$$

If we focus only on the manipulation component of variance, we have, from expression (31):

$$\phi_m\propto\frac{\left(\sum_{i=1}^{n_m}\kappa_{mi}\right)^2}{\kappa_{m1}}$$

We can instead write this as:

$$\phi_m\propto\underbrace{\left(\sum_{i=1}^{n_m}\kappa_m\right)}_{\text{Liquidity}}\underbrace{\frac{\sum_{i=1}^{n_m}\kappa_m}{\kappa_{m1}}}_{\text{Competition}} \quad (32)$$

Thus, ϕ_m assigns higher weights to markets which are more liquid, meaning that the slope of aggregate demand $\sum_{i=1}^{n_m}\kappa_m$ is higher, and more competitive, meaning that the slope of aggregate demand is large relative to the manipulator's slope of demand κ_{m1}.

4.3 Discussion

The model in this section is purposefully stylized in order to illustrate the main intuitions behind the design of aggregated benchmarks. The model can easily be extended to handle multiple manipulators, aggregation functions other than the mean, and correlations

between demand shocks in submarkets. One substantive problem which is more difficult to handle is that this section assumes that auctions in different submarkets are separate, and the outcome in one submarket does not affect agents' utilities in other submarkets. This is unlikely to hold in practice; many large agents are active in submarkets for products of different varieties, locations and dates, and products are likely to be substitutes, so that purchasing quantity in one submarket increases marginal utility in other submarkets. Intuitively, such dependencies should have two opposing effects. First, it is easier for manipulators to move prices, as purchasing large quantities in one submarket will also increase prices in other submarkets. Secondly, manipulation is more costly, as purchasing quantity in a given submarket $\mathfrak{m}$ decreases the manipulator's marginal utility in other submarkets $\mathfrak{m}'$; however, if the manipulator decreases her bids in other submarkets $\mathfrak{m}'$, this will lower aggregate benchmark prices, counteracting her manipulative actions in submarket $\mathfrak{m}$. However, modelling dependencies between submarkets is relatively difficult,[57] and I leave this analysis to future work.

The model of this section is most appropriate for settings in which multiple varieties of a commodity such as oil or gas may be traded in multiple locations. If a benchmark administrator wishes to construct a benchmark representing the world price of oil, she has a choice of how to weight submarkets at different locations, involving different varieties of oil. At present, most aggregated benchmarks are calculated by weighting submarket prices by the volume of trade in a given submarket. This section argues that the benchmark administrator can do better by weighting by *liquidity* and *competitiveness*. Certain submarkets may be consistently liquid and competitive, for example if they are closer to large sources of the commodity, or if multiple large market participants compete. Even if the volume of trade in these submarkets happens to be low in a given period of time, prices in these submarkets may still be informative and robust to manipulation. In many markets, metrics which are correlated with liquidity and competitiveness are observable – for example, regulators often observe the number of large participants in markets, or liquidity measures such as bid-ask spreads. The results of this section suggest that these measured could potentially be used, in place of or in addition to volume, in calculating aggregated benchmarks.

[57] Some recent papers analyzing markets with multiple assets and dependent utility functions are Malamud and Rostek (2017), Wittwer (2018) and Wittwer (2017).

5 The CFTC Commitments of Traders reports

In this section, I analyze the Commodity Futures Trading Commission (CFTC) Commitments of Traders (COT) reports. These data are very limited – they only cover contract markets, with no information about underlying markets, and they track the distribution of contract positions across broad categories of market participants, rather than individual market participants. However, the data gives us a broad overview of how competitive contract markets are, and what the positions of different classes of participants look like in different markets. I demonstrate two stylized facts. First, concentration in contract open interest varies substantially across contract markets. Second, certain classes of market participants are systematically long or short in many contract markets, suggesting that capacity shares and contract positions are correlated, and thus benchmark prices are likely to be biased.

5.1 Data

Contract markets are highly regulated and monitored. Under the Commodity Exchange Act, all commodity futures and options trading must occur on organized exchanges which are registered with the CFTC. The CFTC has authority to set limits on the size of agents' contract positions, and also to track total contract open interest, and the identities and positions of large traders in contract markets. Each Tuesday, the CFTC publishes the Commitments of Traders reports, which contains summary statistics on open interest in different futures and options markets.

The COT reports provide two kinds of information which are relevant for this paper. First, for each contract, the COT reports show the share of gross long and short open interest held by the 4 largest market participants. If we assume that contract position sizes are proportional to demand slopes, then these data allow us to approximately infer the size of the capacity shares of large contract market participants. Second, the reports show the distribution of long and short open interest among different classes of market participants. Different classes of traders likely have different κ_i values: producers and merchants who physically hold and trade the underlying asset likely have low costs of buying or selling the underlying asset, compared to financial speculators who do not usually trade the underlying asset. Thus, price benchmarks at settlement are likely to be biased towards the net contract positions of producers and merchants. The COT reports

thus allow me to predict the likely direction of benchmark price biases in different contract markets.

I analyze long-form Commitments of Traders report data on commodity futures (excluding options) from 2010 to the end of 2016. Details about the COT reports, my data cleaning steps, and definitions of each of the groups of traders can be found in appendix C.1. Results are reported in table 1, and appendix C.2 reports additional results for financial futures.[58]

5.2 Contract market competitiveness

The "Top 4" columns in table 1 show the shares of long and short contract open interest held by the 4 largest market participants in each market. Contract market competitiveness varies substantially across contract markets. The top 4 share is in the range of 20-30% for agricultural products, compared to around 30-60% for energy products; concentration is also high for dairy products and emissions permits.

What matters for contract market manipulability, however, is agents' capacity shares, not their shares of contract open interest. In particular, the manipulation index of definition 3 is based on s_{max}, the capacity share of the largest agent. If we wish to approximately infer s_{max} from top 4 long and short contract shares, there are three issues we need to resolve. First, the reported shares aggregate over multiple delivery months; for the purposes of measuring manipulability, we are interested in concentration in the underlying market during settlement for a contract at a given delivery month. Second, the reported shares represent concentration in contract positions, not concentration in demand slopes. Third, we are interested in the largest agent's capacity share, not the share of the top 4 agents. Since I only observe top 4 shares in the COT reports, I essentially need to assume all three issues away.

Formally, I proceed as follows. First, we need to deal with multiple delivery months. Suppose that, for a given contract, there are T expiration months indexed by t. I assume

[58] A caveat to my analysis is that many of these contract markets are settled by physical delivery of the underlying asset, rather than cash-settled based on price benchmarks; my analysis does not formally apply to physical delivery contract markets. However, physical delivery and cash-settled contract markets largely have the same economic function; the choice of settlement mechanism depends on details such as the cost for delivering the underlying asset – see, for example, Jones (1982), Garbade and Silber (1983), and Paul (1985). The purpose of this section is to show that competition, along with correlations between contract positions and demand slopes, varies across contract markets in practice, and it seems unlikely that these factors behave systematically very differently for physical delivery and cash-settled contracts.

there is a fixed set of long agents $i = 1 \dots n_l$, holding contract positions $y_{tci} > 0$ in each month t, and a fixed set of short agents $j = 1 \dots n_s$ holding contract positions $y_{tcj} < 0$ in each month t. I assume agents' contract positions are proportional across dates.

Assumption 3. *For any agent* i *and any* t, t'*, we have* $y_{tci} \propto y_{t'ci}$*; similarly, for any* j, t, t'*, we have* $y_{tcj} \propto y_{t'cj}$

Thus, any agent who holds a long contract position in one period holds long contract positions in all periods, and likewise for shorts. If we arrange agents in decreasing order of contract position size, the observed aggregate long and short top 4 shares q_{4l}, q_{4s} can be written as:

$$q_{4l} \equiv \frac{\sum_{t=1}^{T} \sum_{i=1}^{4} y_{tci}}{\sum_{t=1}^{T} \sum_{i=1}^{n_{tl}} y_{tci}}, \; q_{4s} \equiv \frac{\sum_{t=1}^{T} \sum_{j=1}^{4} y_{tcj}}{\sum_{t=1}^{T} \sum_{j=1}^{n_{ts}} y_{tcj}} \tag{33}$$

Under assumption 3, q_{4l} and q_{4s} are equal to top 4 shares in any given expiration date t; that is, for any t,

$$q_{4l} = \frac{\sum_{i=1}^{4} y_{tci}}{\sum_{i=1}^{n_{tl}} y_{tci}}, \; q_{4s} = \frac{\sum_{i=1}^{4} y_{tci}}{\sum_{i=1}^{n_{tl}} y_{tci}} \tag{34}$$

Second, we need to infer capacity shares from contract market concentration. I assume that contract positions and demand slopes are proportional.

Assumption 4. *Contract positions are proportional to demand slopes, that is,*

$$|y_{tci}| \propto \kappa_i$$

Under assumptions 3 and 4, top 4 long and short shares q_{4l}, q_{4s} accurately reflect the sum of capacity shares of the 4 largest long and short contract holders on any given date. That is, expression (34) and assumption 4 imply that:

$$q_{4l} = \frac{\sum_{i=1}^{4} \kappa_i}{\sum_{i=1}^{n_l} \kappa_i}, \; q_{4s} = \frac{\sum_{j=1}^{4} \kappa_j}{\sum_{j=1}^{n_s} \kappa_j}$$

Long and short positions must add to 0 across all contract market participants, so $\sum_{i=1}^{n_l} y_{tci} = \sum_{j=1}^{n_{ts}} y_{tcj}$. This implies that $\sum_{i=1}^{n_l} \kappa_i = \sum_{j=1}^{n_{ts}} \kappa_j$, hence,

$$\frac{q_{4l} + q_{4s}}{2} = \frac{\sum_{i=1}^{4} \kappa_i + \sum_{j=1}^{4} \kappa_j}{\sum_{i=1}^{n_l} \kappa_i + \sum_{j=1}^{n_{ts}} \kappa_j} \tag{35}$$

The denominator of the RHS of expression (35), $\sum_{i=1}^{n_l} \kappa_i + \sum_{j=1}^{n_{ts}} \kappa_j$, is the sum of demand slopes of all agents in the market, and the numerator is the sum of slopes of the 4 largest shorts and longs. In words, expression (35) says that the average of q_{4l} and q_{4s} is approximately the total capacity share of 8 agents: the 4 largest short contract holders, and the 4 largest longs.

Finally, we need to infer the share s_{max} of the single largest agent from the share of the 4 largest shorts and longs. I simply assume that s_{max} is half of the average of q_{4l} and q_{4s}.

Assumption 5. *The share of the largest agent is equal to half the share of the 4 largest shorts and longs, that is,*

$$s_{max} = \frac{\max_i \kappa_i}{\sum_{i=1}^{n_l} \kappa_i + \sum_{j=1}^{n_{ts}} \kappa_j} = \frac{1}{2} \frac{\sum_{i=1}^{4} \kappa_i + \sum_{j=1}^{4} \kappa_j}{\sum_{i=1}^{n_l} \kappa_i + \sum_{j=1}^{n_{ts}} \kappa_j} = \frac{q_{4l} + q_{4s}}{4} \tag{36}$$

Appendix C.3 shows that assumption 5 can microfounded by assuming that agents' demand slopes κ_i follow Zipf's law, with an exponent roughly in the range of 1 to 1.5; the demand slope of the largest agent is then roughly 40-50% the sum of demand slopes of the 8 largest agents.

Table 1 shows the inferred values of s_{max}, based on expression (36); s_{max} is estimated to be in the range of 8%-37% across markets. The analysis underlying the manipulation index in subsection 3.7 implies that contract volume can be on the order of $\frac{1}{s_{max}}$ times underlying volume before manipulation begins to substantially distort price benchmarks. For relatively competitive markets, such as those for agricultural commodities, livestock and oil, s_{max} is on the order of $\frac{1}{10}$, so contract volume can be around 10 times underlying volume before manipulation begins to distort benchmarks substantially. For relatively concentrated markets, such as those for dairy products, base and precious metals, electricity, and emissions permits, s_{max} is on the order of $\frac{1}{5}$ to $\frac{1}{3}$, so manipulation could substantially distort benchmarks even if contract volume is not much larger than underlying volume.

5.3 Benchmark price bias

Table 1 also shows the distribution of long and short open interest across four different classes of market participants. In many contract markets – in particular, in most agricultural, livestock and metals products – producers and merchants hold net short contract positions, with short open interest around 20-30% larger than long open interest. Corre-

spondingly, swap dealers and managed money participants are significantly net long in these categories. In the context of my theory, this suggests that price benchmarks are likely to be biased downwards in these markets. As I describe in Appendix C.1, producers and merchants are agents who have commercial uses for the physical underlying commodities: they are, for example, food or energy producers, warehouse or pipeline operators, and other agents who regularly trade the underlying asset. As a result, they are likely to have relatively low costs for trading the underlying asset, and thus high slopes of demand κ_i. In contrast, the swap dealer and managed money categories represent primarily financial participants, who have positions in contract markets but rarely trade the underlying asset. These participants are likely to have high costs of trading the underlying asset, and thus low slopes of demand κ_i.[59]

My theory predicts that the net effect of manipulation on benchmark prices depends on the s_i-weighted sum of contract positions across agents, $\sum_{i=1}^{n} s_i y_{ci}$; if contract positions are correlated with agents' demand slopes κ_i, and thus their capacity shares s_i, then price benchmarks will be biased towards the contract positions of agents with high capacity shares. Thus, price benchmarks will tend to be biased downwards at contract settlements for many of the commodities in the COT data.[60]

5.4 Discussion

My theory predicts that two factors which influence contract market manipulability are competitiveness, and correlations between contract positions and capacity shares. This section uses the CFTC Commitments of Traders reports to show that both factors seem to vary substantially across contract markets in practice.

A somewhat counterintuitive prediction of this paper is that contract market manipulators are more likely to be commercial users of underlying goods, such as food producers

[59]This conclusion would not necessarily hold if financial participants were on average much larger than commercial market participants. However, the data does not appear to support this: short and long open interest are similarly concentrated, and in fact short top 4 shares are somewhat larger than long top 4 shares for a few of the categories in which producers and merchants are substantially net short: corn and wheat, soybean products, cotton, sugar/cocoa/coffee, and previous metals. This suggests that, if anything, commercial market participants are larger and more concentrated than financial participants in these markets.

[60]If demand slopes κ_i are correlated with agents' types, this implies that assumption 4, that agents' contract position sizes are proportional to their demand slopes, is most likely violated. In practice, this likely implies that s_{max} is somewhat higher, so manipulation incentives are somewhat larger, than subsection 5.2 suggests.

or oil pipeline operators, and less likely to be purely financial participants, such as banks or hedge funds. This is because, in my model, trading the underlying asset is the only way to move prices, and commercial market participants likely have lower costs of trading the underlying asset than purely financial participants. There is some qualitative evidence suggesting that this prediction holds in practice; appendix C.4 describes a number of recent cases of contract market manipulation, all of which involve parties that would be characterized as producers and merchants in the COT data. Manipulation cases involving purely financial market participants seem comparatively rare.

In many contract markets, the CFTC sets limits on the size of contract positions that speculators can hold. However, position limits do not apply to agents who have *bona fide* commercial risks to hedge.[61] My theory suggests that commercial market participants are in fact more likely to manipulate than agents who are pure financial participants, implying that current position limit policy may not be optimal for reducing contract market manipulability.[62]

[61] See the CFTC's website on Speculative Limits.

[62] The CFTC's stated purpose for speculative position limits is to "protect futures markets from excessive speculation that can cause unreasonable or unwarranted price fluctuations"; hence position limits do not exist solely to combat manipulation, although according to my theory they can be an effective tool for doing so.

Table 1: Commitments of Traders reports: commodity futures

	Top 4		s_{max}	Prod+Merch		Swap Dealers		Man. Money		Other	
Category	Long	Short		Long	Short	Long	Short	Long	Short	Long	Short
Corn and Wheat	24%	31%	14%	29%	50%	19%	3%	21%	17%	31%	30%
Soybean Products	19%	25%	11%	22%	51%	18%	3%	24%	14%	37%	32%
Cotton	18%	29%	12%	14%	58%	33%	8%	32%	15%	21%	19%
Sugar/Cocoa/Coffee	21%	33%	13%	28%	60%	20%	8%	30%	18%	22%	14%
Livestock	17%	16%	8%	12%	37%	24%	3%	35%	17%	29%	44%
Dairy	52%	50%	25%	61%	48%	3%	3%	2%	2%	33%	47%
Precious Metals	26%	36%	16%	9%	43%	18%	27%	44%	14%	29%	16%
Base Metals	42%	48%	22%	15%	32%	27%	6%	18%	19%	40%	42%
Oil	29%	27%	14%	22%	30%	28%	28%	18%	13%	31%	30%
Natural Gas	34%	37%	18%	35%	23%	27%	43%	20%	17%	18%	17%
Electricity	60%	67%	32%	74%	73%	20%	22%	4%	3%	2%	2%
Lumber	31%	31%	16%	14%	34%	16%	0%	25%	21%	45%	44%
Emissions Permits	72%	75%	37%	47%	30%	10%	40%	7%	0%	35%	29%

Notes. "Top 4" is the fraction of gross long and short open interest held by the 4 largest market participants, defined as q_{l4} and q_{s4} above. s_{max} is the sum of q_{l4} and q_{s4} divided by 4, as in (36). All other columns describe the fraction of long and short open interest held by different classes of market participants. Open interest percentages add to 100%, separately for longs and shorts, across the 4 categories of market participants: "Prod+Merch", "Swap Dealers", "Man. Money", and "Other". I provide descriptions of categories in Appendix C.1. Since the CFTC does not provide a codebook for subgroup codes, I manually label subgroups in the "Description" column.

6 The LBMA Gold Price

In this section, I study the London Bullion Market Association (LBMA) gold price benchmark. Subsection 6.1 introduces the LBMA gold price. Subsection 6.2 discusses the round-level auction data I observe. Subsection 6.3 describes my model for auction participants' utilities, and subsection 6.4 discusses how I map the dynamic model used to set the LBMA gold price to my static auction model. Subsection 6.5 discusses the parametrization and estimation of my model. Subsection 6.6 shows how much benchmark variance would increase if COMEX gold futures contracts were settled based on the LBMA gold price. Proofs, robustness checks, and other details are presented in appendix D.

6.1 Background

The LBMA gold price is an important benchmark price for gold, set twice each business day at 10:30AM and 3:00PM London time. Prior to 2015, the price was set in a private teleconference auction between five members; this took between 10-15 minutes to conclude, and data on auction bidding was not made public. In 2014, the five banks involved in setting the LBMA gold price were accused of manipulation (Reuters Staff, 2014); the UK Financial Conduct Authority fined Barclays $43.9 million for bidding strategically to move benchmark prices, in order to avoid paying USD $3.9m to a customer who held a benchmark-linked option contract FCA (2014). In 2015, the ICE Benchmark Administration (IBA) took over the administration of the LBMA gold price; IBA moved to an electronic auction system, allowing more participants to enter, and began publishing detailed information about bids in intermediate rounds of the auction.

A variety of contracts are settled based on the LBMA gold price. While the popular COMEX gold futures contract (ticker symbol GC) is physically settled, other exchanges have introduced cash-settled gold futures and options contracts tied to the LBMA gold price,[63] and some OTC gold derivatives are settled using the LBMA gold price (FCA, 2014). The LBMA gold price is also used by other market participants, such as miners, central banks, and jewellers, for purposes such as inventory valuation (Aspris et al., 2015).

The LBMA gold price is determined in a multi-round dynamic auction. During each round, IBA publishes a round price, and participants then have 30 seconds to enter how much gold they want to buy or sell at the announced price. If the difference between

[63] Gold Futures and Gold Options: Further Information

buying and selling is within an imbalance threshold – during the time period that my data covers, usually 10,000oz – the auction concludes. Otherwise, IBA adjusts the price in the direction of the imbalance in volume between buy and sell orders, and a new round begins.

6.2 Data

The primary data source I use is the IBA Gold Auction Historical Transparency Reports.[64] The data covers daily morning and afternoon auctions; the full dataset covers 1650 auctions over the period 2015-03-20 to 2018-06-29. Let auctions be indexed by $a \in \{1 \dots A\}$, and suppose that auction a lasts for R_a rounds, indexed by r. For each round r of each auction a, I observe the number of participants, n_{ar}, the round price, p_{ar}, and the total volume of gold that auction participants wish to buy and sell, respectively b_{ar} and b_{ar}.

For estimating my model, I filter to auctions with at least 3 rounds, $R_a \geqslant 3$, as I will estimate slopes of demand by regressing round buy and sell volume on round prices. I also filter to auctions in which the number of participants n_{ar} is constant over the course of the auction. This reduces the estimation sample to 509 auctions. Table 2 shows features of my sample and the full dataset. Most auctions have 6-9 participants and conclude in 3-6 rounds. The range of prices between rounds for any given auction is small, relative to variation in gold prices between auctions: the difference between the highest and lowest round prices is around $1 USD/oz on average. Approximately 168,000oz of gold is traded on average in each auction in my estimation sample.

Define the *volume imbalance* in round r of auction a as the difference between buy and sell volume, that is:

$$i_{ar} \equiv b_{ar} - s_{ar}$$

The auction clearing price, buy and sell volume, and volume imbalance at the final round R_a of auction a are:

$$p_{aR_a}, b_{aR_a}, s_{aR_a}, i_{aR_a}$$

Define the total trade volume at the end of the auction as the sum of buy and sell volume:

$$v_{aR_a} = b_{aR_a} + s_{aR_a}$$

[64] I accessed the data at the ICE website.

For simplicity, when discussing final-round quantities, I will omit the round subscript R_a; thus, I will write p_a, v_a, i_a to mean the final auction price, trade volume, and volume imbalance in auction a. Let n_a represent the number of participants in auction a; this is uniquely defined within my estimation subsample, since I filter to auctions in which participation is constant.

Figure 1 shows a plot of volume imbalance against prices; each data point is an auction round. The x-axis shows $p_{ar} - p_{aR_a}$, the difference between auction prices in round r and in the final round R_a of auction a. The y-axis shows i_{ar}, the difference between buy and sell volume in round r of auction a. There is a clear decreasing trend: when the auction round price is higher, auction participants want to sell more and buy less. The colored lines in the left panel of figure 1 are linear regression lines, separately estimated for auctions with different numbers of participants n_a. Auctions with larger numbers of participants generally have higher slopes of aggregate demand, although the slope when $n_a = 9$ is slightly lower than the slope when $n_a = 8$. Thus, the slope of aggregate auction demand appears to be an increasing but concave function of n_a. This suggests that marginal participants add to the slope of aggregate demand, but later marginal participants contribute less to the demand slope.

The right panel shows total auction volume versus the number of bidders. Unlike the slope of aggregate demand, volume increases approximately linearly in the number of bidders. This suggests that, while later participants may contribute little to the slope of aggregate demand, they have demand shocks of similar size to earlier participants. Appendix D.1 shows that these stylized facts about the relationship between demand slope, auction volume, and participant count are robust to controlling for various other observable features of auctions.

A major weakness of my data is that I only observe the total number of participants, and aggregate buy and sell volume, rather than data on individual participants' bids. One way to deal with this would be to assume that agents are symmetric, and to estimate a single value of demand slope κ and demand shock variance σ_d^2 for all agents. However, the fact that the slope of aggregate demand is a concave function of n suggests that there is some dispersion in agents' demand slopes; agents who participate less often appear to have smaller values of κ_i. Since heterogeneity in agents' κ_i values is an important factor influencing contract market manipulability, I will construct a model that attempts to match the observed heterogeneity in the data, despite not being able to observe participants' identities.

Figure 1: Gold auction demand slope and volume

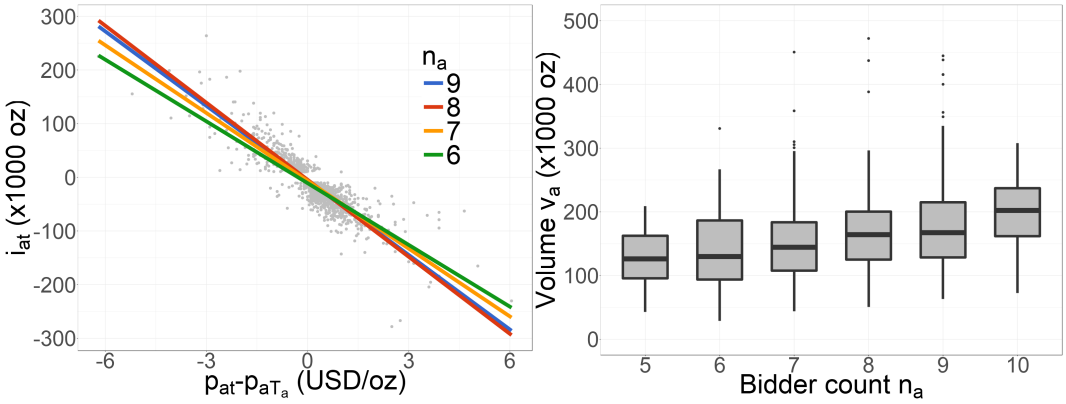

Notes. The left plot shows round volume imbalances and prices, for all non-terminal auction-rounds with $n_a \in \{6, 7, 8, 9\}$ participants. Each data point is an auction-round; final rounds are excluded. The x-axis shows the round price p_{ar} minus the auction clearing price p_{aR_a}, in USD/oz. The y-axis shows the round imbalance i_{ar}, in thousands of ounces of gold. The colored lines are predictions from separate linear regressions for each n_a. There is a narrow band around $i_{ar} = 0$ with no data points because the auction clears when $|i_{ar}|$ is below 10,000oz. The right plot shows boxplots of total auction volume v_a, versus the number of bidders n_a participating in the auction.

Table 2: Gold auction descriptive statistics

	Population mean	Sample mean	SD	P10	P90
Rounds	3.37	4.18	1.52	3	6
Participants	7.38	7.71	1.26	6	9
Price (USD / oz)	1236.98	1223.35	75.11	1118.23	1327.3
Price range (USD / oz)	0.79	1.06	0.8	0.45	1.91
Volume (1000 oz)	148.05	167.79	85.22	90.39	244.78
N	1650	509			

Notes. Each observation is an auction. The "Populationmean " column shows the mean of each variable in the full sample of 1650 auctions, and the "Sample mean" column shows means within my estimation subsample of 509 auctions. The "SD", "P10" and "P90" columns respectively show the standard deviation, 10th percentile, and 90th percentile values of the variable within my estimation subsample. "Rounds" is the number of rounds the auction took to complete, R_a. "Participants" is the number of auction participants, n_a. "Price" is the auction clearing price, p_a. "Price range" is the range of round prices within each auction, that is, the difference between the highest and lowest round prices. "Volume" is v_a, the sum of buy volume and sell volume in the final round of the auction.

6.3 Model

Let $N = \max_a n_a$ be the maximum number of participants across auctions; in my data, $N = 11$. I assume that the observed data is generated by auctions between a fixed set of $N = 11$ participants, indexed by i, with possibly heterogeneous slopes of demand:

$$(\kappa_1 \dots \kappa_N)$$

which are fixed over time. I assume agents receive demand shocks, which are normally distributed with mean 0 and possibly heterogeneous variances:

$$\left(\sigma_{d1}^2 \dots \sigma_{dN}^2\right)$$

where the vector of demand shock variances is also fixed over time.

The assumption that there is a fixed set of agents with stable κ_i is plausible for the LBMA gold price. There is only a small set of agents who are allowed to participate

in LBMA gold auctions,[65] and the relative sizes of participants are probably relatively slow-moving over time.

Suppose auction a has n_a participants; I assume that n_a is exogeneous to auction outcomes.[66] In auction a, the utility of agent i for purchasing z units of gold, when the auction clearing price is p_{ab}, is:

$$U_{ia}(z, p_{ab}) = \pi_a z + \frac{y_{adi} z}{\kappa_i} - \frac{z^2}{2\kappa_i} - p_{ab} z + p_{ab} y_{aci} \tag{37}$$

I assume that demand shocks y_{adi} and contract positions y_{aci} are drawn independently across auctions, and that all agents' utility functions are separable across auctions, so realizations in each auction are independent conditional on agents' demand slopes and demand shock variances. π_a is the mean value of agents for gold in auction a; this varies from auction to auction as market prices of gold change. y_{aci} represents i's contract position; for estimation, I assume that agents do not hold derivative contracts tied to the auction price, that is, $y_{aci} = 0$ in the observed auction data. In the counterfactual exercise of subsection 6.6, I set the variance of y_{aci} according to COMEX gold futures position limits.

I assume that agent i's demand shocks y_{adi} are normally distributed, with mean 0 and variance σ^2_{di}. I only observe aggregate volume, so I cannot tell whether certain auction participants are systematically buyers or sellers of gold across auctions, so I cannot identify mean demand shocks by agent. Since auction participants are gold dealers, this assumption is plausible: dealers may be net buyers or sellers of gold in auctions depending on market conditions. These assumptions will allow me to identify the variance of agents' demand shocks by matching total auction trade volume. The estimated size of demand shocks does not affect my estimates of manipulation-induced benchmark variance: demand shock-induced variance only provides a comparison point, allowing me to ask whether manipulation variance is large relative to variance induced by idiosyncratic differences in participants' desire to buy or sell gold from auction to auction.

When I observe an auction with n_a participants, I need to determine which of the $N = 11$ agents are participating. The crucial assumption I make is that agents participate in order: every auction with n_a participants involves agents $i = 1 \dots n_a$. This is almost

[65] A list of participants is available under the "Current Auction Participants" heading at the ICE website.

[66] Appendix D.1 shows that the average number of participants has changed over time, and that participation is higher on days where gold price changes are larger, but the core stylized facts documented in figure 1 survive controlling for time trends and the size of gold price innovations.

surely violated in practice; however, while stylized, this allows the model to rationalize the observation that, from figure 1, the slope of demand is a concave function of n_a. The model rationalizes this by inferring that demand slopes κ_i are decreasing in i; that is, "later" marginal participants have lower slopes of demand. This is a simple approximate way to capture heterogeneity in agents' demand slopes.[67]

Given these assumptions, the model predicts bid curves and equilibrium prices in each auction as follows. For any n_a, we can calculate the unique equilibrium bids between n_a agents with demand slopes $(\kappa_1 \dots \kappa_{n_a})$, by numerically solving for the exact equilibrium expressions in proposition 2.[68] Equilibrium in auction a is then described by a vector of bid slopes, $(b_{a1} \dots b_{an_a})$, such that bid curves are:

$$z_{Di}(p; y_{ci}, y_{di}) = \frac{b_{ai}}{\kappa_i} y_{adi} + \frac{b_{ai}}{\sum_{j \neq i} b_{aj}} y_{aci} - b_{ai}(p - \pi) \tag{38}$$

and the auction clearing price p_a is:

$$p_a - \pi_a = \frac{1}{\sum_{i=1}^{n} b_{ai}} \left[\sum_{i=1}^{n} \left[\frac{b_{ai}}{\kappa_i} y_{adi} + \frac{b_{ai}}{\sum_{j \neq i} b_{aj}} y_{aci} \right] \right] \tag{39}$$

6.4 The LBMA dynamic auction

The LBMA gold auction differs from the baseline auction model in section 3: rather than a static demand function submission game, it is a dynamic game, in which agents are allowed to submit bids for multiple prices until a approximate market-clearing price is found. In appendix D.2, I construct a simple model of the dynamic auction; I show that there exists an equilibrium in which agents play exactly the linear equilibrium strategies described in subsection 3 in the asymmetric case, announcing their equilibrium demand $z_{Di}(p_{ar})$ when the round price is p_{ar}. This equilibrium may not be unique, but I assume

[67] In a richer model in which each auction with n_a participants may involve a different set of participants, aggregate demand slopes for auctions with $n_a = 6$ will differ depending on the particular set of participants. Thus, my estimated demand slope κ_{n_a} probably capture, roughly speaking, the difference between the average size of participants between auctions with n_a and $n_a - 1$ agents. This likely matches agents' sizes on average, but somewhat underestimates the dispersion of demand slopes κ_i. This in fact has an ambiguous effect on the predicted magnitude of manipulation incentives – if agents' κ_i values are more disperse than my estimated model suggests, "bigger" agents manipulate more and "smaller" agents manipulate less, so the effect on net manipulation is unclear.

[68] Specifically, for any vector $\kappa_1 \dots \kappa_{n_a}$, expression (12) can be solved numerically to find the unique equilibrium value of the aggregate demand slope B, and this can be plugged into expressions (9) and (11) to find equilibrium bid curves.

that agents are playing according to this equilibrium.

Assumption 6. *Agents in the dynamic auction game play strategies corresponding to proposition 2 in the baseline model.*

Under assumption 6, agents' bids in intermediate rounds of the auction reflect their equilibrium demand functions in the baseline model. This justifies interpreting the left panel of figure 1 as a plot of the aution aggregate demand curve. In other words, the volume imbalance during intermediate rounds of the auction should be equal to aggregate demand:

$$i_{ar} = \sum_{i=1}^{n} z_{Di}\left(p_{ar}, y_{adi}\right) = \sum_{i=1}^{n} y_{adi} \frac{b_{ai}}{\kappa_i} - \left(p_{ar} - \pi_a\right) b_{ai} \tag{40}$$

This immediately implies that, for two rounds r and r' within the same auction, we have:

$$i_{ar} - i_{ar'} = \left(\sum_{i=1}^{n} b_{ai}\right)\left(p_{ar} - p_{ar'}\right) \tag{41}$$

Thus, I can measure the slope of aggregate auction demand, $B_a \equiv \sum_{i=1}^{n} b_{ai}$ using any auction a with at least two rounds. In any auction with more rounds, the slope of aggregate demand is overidentified; I estimate the slope of demand $\hat{B}_a$ for auction a as the negative of the coefficient from regressing volume imbalance on prices:

$$\hat{B}_a \equiv -\frac{\sum_{t=1}^{T_a}\left(i_{at} - \bar{i}_a\right)\left(p_{at} - \bar{p}_a\right)}{\sum_{t=1}^{T_a}\left(p_{at} - \bar{p}_a\right)} \tag{42}$$

In appendix D.3, I provide some evidence that assumption 6 is violated in the data. Both buy and sell volume is lower in early rounds of the auction than later rounds, suggesting that agents are shading bids more in earlier rounds. However, appendix D.3 also shows that the slope of demand estimated using only late-round bids, which likely have less shading than early round bids, is quantitatively similar to the slope of demand estimated using all rounds; thus, violations of assumption 6 hopefully have relatively small quantitative effects on results.

6.5 Parametrization and estimation

The unknowns in the model are the vectors of agents' demand slopes $\kappa_1 \ldots \kappa_N$ and demand shock variances $\sigma^2_{d1} \ldots \sigma^2_{dN}$. The moments that I fit are the average estimated

auction demand slopes in expression (42), and auction volumes v_a, for different numbers of auction participants n_a. For any choice of $\kappa_1 \dots \kappa_N$ and $\sigma^2_{d1} \dots \sigma^2_{dN}$, for any n_a, I can numerically solve for the unique equilibrium vector of bid slopes b_{ia} using proposition 2, and thus predict the slope of aggregate demand B_a and total volume v_a.

In the data, I never observe auctions with less than 4 participants or more than 11. Thus, the model is underidentified; I only have 8 observations of average demand slope and volume, corresponding to $n_a \in \{4, 5 \dots 11\}$, and there are 11 parameters κ_i, σ^2_{di} to estimate. To solve this issue, I simply constrain κ_i and σ^2_{di} to be equal for the first four participants, $i \in \{1 \dots 4\}$.[69] I also apply some parametric smoothing in fitting the model: I assume that $\log(\kappa_i)$ and $\log\left(\sigma^2_{di}\right)$ are second-order polynomials in i,[70] and I optimize over polynomial coefficients to minimize the distance between model-predicted and data moments.

Appendices D.4 and D.5 describe further details of moment matching. The left panel of figure 2 shows the average auction slope $\hat{B}_a$ from the data for each n_a, along with predicted bid slopes from the estimated model. The right panel shows the estimated slopes of demand, κ_i. The blue line in the left panel shows that, as noted in the left panel of figure 1, the slope of aggregate demand is an increasing but concave function of the number of participants. As a result, the estimated κ_i values are decreasing; κ_1 is over 10 times as large as κ_7. In the data, the mean slope of aggregate demand $\hat{B}_a$ actually begins to decrease in n_a for $n_a > 8$. This cannot be rationalized in my model; nonetheless, the model-predicted values of B_a lie within the 95% pointwise confidence interval for the mean of $\hat{B}_a$ in the data.

The left panel of figure 3 shows average actual volume v_a for each n, together with volume predicted from the estimated model. The right panel shows the estimated sizes of demand shocks, σ^2_{di}. In the data, volume is increasing close to linearly in n_a. As a result, I infer that demand shocks for marginal participants have similar sizes; that is, σ^2_{di} is relatively similar for different values of i, as seen in the right plot. Again, the model-predicted values for volume lie within the pointwise 95% confidence intervals for the mean of v_a in the data.

[69] This is a crude way to extrapolate, but my results about the size of manipulation- and demand shock-induced variance seem to be robust to different ways to extrapolate; extrapolating by extending the second-order polynomial to $i \in \{1 \dots 4\}$ does not substantially affect my estimates for counterfactual manipulation-induced benchmark variance.

[70] The logs enforce the constraint that κ_i and σ^2_{di} are positive.

Figure 2: Gold auction demand slope and estimated κ_i values

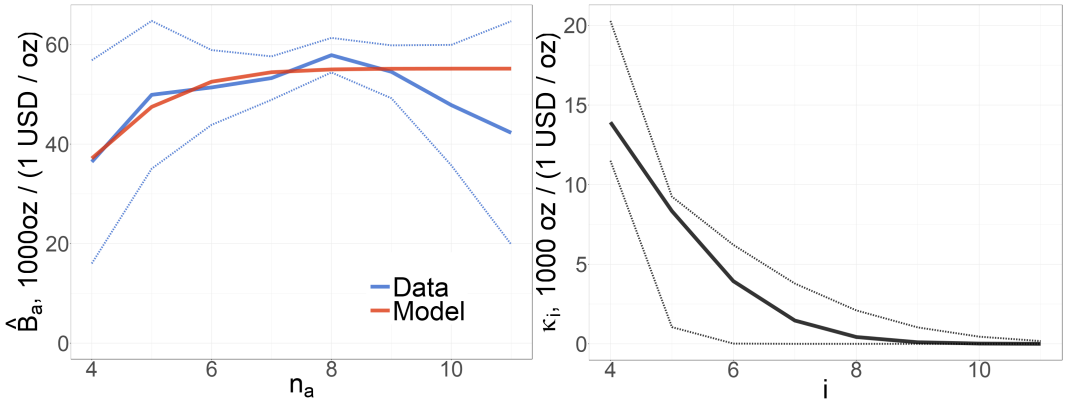

Notes. The left plot shows the average slope of demand estimated from the data, $\hat{B}_a$, for each n in blue, and the model-predicted slope of aggregate auction demand $\hat{B}(n; \theta_\kappa)$ for each n in red. The dotted lines represents 95% pointwise confidence intervals for the mean of $\hat{B}_a$ in the data. The right plot shows the values κ_i as a function of i. The dotted lines show pointwise 95% confidence intervals for the estimated κ_i values from a nonparametric bootstrap with 200 repetitions.

Figure 3: Gold auction volume and estimated $\sqrt{\sigma^2_{di}}$ values

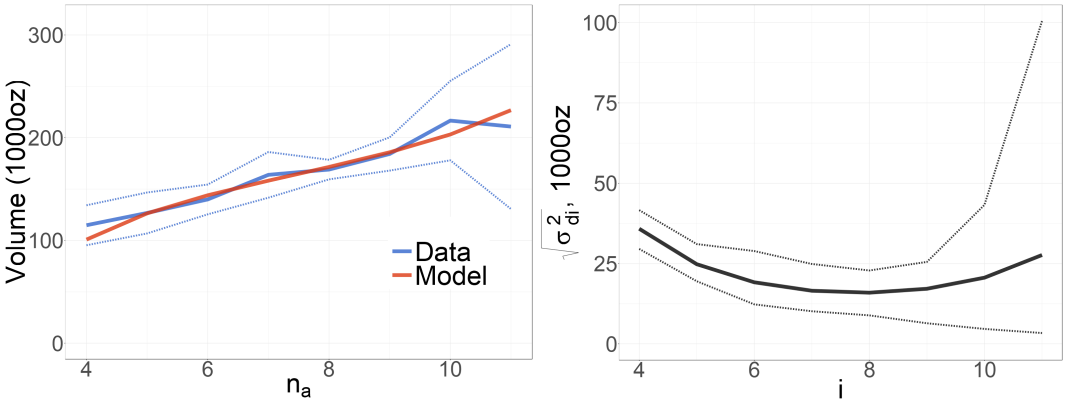

Notes. The left plot shows average volume in the data, v_a, for each n in blue, and the model-predicted volume $\hat{v}(n; \theta_\kappa, \theta_d)$ for each n in red. The dotted lines represents 95% pointwise confidence intervals for the mean of v_a in the data. The right plot shows the values $\sqrt{\sigma^2_{di}}$ as a function of i. The dotted lines show pointwise 95% confidence intervals for the estimated $\sqrt{\sigma^2_{di}}$ values from a nonparametric bootstrap with 200 repetitions.

6.6 Manipulation-induced benchmark variance with cash-settled gold futures

The goal of this section is to estimate, if CME gold futures contracts were settled based on the LBMA gold price, how much manipulation would increase the variance of the LBMA gold price. Taking the variance of expression (39), we have:[71]

$$Var\left(p_a - \pi_a\right) = \frac{1}{\left(\sum_{i=1}^{n} b_{ai}\right)^2}\left[\sum_{i=1}^{n}\left[\underbrace{\frac{b_{ai}^2\sigma_{di}^2}{\kappa_i^2}}_{\text{Demand Shocks}} + \underbrace{\frac{b_{ai}^2\sigma_{ci}^2}{\left(\sum_{j\neq i} b_{aj}\right)^2}}_{\text{Manipulation}}\right]\right] \tag{43}$$

This is close to the expression for benchmark price variance, expression (17) in section 3, but uses the exact equilibrium expressions of proposition 2 rather than the approximation of proposition 3. Similarly to expression (17), this decomposes variance in the LBMA gold price around π_a into components attributable to demand shocks and manipulation.

In order to estimate manipulation variance using expression (43), I must take a stance on the distribution of agents' contract positions. I assume that agents hold normally distributed contract positions, with mean 0 and some variance σ_{ci}^2.[72] I set the standard deviation of contract positions, $\sqrt{\sigma_{ci}^2}$, equal to 150,000 oz, half the CME gold futures position limit.[73] Since mean trade volume in each auction is only around 168,000oz across all auction participants, these contract positions are large relative to the volume of gold traded in auctions.[74] As a comparison point for manipulation-induced benchmark variance, I also calculate the size of innovations in gold prices from auction to auction: the standard deviation of first differences of p_a in the full sample of 1,650 auctions is \$7.08

[71] Since I assume that both demand shocks and contract positions are mean 0, the mean of p_a will be equal to π_a, so benchmarks will not be biased. Separately, there is an additional component of benchmark price variance caused by the fact that the auction clears with some nonzero imbalance, $|i_{aT_a}| \leqslant 10,000$oz. However, this component turns out to be approximately 100 times smaller than the demand shock and manipulation components of variance, so I do not show it in table 3.

[72] The assumption that contract positions have mean 0 is plausible because auction participants are gold dealers, and thus may be long or short gold contracts depending on market conditions.

[73] Position limits can be found on the CME website; the contract position limit is 3000, and each contract represents 100oz of gold.

[74] In dollar terms, a 1SD contract position is worth approximately \$150 million in exposure to gold. However, the maintainence margin requirement for gold futures contracts is only \$3,100 USD per contract, so entering into a 1SD contract position only requires a relatively small \$4.65 million USD margin payment.

USD/oz.[75]

Table 3 shows results. The second column of table 3 shows that the standard deviation of shocks induced by manipulation ranges from \$1.372 to \$2.694 USD/oz. This is small compared to the standard deviation of gold price innovations; the ratio of manipulation to innovation variance, shown in the third column of table 3, ranges from 3.7% to 14.5%, for different values of n_a. We can also compare the sizes of manipulation and demand shock-induced variance; the SD of demand-shock variance, reported in the fourth column, is between \$1.151-\$1.373 USD/oz. The fifth column calculates the ratio of manipulation variance to total variance from manipulation and demand shocks: manipulation variance ranges from 50.0-81.4% of total variance. Thus, settling COMEX gold futures based on the LBMA gold price would increase the variance of the LBMA gold price around 2-4 times. Moreover, column 7 shows that total benchmark variance is a small fraction of the variance in gold price innovations for all values of n_a.

The main takeaway from this section is thus that the LBMA gold price could be used to settle gold futures contracts, without manipulation causing large increases in benchmark variance. There are a number of reasons why manipulation-induced benchmark variance could be even lower than my analysis suggests. First, using the LBMA gold price to settle contracts would attract increased entry into auctions, both from manipulators attempting to profit by moving prices, and from non-manipulative arbitrageurs who profit from arbitrage when manipulators move prices. I study a simple model of costly entry into auctions, demonstrating that manipulation creates incentives for entry, in section 4 of the online appendix. Second, since gold futures contract settlement dates are only precise up to calendar months, rather than settling contracts using a single auction, multiple auctions within a month could be aggregated to produce a less manipulable benchmark. In appendix D.7, I show that an aggregated benchmark averaging prices of all auctions within a month would be much less vulnerable to manipulation than a single auction.

[75]Using the sample of all 1650 auctions, the gold price series p_a does not reject the null (p-value 0.3825) in an augmented Dickey-Fuller test against the hypothesis that the series is stationary. Moreover, an AR(10) regression of first-differences of p_a has no significant coefficients at the 5% level, suggesting that the series p_a is close to a random walk. p_a does not have uniformly sized increments: the standard deviation of price differences from afternoon to morning auctions is \$1.92 USD/oz larger than the SD of differences from morning to afternoon auctions, but I disregard this difference in the following analysis.

Table 3: Components of gold auction price variance

n_a	Manip SD	$\frac{Manip\ var}{Innov\ var}$	Demand SD	$\frac{Manip\ var}{Total\ var}$	$\frac{Total\ var}{Innov\ var}$
4	2.694 (0.388)	14.5% (3.98%)	1.287 (0.209)	81.4% (2.71%)	17.8% (4.86%)
5	1.799 (0.107)	6.4% (0.79%)	1.182 (0.122)	69.9% (3.08%)	9.2% (1.27%)
6	1.502 (0.086)	4.5% (0.56%)	1.151 (0.071)	63.0% (4.06%)	7.1% (0.58%)
7	1.406 (0.141)	3.9% (0.86%)	1.159 (0.049)	59.6% (5.01%)	6.6% (0.86%)
8	1.379 (0.169)	3.8% (1.00%)	1.185 (0.040)	57.5% (5.53%)	6.6% (1.04%)
9	1.373 (0.180)	3.8% (1.05%)	1.223 (0.038)	55.8% (5.84%)	6.7% (1.11%)
10	1.372 (0.185)	3.7% (1.07%)	1.278 (0.070)	53.5% (6.48%)	7.0% (1.17%)
11	1.372 (0.187)	3.7% (1.08%)	1.373 (0.281)	50.0% (9.88%)	7.5% (2.33%)

Notes. Components of gold auction price variance, in units of USD/oz. Values in parentheses are standard errors from a nonparametric bootstrap with 200 repetitions. "Manip SD" is the square root of manipulation variance, calculated using expression (43). $\frac{Manip\ Var}{Innov\ Var}$ is the ratio of manipulation variance to innovation variance, which is \$7.08 USD/oz. "Demand SD" is the square root of demand variance, calculated using expression (43). $\frac{Manip\ Var}{Total\ Var}$ is the ratio of manipulation variance to the sum of manipulation plus demand variance. $\frac{Total\ Var}{Innov\ Var}$ is the ratio of total (manipulation plus demand) variance to innovation variance. All SD numbers are in units of USD/oz.

7 Redesigning the CBOE VIX

In this section, building on the theory of section 4, I design a manipulation-consistent weighting scheme for aggregating S&P500 (SPX) option prices to calculate the VIX volatility index. I show that my manipulation-consistent estimator improves substantially on the VIX formula currently used by CBOE.

Subsection 7.1 describes background on the CBOE VIX. Subsection 7.2 discusses the data I use. Subsection 7.3 describes my model for the fundamental option price function $\pi_t(k)$, that is, the integrand of the integral that VIX is designed to estimate, and the vector of observed option prices $\mathbf{p}_t^{H_t}$, for each date t. Subsection 7.4 describes the parametrization and estimation of the model for $\pi_t(k)$ and $\mathbf{p}_t^{H_t}$. Subsection 7.5 discusses how I use Bayesian quadrature to generate price weighting schemes sensitive to manipulation-induced noise. Subsection 7.6 shows the price weighting schemes generates, and subsection 7.7 shows how much my weighting scheme improves on the current formula used by CBOE to calculate the VIX.

7.0.1 Notation and terminology

In this section, I will often need to refer separately to the function $p(k)$, the vector $(p(k_1)\dots p(k_H))$ of values of $p(\cdot)$ at a length-H collection of strike prices $\{k_1\dots k_h\dots k_H\}$, and square and rectangular matrices of variances and covariances of the vector of values of $p(\cdot)$. Functions will be denoted in regular case, $\pi(k)$. Vectors will be denoted in bold, as in $\mathbf{p}_t^{H_t}$. Subscripts will represent time, and superscripts represent the length of the vector; for example, the vector of prices $\mathbf{p}$ observed at time t on a length-H_t grid of points will be denoted $\mathbf{p}_t^{H_t}$. Vector indexing will be denoted by parentheses, so $\mathbf{p}_t^{H_t}(h)$ is element h of the vector $\boldsymbol{\pi}$. I will often refer to vectors of prices $\mathbf{p}_t^{H_t}$ on a length-H_t grid of strike prices $\{k_{t1}\dots k_{tH_t}\}$; with some abuse of notation, I will use $\mathbf{p}_t^{H_t}(k_{th})$ to mean element h of the vector $\mathbf{p}_t^{H_t}$, corresponding to the price of the option at strike price k_{th}. The transpose of a vector $\boldsymbol{\phi}$ is denoted $\boldsymbol{\phi}^{\mathbf{T}}$ or $(\boldsymbol{\phi})^{\top}$. I omit superscripts occasionally when there is no ambiguity about the length of a vector.

Matrices will be denoted by $\boldsymbol{\Sigma}$. Matrix subscripts denote time and the role that the matrix plays; for example, $\boldsymbol{\Sigma}_{manip,t}$ denotes the time-t covariance matrix of manipulation-induced noise. Superscripts will denote column and row counts; I will write $\boldsymbol{\Sigma}_t^{H_t}$ for a square matrix observed at time t with H_t rows and columns, and $\boldsymbol{\Sigma}_t^{G,H_t}$ for a matrix

observed at time t with G rows and H_t columns. The row i, column j element of a matrix is denoted $\mathbf{\Sigma}_t^{G,H_t}(i,j)$.

Throughout, I use "at the money" and "out of money" with respect to forward prices; thus, a near-the-money option has strike price K_t approximately equal to the forward price F_t, or the ratio $k_t \equiv \frac{K_t}{F_t} \approx 1$, and a far-out-of-money option has k_t much larger or smaller than 1.

7.1 Background

The Chicago Board Options Exchange (CBOE) Volatility Index, or VIX, is a measure of month-ahead volatility implied by the prices of S&P 500 options. It is based on a weighted sum of prices of out-of-money SPX options. Intuitively, the price of options is higher when market participants expect the S&P 500 to be more volatile. Technically, the VIX formula is based on the fact that that variance swaps can be replicated using options; specifically, a portfolio whose return is equal to the risk-neutral expectation of the variance of the return on the S&P 500 can be constructed by buying an infinite number of out-of-money options, with portfolio weights inversely proportional to the square of the option strike price.[76] This result holds so long as the price of the underlying security evolves continuously, hence the VIX formula measures the variance of the underlying process under much more general assumptions than, for example, Black-Scholes implied volatilities. In 2004, CBOE introduced cash-settled VIX futures contracts, and in 2006 CBOE introduced VIX options. VIX derivatives are now widely traded; open interest in VIX futures in most months of 2017 and 2018 peaked at over 200,000 contracts, which represents over $3 billion dollars of exposure to VIX.[77]

The value of VIX for determining contract payments is calculated based on exchange opening auctions prices of S&P 500 options 30 days from expiration. The VIX is calculated using a Riemann sum over option prices, essentially assuming that all option prices are observed noiselessly; CBOE uses all observed out-of-money option prices in VIX calculation, until it reaches two consecutive options with no bids. This implies that the VIX settlement formula often assigns high weight to far-out-of-money options, which

[76] Demeterfi et al. (1999) discuss the theory of volatility and variance swaps, the CBOE VIX White Paper discusses the current VIX calculation methodology in detail, and Carr and Wu (2006) compares the current VIX methodology to the original methodology based on implied volatilities of close-to-the-money options.

[77] The mean value of the VIX is around 16.8, and the contract multiplier for VIX futures contracts is $1000USD.

are very thinly traded, and thus in principle prone to manipulation. Griffin and Shams (2018) show that far-out-of-money SPX options are heavily traded on VIX settlement dates, roughly proportionally to their weight in VIX calculation, suggesting that agents are attempting to move the VIX settlement value with their trades and influence VIX contract payoffs. In April 2018, the VIX spiked upwards on the contract settlement date, caused by a large bid for far-out-of-money put options.[78] A whistleblower has also claimed that VIX has been manipulated (Louis and Gammeltoft, 2018), and an investigation of manipulation is ongoing.

7.2 Data

7.2.1 Option prices and variance swaps

Suppose we observe put and call option prices for the S&P500 at a continuum of strike prices, without noise from manipulation or demand shocks. Let F_t represent the time t forward price of the underlying S&P500 index, and let $C_t(K)$ and $P_t(K)$ represent prices of put and call options at strike price K. Let $O_t(K)$ represent the price of the out-of-money option at strike price K, that is:

$$O_t(K) = \begin{cases} C_t(K) & K > F_t \\ P_t(K) & K < F_t \end{cases}$$

I will estimate volatility using exchange closing prices of options a day before VIX settlement dates; the time to maturity of these options is $T = \frac{30.718}{365}$ years, or somewhat less than 31 days.[79] Let r_t represent the 31-day interest rate at date t. I first normalize all option prices and strike prices by the forward index level. Thus, define:

$$k \equiv \frac{K}{F_t}, \; o_t(k) \equiv \frac{O_t(F_t k)}{F_t} \tag{44}$$

Define the normalized option prices:

$$\pi_t(k) = \frac{2}{T} e^{T r_t} \frac{o_t(k)}{k^2} \tag{45}$$

[78] Odd Spike in Wall Street Fear Gauge Awakens Manipulation Debate

[79] Option Metrics reports SPX option closing prices, at 3:15pm the day before VIX settlement. VIX settlement occurs at 8:30am, exactly 30 days before SPX option settlement. There are 17.25 hours between 3:15pm and 8:30am, so the time to maturity of SPX options in my dataset is $\left(\frac{1}{365}\right)\left(30 + \frac{17.25}{24}\right) \approx \frac{30.718}{365}$.

Since k is the strike price normalized by the forward price, the function $\pi_t(k)$ can be compared across dates with different levels of the forward price F_t. If the risk-neutral distribution of SPX returns is unchanged, the function $\pi_t(k)$ observed at normalized strike price k should likewise be unchanged. Appendix E.1 follows Carr and Wu (2006) to show that the risk-neutral expectation S&P500 return variance, the quantity that VIX^2 is designed to estimate, is:

$$V_t^2 = \int_0^\infty \pi_t(k)\, dk \tag{46}$$

Thus, V_t^2 is simply an integral over out-of-money option prices $o_t(k)$, adjusted by strike prices, interest rates, and time to expiration, using expression (45). The CBOE VIX is designed to estimate V_t^2.[80] There are two difficulties in estimating V_t^2 from observed option prices. First, on any given date t, we can only observe a finite-length vector $\mathbf{p_t^{H_t}}$ of option prices, at a finite number of strike prices. Second, observed option prices $\mathbf{p_t^{H_t}}$ are not the fundamental option prices $\pi_t(k)$, but are subject to noise from demand shocks and manipulation.

7.2.2 Observed option price vectors $\mathbf{p_t^{H_t}}$

The main data source I use is Option Metrics data on daily SPX option closing prices prices and trade volume, from 2010-05-18 to 2017-12-19. I restrict my sample to monthly SPX options 31 days before expiration; this is to avoid using option prices on VIX settlement dates, since Griffin and Shams (2018) provides evidence that SPX option prices are manipulated at VIX settlement. I also use Option Metrics data on the daily zero coupon yield curve to get 31-day interest rates r_t.

Let t index the 92 dates at which we observe option prices. For each date t option, the time to maturity is $T = \frac{30.718}{365}$, in units of years. On each date t, we observe vectors of prices of call options $\mathbf{C_t^{H_t}}(K_h)$ and $\mathbf{P_t^{H_t}}(K_h)$, at strike prices K_h on a size H_t grid of strike prices, $K_{t1} \ldots K_{tH_t}$. Observed prices may be subject to noise from demand shocks and manipulation. As I describe in appendix E.4, for each date t, I estimate forward prices $\tilde{F}_t$ from option prices $\mathbf{C_t^{H_t}}$ and $\mathbf{P_t^{H_t}}$ using put-call parity. This allows me to calculate the

[80] A minor difference to my definition is that CBOE uses option prices exactly 30 days before expiration, but as the definition of $\pi_t(k)$ normalizes by time to maturity, this only affects the integral in (46) through the term structure of volatility, which is likely to have small effects at short time spans.

vector of out-of-money option prices $\mathbf{O_t^{H_t}}$:

$$\mathbf{O_t^{H_t}}(K_{th}) = \begin{cases} \mathbf{C_t^{H_t}}(K_{th}) & K_h > \tilde{F}_t \\ \mathbf{P_t^{H_t}}(K_{th}) & K_h < \tilde{F}_t \end{cases}$$

As in expression (44), I first normalize all strike prices K_h and option prices $\mathbf{O_t^{H_t}}(K_{th})$ by the option-implied forward price $\tilde{F}_t$:

$$k_{th} \equiv \frac{K_{th}}{\tilde{F}_t},\ \mathbf{o_t^{H_t}}(k_{th}) = \frac{\mathbf{O_t^{H_t}}(K_{th})}{\tilde{F}_t} \tag{47}$$

Then, following expression (45), I further normalize option prices by interest rates, expiration dates, and the squared normalized strike price. I will call the observed vector of normalized option prices $\mathbf{p_t^{H_t}}$, where the k_{th} element of $\mathbf{p_t^{H_t}}$ is defined as:

$$\mathbf{p_t^{H_t}}(k_{th}) = \frac{2}{T} e^{Tr_t} \frac{\mathbf{o_t^{H_t}}(k_{th})}{k_{th}^2} \tag{48}$$

The observed normalized option prices $\mathbf{p_t^{H_t}}(k_{th})$ can differ from the fundamental prices $\pi_t(k_{th})$ because of demand shocks and manipulation; I will describe my model of $\pi_t(k_{th})$ and $\mathbf{p_t^{H_t}}(k_{th})$ in subsection 7.3 below.

Since far-out-of-money option prices are not always observed, I will approximate expression (46) using:

$$V_t^2 = \int_{0.6}^{1.1} \pi_t(k)\, dk \tag{49}$$

Thus, I will discard all observed option prices $\mathbf{p_t^{H_t}}(k_{th})$ with normalized strike prices $k_{th} < 0.6$ and $k_{th} > 1.1$, that is, with normalized strike prices lower than $0.6\times$ or greater than $1.1\times$ the option-implied forward price $\tilde{F}_t$. My final estimation sample consists of 14,984 observed option prices at 92 dates. The average number of option prices I observe on each date is 163, with maximum 221 and minimum 103.[81] Further details about the dataset and my cleaning steps are available in appendix E.2.

Figure 4 plots the observed vectors of prices, $\mathbf{p_t^{H_t}}$, for all 92 dates, highlighting in color the observations from January of each year from 2011 to 2017. The values of $\mathbf{p_t^{H_t}}(k_{th})$

[81] The number of prices per day varies over time because the average spacing between strike prices at which SPX options are traded decreases over the sample time period.

at different strike prices k_{th} tend to move upwards and downwards together. However, there are small but noticeable differences in the shape of $\mathbf{p}_t^{H_t}$ at different times. For example, $\mathbf{p}_t^{H_t}$ in January of 2012 is low near $k = 1$, but is somewhat higher for low and high of k than $\mathbf{p}_t^{H_t}$ in 2013 and 2014; this suggests that market participants in 2012 assigned somewhat higher probability to large jumps in the price of the underlying. Another observation is that the $\mathbf{p}_t^{H_t}$ vectors are not perfectly smooth in the strike price k – especially in the tails, there is some noticeable waviness in the functions, which appears to be localized to one or a few strike prices. This can be rationalized in the context of my model as noise in prices induced by demand shocks, which are independent across strike prices.

The goal of the benchmark designer is to estimate V_t^2, defined in expression (49), using the vector of observed prices $\mathbf{p}_t^{H_t}$. The CBOE VIX estimates 100 times the square root of V_t^2; since the integral (49) is linear in prices, I will attempt to construct benchmarks that estimate V_t^2, rather than its square root. I will also consider only benchmarks which are affine functions of option prices. Thus, I consider benchmarks P_b which estimate V_t^2 as defined in expression (49), characterized by some constant C and a vector of submarket weights $\boldsymbol{\phi}_t^{H_t}$, such that:[82]

$$P_b = C + \sum_{h=1}^{H_t} \boldsymbol{\phi}_t^{H_t}(k_{th})\, \mathbf{p}_t^{H_t}(k_{th}) \tag{50}$$

In vector notation, this is:

$$P_b = C + \left(\boldsymbol{\phi}_t^{H_t}\right)^T \mathbf{p}_t^{H_t}$$

The benchmark designer's problem is to choose C and $\boldsymbol{\phi}_t^{H_t}$ to estimate the V_t^2 integral, expression (49), given the fact that option prices $\mathbf{p}_t^{H_t}$ are observed only at a fininte number of points, and are subject to noise from demand shocks and manipulation.

7.3 Model

In subsection 7.3.1, I model the fundamental option price function $\pi_t(k)$ as a Gaussian process; this is a flexible statistical model of prices, which is agnostic to the underlying

[82]Note that I allow weights $\phi_t(k)$ to differ by date. As I discuss below, CBOE's current weights, which are based on the Riemann sum formula, also vary by date, as they depend on the spacing between SPX options at VIX settlement date, which varies over time.

Figure 4: Distribution of observed option prices

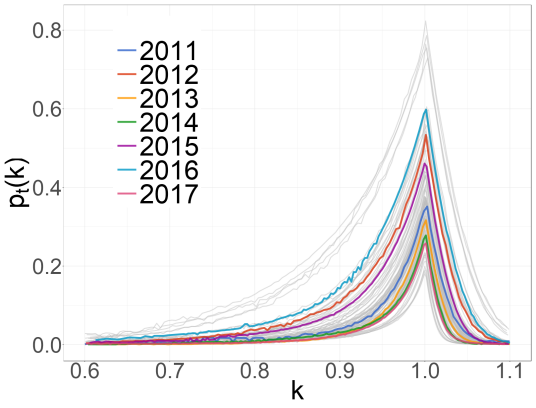

Notes. Each lines is an observation of $\mathbf{p_t^{H_t}}$ on a single date t, with the normalized strike price on the x-axis. Colored lines are observations of $\mathbf{p_t^{H_t}}$ in January of the year indicated on the legend.

economic process generating $\pi_t(k)$. In subsection 7.3.2, I model $\mathbf{p_t^{H_t}}$ as determined in an auction between symmetric agents, so that $\mathbf{p_t^{H_t}}(k_{th})$ is equal to $\pi_t(k_{th})$ plus noise from demand shocks and manipulation.

7.3.1 Fundamental option price function $\pi_t(k)$

Figure 4 suggests that option prices at different strikes are very correlated, but also that the mean and variance of option prices depends on the normalized strike price k. Thus, I need a model of the fundamental option price process $\pi_t(k)$ which can match these features of the data. Instead of imposing an economic model on option prices, I use a flexible nonparametric statistical model for $\pi_t(k)$. I assume that $\pi_t(k)$ is a Gaussian process: the function $\pi_t(k)$ is a random variable, independently drawn over dates t, with the property that the vector $\boldsymbol{\pi_t^H} = \boldsymbol{\pi_t^H}(k_1) \ldots \boldsymbol{\pi_t^H}(k_H)$ of values of $\pi_t(k)$ observed at any finite set of points $\{k_1 \ldots k_h \ldots k_H\}$, follows a multivariate normal distribution. The mean of $\pi_t(k)$ is some function $\mu(k)$, and the covariance matrix of the vector $\boldsymbol{\pi_t^{H_t}} = \boldsymbol{\pi_t^{H_t}}(k_1) \ldots \boldsymbol{\pi_t^{H_t}}(k_H)$ is described by some *covariance function* $\mathrm{Cov}(k_h, k_{h'})$, which may depends on strike prices k_h and $k_{h'}$. Let $\boldsymbol{\Sigma_{\pi,t}^{H_t}}$ be the covariance matrix of $\boldsymbol{\pi^{H_t}}$; the (h, h') element of $\boldsymbol{\Sigma_{\pi,t}^{H_t}}$ is:

$$\boldsymbol{\Sigma_{\pi,t}^{H_t}}(h, h') = \mathrm{Cov}(k_h, k_{h'}) = \sigma_\pi(k_h)\,\sigma_\pi(k_{h'}) \exp\left(-\frac{(k_h - k_{h'})^2}{2l^2}\right) \tag{51}$$

This is a valid covariance function, since Rasmussen and Williams (2005, pg. 95) shows that a rescaling of a covariance function is another covariance function. In the Gaussian process literature, the function

$$\exp\left(-\frac{(k_h - k_{h'})^2}{2l^2}\right) \tag{52}$$

is called the *squared exponential* covariance function. I multiply this by scaling functions $\sigma_\pi(k_h)\,\sigma_\pi(k_{h'})$ which depend on the strike prices $k_h, k_{h'}$, to fit the observation in figure 4 that the variance of $\mathbf{p}_t^{\mathbf{H}_t}(k_{th})$ is higher for values of k_{th} closer to 1. $\sigma_\pi(k)$ can be interpreted as a "signal variance" function – if $\sigma_\pi(k)$ is high at a given strike price k, then $\pi_t(k)$ has high variance at k, and thus strike price k contributes substantially to variance in the V_t^2 integral, expression (49). Benchmarks should thus assign high weight to strike prices with high $\sigma_\pi(k)$ relative to noise.

The correlation term in expression (51), that is, expression (52), implies that $\pi_t(k), \pi_t(k')$ are correlated when k, k' are close; that is, options at close strike prices have correlated prices. The "length-scale" parameter l determines the decay rate of these price correlations. If l is large, covariances decay slowly in the difference between strike prices $(k_h - k_{h'})^2$, so option prices across the entire range of strikes tend to move together. If l is small, covariances decay quickly in $(k_h - k_{h'})^2$, hence innovations in $\pi_t(k)$ are more localized.

The primary technical motivation for modelling option prices as Gaussian processes is that this leads to a particularly simple noise-sensitive integration procedure; as I discuss in subsection 7.5 below, posterior inference for integrals of Gaussian processes with noisily observed inputs reduces to a simple set of OLS-like matrix equations, which produces quadrature rules which are affine in option prices. The cost of this analytical simplicity is that Gaussian processes are an unrealistic model of option prices. I do not impose any economic restrictions on prices; under the Gaussian process model, option prices may be nonconvex in strike prices, or they may even be negative. Moreover, the Gaussian process model forces the marginal distribution of option prices at any given strike price to be normal; in practice, figure 4 suggests that the distribution of option prices is noticeably skewed upwards. In principle, it may be possible to preserve some of the analytical simplicity of the Gaussian process model while imposing theory-implied restrictions on option price functions, but I leave this to future work.[83]

[83]Rasmussen and Williams (2005) discusses the use and theoretical underpinnings of Gaussian processes in detail. My model also disregards the possibility of intertemporal correlation in $\pi_t(k)$. In principle, a fully

7.3.2 Observed option price vector $\mathbf{p_t^{H_t}}$

At each date t, we observe option prices at an exogeneous grid of normalized strike prices, $k_1 \dots k_{th} \dots k_{tH_t}$. While the model of the fundamental option prices $\pi_t(k)$ is particular to the VIX setting, I model observed prices $\mathbf{p_t^{H_t}}(k_{th})$ essentially identically to the model of section 4. For each strike price k_{th}, there are $n-1$ regular traders indexed by i. If a regular trader purchases $z_t(k_{th})$ options at price $\mathbf{p_t^{H_t}}(k_{th})$, her utility is:[84]

$$U_{reg}\left(\mathbf{p_t^{H_t}}(k_{th}), z_t(k_{th})\right) = \pi_t(k_{th}) z_t(k_{th}) + \frac{y_{tdi}(k_{th}) z_t(k_{th})}{\kappa(k_{th})} - \frac{z^2}{2\kappa(k_{th})} - \mathbf{p_t^{H_t}}(k_{th}) z(k_{th}) \tag{53}$$

As in the analysis of aggregated benchmarks in section 4, I assume that demand shocks $y_{tdi}(k_{th})$ are normally distributed, with mean 0, and uncorrelated across strike prices. There is a single manipulator, $i = 1$, holding y_{tc} contracts, with monetary payment $P_b y_{tc}$.[85] I assume that the manipulator's contract position y_{tc} is normally distributed, with mean 0 and variance σ_c^2. The manipulator's utility is the sum of utilities across options at all strike prices, plus her contract payoffs:

$$\begin{aligned} U_{manip}\left(\mathbf{p_t^{H_t}}(k_{th}), z(k_{t1}) \dots z(k_{tH_t})\right) = \\ \left[\sum_{h=1}^{H} \pi_t(k_{th}) z(k_{th}) + \frac{y_{td1}(k_{th}) z(k_{th})}{\kappa(k_{th})} - \frac{z^2}{2\kappa(k_{th})} - \mathbf{p_t^{H_t}}(k_{th}) z(k_{th})\right] + \\ y_{tc}\left[C + \sum_{h=1}^{H_t} \boldsymbol{\phi}_t^{H_t}(k_{th}) \mathbf{p_t^{H_t}}(k_{th})\right] \end{aligned} \tag{54}$$

where $\bar{\mathbf{y}}_{\mathbf{dt}}^{\mathbf{H_t}}(k_{th}) = \sum_{i=1}^{n} y_{tdi}(k_{th})$ is the sum over all agents' demand shocks in the sub-market at strike price k_{th}. Thus, the observed option price $\mathbf{p_t^{H_t}}(k_{th})$ differs from the

Bayesian estimate of $\pi_t(k)$ should incorporate all prior information about $\pi_t(k)$; VIX is autocorrelated at the monthly frequency, so the value of VIX at time $t-1$ is informative about the value of VIX today. I ignore autocorrelation, deriving the Bayesian posterior mean of V_t^2 as if observations of $\pi_t(k)$ were independent over time. I ignore intertemporal correlations partially for analytical simplicity, but also because we wish to construct a point-in-time estimator of volatility. Using a fully Bayesian estimator which accounts for intertemporal correlations in VIX would mechanically smooth the benchmark over time, which may be practically undesirable.

[84] I represent $z_t(k_{th}), y_{tdi}(k_{th}), \kappa(k_{th})$ in regular case, so they can be interpreted as functions, though the distinction between functions and vectors here is unimportant, as these expressions are never observed. They are only needed to derive the expression for observed prices, expression (55).

[85] Since agents are symmetric, the contract position of the manipulator could also be thought of as the net contract position across all market participants, in a richer model where all agents are allowed to hold contracts and manipulate.

fundamental option price $\pi_t(k_{th})$ at strike k_{th} because of demand shocks and manipulation. Solving for equilibrium is isomorphic to the analysis of aggregated benchmarks in section 4; in appendix E.3, I show that the equilibrium option price at strike k_{th} is:

$$p_t^{H_t}(k_{th}) = \pi_t(k_{th}) + \underbrace{\frac{\bar{y}_{dt}^{H_t}(k_{th})}{n\kappa_t^{H_t}(k_{th})}}_{\text{Demand Shocks}} + \underbrace{\frac{\phi_t^{H_t}(k_{th})}{n(n-2)\kappa_t^{H_t}(k_{th})} y_{tc}}_{\text{Manipulation}}, \; h \in \{1 \dots H_t\} \tag{55}$$

Expression (55) implies that errors in prices generated by demand shocks and manipulation will have different covariance structures. Letting $\Sigma_{demand,t}^{H_t}$ represent the covariance matrix of the demand shock term in expression (55); I assume that the h, h' element of $\Sigma_{demand,t}^{H_t}$ is:

$$\Sigma_{demand,t}^{H_t}(h, h') \equiv Cov\left[\frac{\bar{y}_{dt}^{H_t}(k_{th})}{n\kappa_t^{H_t}(k_{th})}, \frac{\bar{y}_{dt}^{H_t}(k_{th'})}{n\kappa_t^{H_t}(k_{th'})}\right] = \delta_{hh'}\sigma_d^2(k)$$

$\delta_{hh'}$ is the delta function, equal to 1 if $h = h'$ and equal to 0 otherwise; thus, $\Sigma_{demand,t}^{H_t}$ is a diagonal matrix. $\sigma_d^2(k)$ is a flexible function which determines how the variance of demand shocks depends on the strike price k. In the gold application of section 6, I estimated $\sigma_d^2(k)$ using data on total auction volume; since I do not observe agents' behavior in the VIX settlement auctions of SPX options, I do not use the same identification strategy here. Instead, I identify demand shocks relying entirely on the assumption that demand shocks are independent across strike prices. In practice, the estimated value of $\sigma_d^2(k)$ is very low across the entire range of strike prices; thus, essentially all variation in the observed data is attributed to changes in $\pi_t(k)$.

Manipulation, in contrast to demand shocks, induces errors which are perfectly correlated across strike prices. Letting $\Sigma_{manip,t}^{H_t}$ represent the covariance matrix of the manipulation error term in expression (55), the h, h' element of $\Sigma_{manip,t}^{H_t}$ is:

$$\Sigma_{manip,t}^{H_t}(h, h') \equiv Cov\left[\frac{\phi_t^{H_t}(k_{th})}{n(n-2)\kappa_t^{H_t}(k_{th})} y_{tc}, \frac{\phi_t^{H_t}(k_{th})}{n(n-2)\kappa_t^{H_t}(k_{th})} y_{tc}\right] = \left(\frac{1}{n^2(n-2)^2}\right)\left(\frac{\phi_t^{H_t}(k_{th})\,\phi_t^{H_t}(k_{th'})}{\kappa_t^{H_t}(k_{th})\,\kappa_t^{H_t}(k_{th'})}\right)\sigma_c^2$$

As shown in section 4, manipulation noise is thus larger at strike prices with higher weight

$\boldsymbol{\phi}_t^{H_t}(k_{th})$, since the manipulator preferentially trades in submarkets with higher weights, and at strike prices lower slopes of demand $\boldsymbol{\kappa}_t^{H_t}(k_{th})$, since the manipulator's trades move prices more. Manipulation noise is also influenced by submarket competitiveness; however, I do not have data on the number of participants in each submarket, so I must assume that n is equal across submarkets.

The dependence of $\boldsymbol{\Sigma}_{manip,t}^{H_t}$ on the submarket weight vector $\boldsymbol{\phi}_t^{H_t}$ creates a fixed-point problem: given some manipulation noise matrix $\boldsymbol{\Sigma}_{manip,t'}^{H_t}$ if we choose $\boldsymbol{\phi}_t^{H_t}$ optimally to trade off signal and noise across submarkets, this affects agents' manipulation decisions, and thus changes the manipulation noise structure $\boldsymbol{\Sigma}_{manip,t}^{H_t}$. I discuss this further in subsection 7.5; essentially, I will require the weight vector $\boldsymbol{\phi}_t^{H_t}$ to be consistent with the manipulation noise structure $\boldsymbol{\Sigma}_{manip,t}^{H_t}$ that it induces.

All terms on the RHS of expression (55) for the price vector $\mathbf{p}_t^{H_t}(k_{th})$ are normal; hence, for any date t, the vector $\mathbf{p}_t^{H_t}(k_{th})$ is a finite-dimensional multivariate normal random variable, with covariance matrix:

$$\boldsymbol{\Sigma}_{p,t}^{H_t} = \boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t} + \boldsymbol{\Sigma}_{manip,t}^{H_t} \tag{56}$$

This implies that the parameters

$$\mu(k), \sigma_\pi(k), \sigma_d^2(k), l$$

of the Gaussian process for $\pi_t(k)$ and $\mathbf{p}_t^{H_t}(k_{th})$ can be estimated from the observed price vectors $\mathbf{p}_t^{H_t}(k_{th})$ using maximum likelihood.

Since I use SPX option prices 31 days before maturity, observed option prices should be unaffected by manipulation; thus, the observed price vectors $\mathbf{p}_t^{H_t}$ should reflect only variation in $\boldsymbol{\pi}_t$ and demand shocks. My estimation strategy is thus essentially the same as that of section 6: I estimate parameters of the model using data where there should be no manipulation, then use the estimated parameters to counterfactually predict the amount by which manipulation would increase benchmark variance for any conjectured size of contract positions σ_c^2.

7.4 Parametrization and estimation

The data I use is significantly less detailed than the gold auction data used in section 6, so I need to impose two important simplifying assumptions. First, as discussed in subsection

7.3.2, I do not observe the number or size of participants, so I assume that agents are symmetric, and set $n = 10$ for all strike prices. Some of the most competitive contract markets in the CFTC Commitments of Traders data in section 3 have approximately $s_{max} = \frac{1}{10}$, so this is plausibly a high estimate. Second, I do not observe data on VIX auctions at settlement, so I measure slopes of demand $\boldsymbol{\kappa}_t^{H_t}(k_{th})$ using limit order book quote data, as if it were auction data. It is known in the empirical market microstructure literature that using limit order book quote data tends to overestimate price impact.[86]

The parameters of the model are:

$$\mu(k), \sigma_\pi(k), \sigma_d^2(k), l, \boldsymbol{\kappa}_t^{H_t}(k_{th}), \sigma_c^2$$

In subsection 7.4.1, I discuss maximum likelihood estimation of the Gaussian process parameters

$$\mu(k), \sigma_\pi(k), \sigma_d^2(k), l$$

using observed option prices $\mathbf{p}_t^{H_t}$. In subsection 7.4.2, I discuss estimation of demand slopes $\boldsymbol{\kappa}_t^{H_t}(k_{th})$ using best bid and ask prices and volumes. In subsection 7.4.3, I discuss estimation of contract position size σ_c^2 using data on open interest in VIX futures at expiration. Further details on estimation are available in various subsections of appendix E.

7.4.1 Maximum likelihood estimation of Gaussian process parameters

The unknowns of the Gaussian process model for the fundamental option price function $\pi_t(k)$ and the observed option price vectors $\mathbf{p}_t^{H_t}$, are the mean function, $\mu(k)$, the covariance scaling function $\sigma_\pi(k)$, the demand shock noise scaling function $\sigma_d^2(\cdot)$, and the covariance length-scale parameter l. Since I use option prices $\mathbf{p}_t^{H_t}$ measured 31 days before expiration, my measurements are a day before VIX settlement dates, so option prices should reflect only demand shocks and not manipulation. Thus, the covariance matrix of the price vector $\mathbf{p}_t^{H_t}$ is:

$$\boldsymbol{\Sigma}_{p,t}^{H_t} = \boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t}$$

[86] If actual bidding data from VIX settlement auctions were available, I could instead use bidding data to estimate agents' slopes of demand for options, as I do in section 6.

I model the functions $\mu(k), \sigma_\pi(k), \sigma_d^2(\cdot)$ using b-splines and i-splines; I describe the parametrization of the model in detail in appendix E.5. Let $\theta_\mu, \theta_\pi, \theta_d$ respectively represent the spline parameters for $\mu(k), \sigma_\pi(k)$ and $\sigma_d^2(k)$, and let $\Theta = \theta_\mu, \theta_\pi, \theta_d, l$ represent the full vector of parameters characterizing the Gaussian process for $\mathbf{p}_t^{H_t}$ and $\pi_t(k)$. Any given set of parameters Θ determines the functions $\mu(k), \sigma_\pi(k)$ and $\sigma_d^2(k)$, which then determines predicted means $\boldsymbol{\mu}^{H_t}$ and covariance matrices $\boldsymbol{\Sigma}_{p,t}^{H_t}$ for the vector $\mathbf{p}_t^{H_t}$ on any date t. Since $\mathbf{p}_t^{H_t}$ is multivariate normal, I can thus calculate the likelihood of any $\mathbf{p}_t^{H_t}$ given $\boldsymbol{\mu}^{H_t}, \boldsymbol{\Sigma}_{p,t}^{H_t}$, and I can fit Θ to the observed vectors of prices $\mathbf{p}_t^{H_t}$ using maximum likelihood. I derive the Gaussian process log likelihood function in appendix E.6, and I discuss numerical details of optimization in appendix E.7.

Figure 5 shows the estimated $\mu(k)$ function, in black. The green lines show the 95% pointwise Bayesian credible intervals:

$$\mu(k) \pm 2\sqrt{\boldsymbol{\Sigma}_{p,t}^{H_t}(k,k)}$$

This captures where the function values should fall 95% of the time. The model fits the data fairly well, with the caveats that observed option prices have positive skew and are always nonnegative; both of these facts cannot be rationalized in the Gaussian process model. The estimated signal variance function $\sigma_\pi(k)$ scales roughly proportionately with the mean function $\mu(k)$, being large near $k = 1$ and small for further values of k, where the normalized price $\pi_t(k)$ is also lower. This implies that most of the variation in the V_t^2 integral – that is, most of the "signal" value in option prices – comes from close-to-the-money options, with k close to 1.

The estimated value of l is 0.0425. In comparison, the average spacing between options is approximately 0.003; thus, the price of a given option is correlated fairly strongly with approximately 10 neighboring options. This is intuitive, given that observed option price surfaces are smooth. This implies that, for the purposes of estimating the integral of $\pi_t(k)$, options at close strike prices are strongly correlated and contain similar information about volatility, so manipulability can be decreased by downweighting options with low liquidity and increasing weights on neighboring options to compensate.

The estimated demand shock variance, $\sigma_d^2(k)$, is very small, only approximately 2% of the variance in $\pi_t(k)$ across strike prices. As I discuss in subsection 7.4, I estimate $\sigma_d^2(k)$ relying on the strong assumption that demand shocks are independent across strike prices; the result is that $\sigma_d^2(k)$ rationalizes the small amount of non-smoothness in

Figure 5: Gaussian process model fit

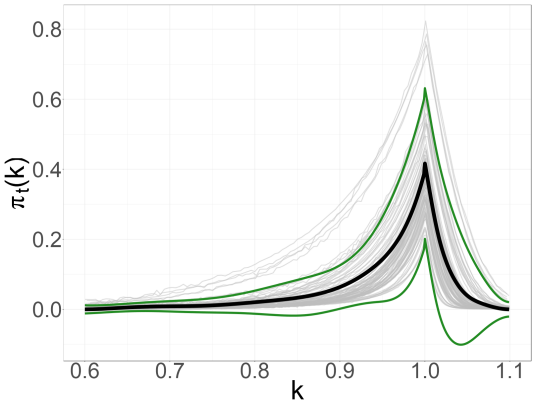

Notes. As in figure 4, each of the 92 gray lines is an observation of $\mathbf{p_t^{H_t}}$ on a single date t. The black line is the function $\mu(k)$ and the green lines are $\mu(k) \pm 2\sqrt{\mathbf{\Sigma_{p,t}^{H_t}}(k,k)}$; that is, they are pointwise intervals that $p_t(k)$ should fall into approximately 95% of the time.

observed option prices $p_t(k)$, visible in figure 4. The exercise of comparing manipulation to demand variance is thus not very meaningful in this setting, as demand variance is negligible, and as I will show, manipulation variance is substantial.

7.4.2 Estimating demand slopes $\boldsymbol{\kappa_t^{H_t}}(k_{th})$

I use iVolatility data on minute-level bid and ask price and volume, for all VIX settlement dates in 2016 and 2017, to estimate $\boldsymbol{\kappa_t^{H_t}}(k_{th})$.[87] In the application to gold in section 6, I observe auction bids, hence I can estimate the slope of aggregate demand directly using bidding data. I do not have SPX opening auction data, hence cannot apply a similar approach here. Instead, I estimate $\boldsymbol{\kappa_t^{H_t}}(k_{th})$ using top-of-book bid and ask prices and volumes, as if they represented auction data. For each option on each trading day 31 days before VIX expiration, I calculate the average bid plus ask volume, and the volume-weighted average bid-ask spread, through the course of the trading day; I treat these as if they were observations of the aggregate demand curve in an auction between $n = 10$ symmetric participants, allowing me to use their ratio to measure the slope of aggregate demand at strike k_{th}. This is likely to underestimate the slope of aggregate demand and

[87]The iVolatility data goes further back, but estimated demand slopes are very low in 2015 and earlier dates, suggesting data issues; thus I focus on 2016 and 2017.

Figure 6: Estimated demand slopes

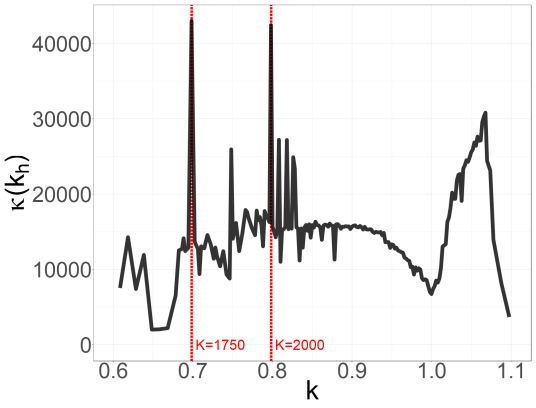

Notes. Estimated slope of demand $\kappa_{dt}^{H_t}(k_{th})$ versus strike price k_{th} for the single date 2017-09-19. Details of estimation are discussed in appendix E.8.

thus overestimate price impact, since market depth implied by limit order book quotes tends to be lower than market depth implied by other methods, such as price impact regressions (Hasbrouck, 2007, ch. 9).

Further details of $\kappa_t^{H_t}(k_{th})$ estimation are described in appendix E.8. Figure 6 shows the estimated $\kappa_t^{H_t}(k_{th})$ vector, for the date 2017-09-19. The function is relatively non-smooth; as the two dotted red lines show, $\kappa_t^{H_t}(k_{th})$ is discontinuously higher at round-number strike prices. This seems to be because because market participants to favor trading round-number strike prices, so bid-ask spreads are lower and bid and ask volume is higher at round-number strike prices. Another observation is that estimated aggregate demand slopes are in fact lower for close-to-the-money options, with k near 1, compared to options with strike prices somewhat further from 1. As appendix E.8 shows, this is because average prices of close-to-the-money options, with $k \approx 1$, are much higher than those of far-out-of-money options; while bid-ask spreads in percentage terms are much smaller for close-to-the-money options, bid-ask spreads in dollar terms are actually higher, so I infer the slope of demand $\kappa_t^{H_t}(k_{th})$ to be somewhat lower for close-to-the-money options.

7.4.3 Estimating VIX position size

Unlike gold futures contracts, there is no position limit for VIX futures contracts.[88] I will instead calibrate contract position size σ_c^2 using data on VIX futures open interest on VIX settlement dates. The CBOE website has data on VIX futures open interest on settlement dates from January 2016 to December 2017; I set the standard deviation of the manipulator's contract positions, $\sqrt{\sigma_c^2}$ equal to 10% of VIX futures open interest at settlement. Since VIX contracts are based on 100 times the square root of the V_t^2, and since contract multipliers for SPX options and VIX futures differ, there are a number of steps required to convert VIX futures contract positions into contracts on the price benchmark P_b; I describe details of the conversion in appendix E.9. The result is that I set:

$$\sqrt{\sigma_c^2} = 14,450,248$$

7.5 Posterior inference on the V_t^2 integral

Any vector of observed prices $\mathbf{p_t^{H_t}}$, together with the Gaussian process prior for $\pi_t(k)$, generates a posterior distribution over the function $\pi_t(k)$. A natural class of estimators for the integral V_t^2 is the expectation of the integral in expression (49) over the posterior for $\pi_t(k)$; that is,

$$E\left[\int_{0.6}^{1.1} \pi_t(k)\, dk \mid \mathbf{p_t^{H_t}}\right] \tag{57}$$

In the literature, the problem of estimating the integral of a Gaussian process posterior is referred to as a *Bayesian quadrature* problem. I briefly discuss background and related literature on Bayesian quadrature in appendix E.10. Expression (57) appears to be a complex object, but because the integral is a linear operator, the expectation of the integral is simply the integral of the posterior expectation of $\pi_t(k)$ given $\mathbf{p_t^{H_t}}$; that is:

$$E\left[\int_{0.6}^{1.1} \pi_t(k)\, dk \mid \mathbf{p_t^{H_t}}\right] = \int_{0.6}^{1.1} E\left[\pi_t(k) \mid \mathbf{p_t^{H_t}}\right] dk \tag{58}$$

The posterior expectation can then be estimated on a grid of points using standard Bayesian inference formulas, and then the RHS of expression (58) can be approximated

[88] As discussed in the VIX futures Contract Specifications, market participants who hold large contract positions are subject to position accountability requirements, but not to explicit position limits.

using a Riemann sum. Fix a uniformly spaced grid of evaluation points $\{k_1 \dots k_G\}$ on the interval $[0.6, 1.1]$, and define the sum vector:

$$\mathbf{S}^{G} \equiv \left(\frac{1.1 - 0.6}{G}\right) \mathbf{1}^{G}$$

In appendix E.11, I show that the Riemann sum approximation for posterior integral expression (57) on the grid $\{k_1 \dots k_G\}$ takes the following form:

$$\mathsf{E}\left[\int_{0.6}^{1.1} \pi_t(k)\, dk \mid \mathbf{p}_t^{H_t}\right] \approx \left(\mathbf{S}^{G}\right)^{\top} \left(\boldsymbol{\mu}^{G} - \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1} \boldsymbol{\mu}^{H_t} + \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1} \mathbf{p}_t^{H_t}\right) \quad (59)$$

where $\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}$ is a rectangular matrix of covariances of $\mathbf{p}_t^{H_t}, \boldsymbol{\pi}_t^{G}$, with elements:

$$\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}(g,h) = \sigma_\pi(k_{th})\, \sigma_\pi(k_g) \exp\left(-\frac{(k_{th} - k_g)^2}{2l^2}\right)$$

Expression (59) can be interpreted as follows. The right term in brackets in expression (59),

$$\boldsymbol{\mu}^{G} - \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1} \boldsymbol{\mu}^{H_t} + \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1} \mathbf{p}_t^{H_t}$$

calculates the posterior expectation vector $\mathsf{E}\left[\boldsymbol{\pi}_t^{G} \mid \mathbf{p}_t^{H_t}\right]$ on the grid G, and the sum vector $\left(\mathbf{S}^{G}\right)^{\top}$ takes the Riemann sum over the posterior mean. The matrix $\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}$ multiplying $\mathbf{p}_t^{H_t}$ is a Bayesian linear regression projection matrix, which determines how much the posterior mean vector $\mathsf{E}\left[\boldsymbol{\pi}_t^{G} \mid \mathbf{p}_t^{H_t}\right]$ is updated, given $\mathbf{p}_t^{H_t}$. Since the matrix $\boldsymbol{\Sigma}_{p,t}^{H_t}$ contains information about noise in observations of $\mathbf{p}_t^{H_t}$, the inverse $\left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}$ lowers the projection weights on values of $\mathbf{p}_t^{H_t}$ which are noisily observed. The terms $\boldsymbol{\mu}^{G} - \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1} \boldsymbol{\mu}^{H_t}$ shrink the estimate of $\mathsf{E}\left[\boldsymbol{\pi}_t^{G} \mid \mathbf{p}_t^{H_t}\right]$ towards the prior mean function, $\mu(k)$.

Expression (59) is an affine function of submarket prices, so it falls in the class of estimators described by expression (50) above. The submarket weights are the coefficient vector on $\mathbf{p}_t^{H_t}$ in expression (59), which is:

$$\left(\boldsymbol{\phi}_t^{H_t}\right)^T = \left(S^G\right)^T \left(\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\right) \tag{60}$$

Expression (60) describes Bayesian quadrature weights for any fixed variance structure $\boldsymbol{\Sigma}_{p,t}^{H_t}$. However, Bayesian quadrature under manipulation is a fixed-point problem. The submarket weight vector $\boldsymbol{\phi}_t^{H_t}$ depends on the covariance structure of $\mathbf{p}_t$, $\boldsymbol{\Sigma}_{p,t}^{H_t}$, through expression (60); simultaneously, $\boldsymbol{\phi}_t^{H_t}$ itself affects the extent of manipulation, changing the covariance structure $\boldsymbol{\Sigma}_{p,t}^{H_t}$, through expression (56). I define *manipulation-consistent Bayesian quadrature weights* as a weight vector $\boldsymbol{\phi}_t^{H_t}$ which simultaneously satisfies both conditions.

Definition 4. A vector of weights $\boldsymbol{\phi}_t^{H_t}$ are *manipulation-consistent Bayesian quadrature weights*, together with covariance structure $\boldsymbol{\Sigma}_{p,t}^{H_t}$, if they satisfy:

$$\left(\boldsymbol{\phi}_t^{H_t}\right)^T = \left(S^G\right)^T \left(\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t} \left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\right) \tag{61}$$

$$\boldsymbol{\Sigma}_{p,t}^{H_t}\left(h, h'\right) = \sigma_\pi\left(k_{th}\right)\sigma_\pi\left(k_{th'}\right)\exp\left(-\frac{\left(k_{th}-k_{th'}\right)^2}{2l^2}\right) + \delta_{hh'}\sigma_d^2\left(k_{th}\right) + \left(\frac{\boldsymbol{\phi}_t^{H_t}\left(k_{th}\right)}{n\left(n-2\right)\boldsymbol{\kappa}_t^{H_t}\left(k_{th}\right)}\right)\left(\frac{\boldsymbol{\phi}_t^{H_t}\left(k_{th'}\right)}{n\left(n-2\right)\boldsymbol{\kappa}_t^{H_t}\left(k_{th'}\right)}\right)\sigma_c^2 \tag{62}$$

In words, Definition 4 says that weights $\boldsymbol{\phi}_t^{H_t}$ must be the correct Bayesian quadrature weights under the covariance structure $\boldsymbol{\Sigma}_{p,t}^{H_t}$ that they induce. For any given collection of primitives, I numerically find a vector of weights $\boldsymbol{\phi}_t^{H_t}$ which satisfied Definition 4. Details of the solution algorithm are discussed in appendix E.12. I could not prove that there is a unique vector $\boldsymbol{\phi}_t^{H_t}$ which satisfies expressions (61) and (62) of Definition 4 for any given primitives, but in practice, for the parameter settings I use, the final value of $\boldsymbol{\phi}_t^{H_t}$ does not seem to be affected by the choice of starting point.

7.6 Estimated Bayesian quadrature weights

I calculate manipulation-consistent Bayesian quadrature weights for all dates in 2016 and 2017. For simplicity, I will henceforth refer to manipulation-consistent Bayesian quadrature weights simply as "Bayesian quadrature weights" or "BQ weights". To

illustrate the qualitative behavior of these weights, figure 7 shows estimated quadrature weights on 2017-09-19. Bayesian quadrature weights depend on contract position size, σ_c^2; I calculate three sets of weights for each date. First, I assume $\sqrt{\sigma_c^2} = 0$, so there is no manipulation. Second, I set $\sqrt{\sigma_c^2} = 1,445,024.8$, which is 10% of the calibrated standard deviation of contract positions from subsection 7.4.3, to generate weights which reflect relatively low levels of manipulation; this serves to illustrate how the Bayesian quadrature methodology works under relatively low levels of manipulation-induced noise. Third, I set $\sqrt{\sigma_c^2} = 14,450,248$, which is the full calibrated standard deviation of contract positions.

The top panel of figure 7 shows the two main drivers of BQ weights. The purple line, labelled "signal variance", shows the diagonal elements of the $\mathbf{\Sigma}_{\pi,t}^{\mathbf{H}_t}$ matrix, or the marginal variance of the $\pi_t(k)$ as a function of k. Signal variance is highest at near-the-money strike prices, $k \approx 1$, and decreases rapidly for far-out-of-money options, with k far away from 1. Intuitively, this is because, from figure 4, the variance in normalized option prices $\mathbf{p}_t^{\mathbf{H}_t}(k_{th})$ is higher for options with strike prices k_{th} close to 1. Optimal benchmark weights should thus be higher for strike prices k close to 1, as these are most informative about the V_t^2 integral. The blue line shows the vector of estimated demand slopes, $\boldsymbol{\kappa}_t^{\mathbf{H}_t}(k_{th})$, at different strike prices. Manipulation-induced price variance is lower at strike pries with higher demand slopes, so manipulation-consistent benchmark weights should be higher at these strike prices. Thus, we expect manipulation-consistent benchmark weights to be correlated with both the blue and purple lines in the top panel of figure 4.

The bottom panel of figure 7 plots four different sets of quadrature weights. The blue line shows quadrature weights calculated using the CBOE VIX formula, which approximates the integral in expression (46) using a Riemann sum:

$$\frac{2}{T}e^{RT}\sum_i \frac{o(k_{th})}{k_{th}^2}\Delta k_{th} \tag{63}$$

where Δk_{th} is the average distance between strike price k_{th} and its immediate neighbords. The CBOE weights are not smooth because spacing between strike prices is uneven, so the Riemann sum formula (63) assigns higher weight to options with further neighboring strike prices. Options with normalized strike prices prices far away from 1 are spaced further apart, so the CBOE VIX weights are in fact higher for far-out-of-money options. Thus, the CBOE VIX estimator is highly vulnerable to manipulation because it assigns high benchmark weight to many options with low signal value and high manipulation

variance.

The red line shows Bayesian quadrature weights in the absence of manipulation; I will refer to these as "no-manipulation BQ weights". BQ weights roughly have the same shape and magnitude as the CBOE VIX weights, suggesting that the BQ methodology is not artificially distorting quadrature weights. The main difference is that BQ weights are smoother than Riemann sum weights; this is because the estimated Gaussian process length-scale parameter l is relatively large, implying that the function $\pi_t(p)$ is smooth, so prices of close options will be highly correlated, and thus it is not necessary for the weight vector to vary at high frequency.

The yellow line shows BQ weights when $\sqrt{\sigma_c^2}$ is set to $1,445,024.8$, 10% of the full calibrated size of contract positions, to illustrate how BQ weights behave when manipulation variance is relatively low. These weights are less smooth than no-manipulation BQ weights, and put low weight at values of k far away from 1. Looking at the top panels of figure 4, the 10% BQ weights are strongly correlated with the vector $\boldsymbol{\kappa}_t^{H_t}(k_{th})$; as expected, manipulation-sensitive weights are higher for strike prices with higher slopes of demand.

The green line shows BQ weights, setting $\sqrt{\sigma_c^2}$ to $14,450,248$, which is 100% of its calibrated value. These weights resemble the 10% BQ weights, but assign much lower mass to values of k far away from 1.[89] Figure 7 thus shows that Bayesian quadrature qualitatively behaves as we would expect a noise-sensitive weighting rule to behave. In the absence of noise, BQ weights are very similar to Riemann sum weights. When manipulation noise is taken into account, BQ weights are correlated with both the blue and purple lines in the top panel of figure 4: they are higher at strike prices with high signal value and high demand slopes.

7.7 Signal and manipulation variance under alternative volatility estimators

For any weight vector $\boldsymbol{\phi}_t^{H_t}$, I can calculate total variance in the resultant benchmark P_b, and decompose it into components attributable to signal variance – that is, shifts in the

[89]Since the BQ estimator is Bayesian, the low weights do not imply that prices at these strike prices are biased towards 0 – the estimator effectively "shrinks" prices at these strikes towards the historical prior level of prices. Bayesian shrinkage allows the estimator to adjust sensitivity to options at different strikes, without systematically biasing P_b upwards or downwards relative to V_t^2; P_b is always biased, but it is biased towards the historical prior value of V_t^2.

Figure 7: Signal variance, demand slopes, and benchmark weights

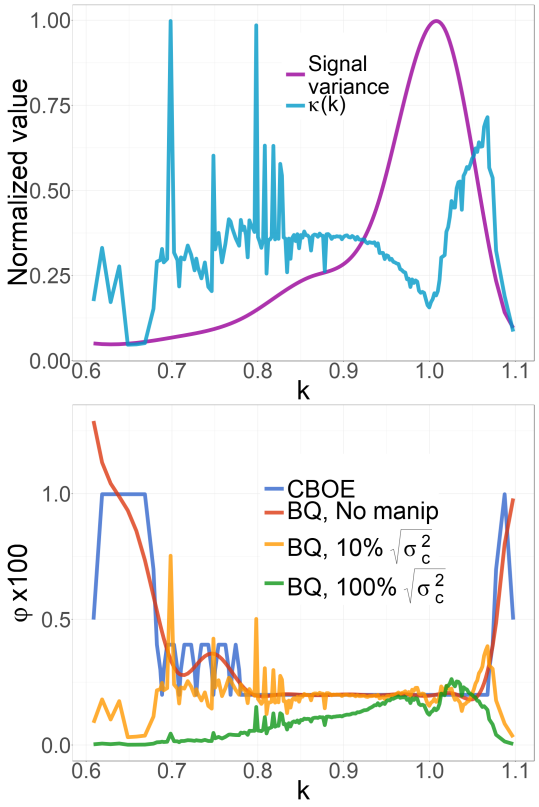

Notes. Both plots show quantities for a single date, 2017-09-19. The top panel shows signal variance, which is the diagonal of the $\mathbf{\Sigma}_{\pi,t}^{H_t}$ matrix, and the estimated slope of demand $\boldsymbol{\kappa}_t^{H_t}(k_{th})$. Both series are normalized so the maximum value is equal to 1. The bottom panel shows different sets of quadrature weights. "CBOE" (blue) represents Riemann sum weights, calculated using expression (63). "BQ, No manip" (red) represents Bayesian quadrature weights assuming no manipulation, so $\sigma_c^2 = 0$. The yellow line represents BQ weights for $\sqrt{\sigma_c^2} = 1,445,024.8$, 10% of the full calibrated size of contract positions, and the green line represents BQ weights for $\sqrt{\sigma_c^2} = 14,450,248$, the full calibrated size of contract positions.

fundamental option price function $\pi_t(k)$ – and manipulation-induced noise.[90] Specifically, from expressions (50) and (56), benchmark variance is:

$$Var(P_b) = \left(\phi_t^{H_t}\right)^T \Sigma_{p,t}^{H_t} \phi_t^{H_t} = \underbrace{\left(\phi_t^{H_t}\right)^T \Sigma_{\pi,t}^{H_t} \phi_t^{H_t}}_{Signal} + \underbrace{\left(\phi_t^{H_t}\right)^T \Sigma_{manip,t}^{H_t} \phi_t^{H_t}}_{Manipulation} \tag{64}$$

The "signal" term, $\left(\phi_t^{H_t}\right)^T \Sigma_{\pi,t}^{H_t} \phi_t^{H_t}$, measures how sensitive the benchmark P_b is is to movements in $\pi_t(k)$. Since Bayesian estimators shrink estimates towards the prior, the signal variance is lower under a Bayesian estimator than under a non-Bayesian estimator such as the actual VIX. The "manipulation" term, $\left(\phi_t^{H_t}\right)^T \Sigma_{manip,t}^{H_t} \phi_t^{H_t}$, measures the amount of variance in P_b caused by manipulation. We wish to design volatility estimators which have high signal variance and low manipulation variance.

Table 4 shows the size of the two components of variance, under three different sets of weights: the CBOE weights, no-manipulation BQ weights, and 100% BQ weights.[91] I average signal and manipulation variance over 21 settlement dates, dropping 3 dates where manipulation variance is unusually high.[92]

Under CBOE weights, shown in the first row of the table, the square root of signal variance is approximately 1, and the square root of manipulation variance is approximately 4.74. Manipulation variance is thus 2254% of signal variance; that is, under my parameter estimates, the vast majority of variance in the CBOE VIX is predicted to be caused by manipulation. Under the no-manipulation BQ weights, shown in the second row of the table, results are similar – signal variance is essentially the same as under the CBOE weights, and manipulation variance is even larger.

Under the manipulation-consistent BQ weights, shown in the third row, signal sensitivity decreases approximately 30%, to 68.8% of the original signal variance under CBOE weights; however, manipulation variance decreases approximately 140 times, to only 16% of signal variance under CBOE weights. Thus, a simple redesign of the CBOE VIX, assigning less weight to far-out-of money options, can dramatically reduce the manipulability of VIX, while preserving most of its signal sensitivity.

[90] Under my estimates, demand shocks are negligible, so I disregard demand shock variance in the following analysis.

[91] The 10% BQ weights are only calculated for illustrative purposes, and I do not show them in table 4.

[92] This does not substantially affect results – my estimator improves the signal-to-noise ratio even more if outlier dates are included in estimation.

Table 4: Signal and manipulation variance under different benchmark weights

Type	Signal SD	$\frac{Sig\ var}{CBOE\ sig\ var}$	Manip SD	$\frac{Manip\ var}{CBOE\ sig\ var}$
CBOE VIX	1.00	100.0%	4.74	2254%
BQ, No manip	1.00	100.1%	5.45	2978%
BQ, Manipulation	0.83	68.8%	0.40	16%

Notes. Rows describe which weight vector is being used. Column 1 shows the standard deviation of signal variance, the square root of the left term in expression (64), for each of the three weight vectors. Column 2 shows the ratio of signal variance under the given weight vector to signal variance under CBOE weights (the square of the first row of the first column). Column 3 shows the standard deviation of manipulation variance, the square root of the right term in expression (64). Column 4 shows the ratio of manipulation variance to CBOE-weights signal variance – that is, the ratio of the square of the third column to the square of the first row of the first column. All quantities are calculated as averages of values for all dates from 2016-2017, excluding three outlier dates where manipulation variance is unusually high.

8 Discussion

8.1 Additional theoretical results

The online appendix contains a number of additional theoretical extensions of my results, generalizing and extending some of the qualitative intuitions in the main text. Section 1 of the online appendix shows that proposition 1, which shows that the pass-through of contract positions into bid curves is $\frac{\kappa}{\kappa+d}$, approximately generalizes to arbitrary smooth nonlinear utility and residual supply functions. I then show how utility functions and manipulation incentives can be nonparametrically estimated using auction bidding data. Section 2 of the online appendix endogenizes agents' contract positions, assuming that agents have exogeneously determined exposures to π, and hedge against uncertainty in π by writing contracts based on auction prices. I show how manipulation decreases welfare by reducing the effectiveness of auction prices for hedging, as well as decreasing allocative efficiency in the auction.

Section 3 of the online appendix shows that, if agents' values are interdependent rather than private, agents will shade bids more in equilibrium; this decreases the slope of residual supply and increases manipulation incentives. However, as the size

of contract positions increases, the extent of adverse selection decreases, increasing the slope of residual supply and decreasing manipulation incentives. Section 4 of the online appendix endogenizes agents' decisions to participate in auctions; I show that increasing contract volume can in fact decrease benchmark variance by increasing auction entry, both from manipulators and non-manipulators. Section 5 of the online appendix shows that manipulation incentives are higher when agents collude; colluding agents' bids are influenced more by each unit of contract they collectively hold, since each bidder internalize the effects of prices on all other bidders' contract positions.

8.2 Mechanisms for benchmark setting

This paper uses uniform-price double auctions as a reduced-form model of price benchmarks. A natural question is the extent to which these results apply to other benchmark-setting mechanisms. Some benchmarks, such as VIX and the LBMA gold price discussed in this paper, are determined using actual auctions. Derivative contracts for many equity indices are also settled based on exchange opening or closing auction prices.[93] Other benchmark-setting mechanisms may produce outcomes similar to uniform-price double auctions. The WM/Reuters FX fixing[94] and the ISDAFIX interest rate swap benchmark[95] (now the ICE swap rate) are set using exchange prices within a few minutes; if agents submitted fixed bid curves at the start of the fixing period and do not adjust bids through the course of the auction, outcomes will coincide with static uniform-price double auction outcomes.[96] Some benchmarks for commodities such as oil and gas are set using volume-weighted average prices in specific geographical locations, over relatively short time spans; if the underlying goods are relatively homogeneous, market outcomes can be approximated by supply function competition between dealers (Klemperer and Meyer, 1989), which is equivalent to a uniform-price double auction.

Other benchmarks are less well approximated by auctions. Some benchmarks are based on trades of underlying assets in markets with large search or transportation frictions. For example, the CME Feeder Cattle Index is based on US-wide cattle trade prices; the price of cattle traded in New York on any given day may differ substantially from the price of cattle

[93] Understanding the Special Opening Quotation (SOQ)

[94] WM/Reuters FX Benchmarks

[95] ICE Swap Rate

[96] Since the WM/Reuters fixing period is only a few minutes long, there is likely little arrival of information during the fixing, and thus no fundamental reason for participants to change their bids over time. However, there may be strategic reasons to adjust bids dynamically; Du and Zhu (2017) argue that the frequency of auctions does in fact affect market outcomes even in the absence of information arrival.

traded in California. Other markets may have nontrivial network structure; a number of recent papers analyze trading networks with central dealers trading with peripheral counterparties (Wang, 2016; Duffie and Wang, 2016). In these markets, the ability of different agents to influence price benchmarks may vary depending on their network positions and resultant bargaining power. Other benchmarks are not based on prices of verifiable trades, but rely on market participants to self-report trades or potential trades; for example, LIBOR is based on banks' announcements of their borrowing costs,[97] and some natural gas benchmarks are based on reports of trades which are rarely verified.[98] Such benchmarks can be manipulated without trading the underlying asset, as agents can simply make reports to move benchmarks in the desired direction. Auctions are potentially a reasonable model for market structures in which the primary distortion is market power. Settings in which there are significant frictions other than market power, such as search frictions, network structure, or concerns about trade falsification, are less well-approximated by my model.

One possible normative interpretation of my results is that auctions are a robust way to determine price benchmarks for contract settlement, and thus should be used more broadly. The analysis of this paper shows that prices from competitive auctions can be used to settle large volumes of derivative contracts, even if they induce relatively small volumes of trade in underlying assets. Auctions have a number of other practical benefits. They are anonymous, so any offer to buy or sell can be taken up by any other agent, so agents cannot distort benchmarks by trading at artificially high or low prices with favored counterparties. They are straightforwards to run; unlike average price benchmarks, they do not require costly infrastructure to track all trades, and thus cannot be manipulated by trade falsification, as is common in markets with average price benchmarks.[99] Regulators could in principle encourage broader use of auctions for benchmark setting; even assets which are usually traded in OTC markets, such as gold, could be traded in periodic auctions in order to determine benchmark prices. However, I have not formally compared auctions to other mechanisms in this paper, so these benefits are still conjectural; future work is needed to formally evaluate the benefits and flaws of different mechanisms for

[97] ICE LIBOR

[98] CFTC Press Release 5409-07

[99] As an example, the CFTC brought charges against many companies for false reporting of natural gas trades to price index compilers; see CFTC Press Release 5300-07.

setting price benchmarks.

8.3 Applications and extensions

My results could be used to analyze other settings in which market prices are used as benchmarks for contract settlement. Duffie (2018*a*), Duffie (2018*b*), and Zhu (2018) have proposed using a double auction to determine the conversion rate between LIBOR and SOFR, and to allow agents to contract to convert LIBOR to SOFR at the auction rate. Contracts written on the auction price give agents incentives to manipulate the auction; however, the results of this paper show that large volumes of contracts can be settled based on small volumes of realized auction trade, so long as slopes of demand are high and the auction is competitive.

There are some derivative contract markets for which the underlying benchmark is not a market price. Credit default swaps are contracts with payouts linked to whether a debtor defaults on her obligation; thus the underlying is not a market with well-defined asset prices, but the whether a single debtor is able to remain solvent. Recently, there have been a number of cases of CDS manipulation, in which funds holding CDS contracts on a given firm's debt extends credit to the firm, with the condition that the firm must make nominally late payments on their debt. This technically causes the firm to default on its debt, triggering the payout on the CDS contract (Levine, 2017). Another example is sports betting, in which the underlying is a sports contest: match fixing can be thought of as manipulation by bribing athletes to underperform. In prediction markets, derivative contracts are written on events such as the weather,[100] election results, and even terrorist attacks and assassinations,[101] all of which are in principle manipulable.

My results show that futures contracts written on market prices are fairly difficult to manipulate; it is not clear that contracts based on other underlying events are similarly robust. The cost of bribing a single athlete to throw a match, or to orchestrate a terrorist attack to profit on bets made in prediction markets, does not clearly decrease as the number of competing manipulators becomes large. Roughly speaking, the robustness of markets as underlying events seems to arise from the a kind of "smoothness" property which allows competition to mitigate manipulation – any purchases made by one agent to raise prices can be counteracted by sales made by other agents. Other underlying events seem not to have this property, and an interesting direction for future research would be

[100] US Monthly Weather Cooling Degree Day (CDD) Contract Specs
[101] CNN: Amid furor, Pentagon kills terrorism futures market

to identify broader classes of underlying events for which competition can mitigate the effects of manipulation.

9 Conclusion

> *"Market power is a structural issue to be remedied, not by behavioral prohibitions"*
>
> – 114 FERC ¶61,165

Contract market manipulation is a result of market power in underlying markets, and it should be regulated using a primarily structural approach. Imperfectly competitive industries are not regulated by prosecuting firms on a case-by-case basis for pricing above marginal cost; regulators instead aim to maintain sufficiently high levels of industry competition that the market disciplines firms' day-to-day pricing behavior, without the need for case-by-case regulatory intervention. It is an empirical question whether a given market is sufficiently concentrated to merit regulatory intervention, and economists have developed quantitative models of equilibrium outcomes under imperfect competition to predict the likely impacts of market power in any given setting. Simple metrics such as top 4 shares and Herfindahl-Hirschman indices can be used as first-pass, diagnostic indicators to monitor industry competitiveness. If an industry is concentrated enough to warrant further analysis and possible policy intervention, more precise methods, such as differentiated-product demand systems, can be used to predict firm markups and welfare losses from market power. These metrics and methods are crucial for informing the application of policy tools, such as blocking mergers and forcing divestitures, for structurally controlling industry competitiveness.

Regulators have many tools for controlling the structure of derivative contract markets, but lack theory-based metrics and methods for quantifying contract market manipulability. The unique two-layered architecture of derivative contract markets precludes the use of standard measures of industry competitiveness. Derivative contract markets tend to be much larger than underlying markets, so a naive approach limiting contract markets to the size of underlying markets would effectively regulate all markets out of existence. What is needed is a theoretical framework which allows derivative contract markets to be much larger than underlying markets, but which can tell when contract markets are too large.

This paper attempts to construct such a framework. I show that manipulation incentives are small in competitive markets, so contract markets can be much larger than

underlying markets without creating unduly large incentives for manipulation, but they cannot be too large. The theory is quantitatively precise: agents' manipulation incentives in equilibrium are approximately equal to their capacity shares, which can be estimated or approximated in many markets. I develop a simple manipulation index, which can be calculated using very limited data on contract markets; roughly speaking, the manipulation index implies that contract markets can be order $\frac{1}{s_{max}}$ times as large as underlying markets before manipulation begins to substantially increase benchmark variance. The manipulation index can play a role similar to that of the Herfindahl-Hirschman index in antitrust: it is a simple metric which, while imperfect, can be used as a diagnostic tool to find markets which may be vulnerable to manipulation. I also show that manipulation-induced bias and variance in price benchmarks can be estimated using data on agents' demand slopes and contract positions; these can also be measured or approximated by regulators in many contract markets. This paper has used these tools to study a number of empirical settings; a natural direction for further research is to apply these tools to predict the size of manipulation-induced distortions in other contract markets.

I do not take a stance on how regulators should optimally intervene in manipulable contract markets; I cannot address this question, as I do not explicitly model the welfare effects of contract markets in this paper. A natural question would be to determine the welfare-optimal way to limit manipulation in contract markets, and to see whether this is well-approximated by existing tools such as contract position limits. Another limitation of this paper is that I take a comparative statics approach, measuring manipulation-induced distortions fixing the set of market participants and the distribution of their contract positions. This approach is appropriate for determining whether regulators should intervene in a market given current equilibrium outcomes. However, regulation also has long-run effects as agents enter and markets evolve in response to policy interventions, and modelling the dynamic evolution of contract markets may be important for understanding these longer-term effects.

These are important directions for future work, but are beyond the scope of the current paper. This paper constructs metrics and methods to quantify the manipulability of derivative contract markets; it is hoped that these metrics can be used as the basis for a structural, theory-based approach to regulating manipulation in derivative contract markets.

References

Abrantes-Metz, Rosa M., Michael Kraten, Albert D. Metz, and Gim S. Seow. 2012. "Libor manipulation?" *Journal of Banking & Finance*, 36(1): 136–150.

Andersen, Torben G., Oleg Bondarenko, and Maria T. Gonzalez-Perez. 2015. "Exploring Return Dynamics via Corridor Implied Volatility." *The Review of Financial Studies*, 28(10): 2902–2945.

Aspris, Angelo, Sean Foley, Fergus Grattan, and Peter O'Neill. 2015. "Towards a New Fix: Assessing the new FIX regimes for metals trading." Working paper.

Baer, Julius B., and Olin Glenn Saxon. 1949. *Commodity Exchanges And Futures Trading*. Harper And Brothers Publishers.

Baldauf, Markus, Christoph Frei, and Joshua Mollner. 2018. "Contracting for Financial Execution." Social Science Research Network SSRN Scholarly Paper ID 3177283, Rochester, NY.

Carr, Peter, and Liuren Wu. 2006. "A Tale of Two Indices." *The Journal of Derivatives*, 13(3): 13–29.

CFTC. 2008. "Staff Report on Commodity Swap Dealers & Index Traders with Commission Recommendations."

Chen, Yiling, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. 2010. "Gaming Prediction Markets: Equilibrium Strategies with a Market Maker." *Algorithmica*, 58(4): 930–969.

Coulter, Brian, Joel Shapiro, and Peter Zimmerman. 2018. "A Mechanism for LIBOR." *Review of Finance*, 22(2): 491–520.

Demeterfi, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou. 1999. "A Guide to Volatility and Variance Swaps." *The Journal of Derivatives*, 6(4): 9.

Diaconis, Persi. 1988. "Bayesian Numerical Analysis." In *S.S. Gupta and J. Berger, Eds., Statistical Decision Theory and Related Topics IV, Vol. 1.*. 163–175. New York:Springer-Verlag.

Duffie, Darrell. 2018*a*. "Compression Auctions, with an Application to LIBOR-SOFR Swap Conversion."

Duffie, Darrell. 2018*b*. "Notes on LIBOR Conversion."

Duffie, Darrell, and Chaojun Wang. 2016. "Efficient Contracting in Network Financial Markets."

Duffie, Darrell, and Haoxiang Zhu. 2017. "Size Discovery." *The Review of Financial Studies*, 30(4): 1095–1150.

Duffie, Darrell, and Piotr Dworczak. 2018. "Robust Benchmark Design." National Bureau of Economic Research Working Paper 20540.

Duffie, Darrell, Piotr Dworczak, and Haoxiang Zhu. 2017. "Benchmarks in Search Markets." *The Journal of Finance*, 72(5): 1983–2044.

Du, Songzi, and Haoxiang Zhu. 2012. "Ex post equilibria in double auctions of divisible assets." Working paper.

Du, Songzi, and Haoxiang Zhu. 2017. "What is the Optimal Trading Frequency in Financial Markets?" *The Review of Economic Studies*, 84(4): 1606–1651.

Dutt, Hans R., and Lawrence E. Harris. 2005. "Position limits for cash-settled derivative contracts." *Journal of Futures Markets*, 25(10): 945–965.

FCA. 2014. "Barclays fined £26m for failings surrounding the London Gold Fixing and former Barclays trader banned and fined for inappropriate conduct."

Financial Stability Board. 2014. "Final Report on Foreign Exchange Benchmarks."

Financial Stability Board. 2018. "Final Report of the Market Participants Group on Reforming Interest Rate Benchmarks."

Fischel, Daniel R., and David J. Ross. 1991. "Should the Law Prohibit "Manipulation" in Financial Markets?" *Harvard Law Review*, 105(2): 503.

Gandhi, Priyank, Benjamin Golez, Jens Jackwerth, and Alberto Plazzi. 2015. "LIBOR Manipulation : Cui Bono?" Working paper.

Garbade, Kenneth D., and William L. Silber. 1983. "Cash settlement of futures contracts: An economic analysis." *Journal of Futures Markets*, 3(4): 451–472.

Ghahramani, Zoubin, and Carl E. Rasmussen. 2003. "Bayesian monte carlo." 505–512.

Griffin, John M, and Amin Shams. 2018. "Manipulation in the VIX?" *The Review of Financial Studies*, 31(4): 1377–1417.

Hanson, Robin, and Ryan Oprea. 2008. "A Manipulator Can Aid Prediction Market Accuracy." *Economica*, 76(302): 304–314.

Hasbrouck, Joel. 2007. *Empirical market microstructure: The institutions, economics, and econometrics of securities trading*. Oxford University Press.

Huang, Eric H., and Yoav Shoham. 2014. "Price Manipulation in Prediction Markets: Analysis and Mitigation." *AAMAS '14*, 213–220. Richland, SC:International Foundation for Autonomous Agents and Multiagent Systems.

Jones, Frank J. 1982. "The economics of futures and options contracts based on cash settlement." *Journal of Futures Markets*, 2(1): 63–82.

Klemperer, Paul D., and Margaret A. Meyer. 1989. "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica*, 57(6): 1243–1277.

Kumar, Praveen, and Duane J. Seppi. 1992. "Futures Manipulation with "Cash Settlement"." *The Journal of Finance*, 47(4): 1485–1502.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica*, 53(6): 1315–1335.

Leising, Matthew. 2017. "Rate Benchmark Scandal Hits $570 Million in Fines as RBS Settles." *Bloomberg.com*.

Levine, Matt. 2014. "Banks Manipulated Foreign Exchange in Ways You Can't Teach." *Bloomberg.com*.

Levine, Matt. 2015. "Some Banks Paid Some More Fines for FX Rigging." *Bloomberg.com*.

Levine, Matt. 2017. "Blackstone May Do Its Cleverest CDS Trade Again." *Bloomberg.com*.

Louis, Brian, and Nikolaj Gammeltoft. 2018. "VIX Manipulation Costs Investors Billions, Whistle-Blower Says." *Bloomberg.com*.

Malamud, Semyon, and Marzena Rostek. 2017. "Decentralized Exchange." *American Economic Review*, 107(11): 3320–3362.

Markham, Jerry. 2014. *Law Enforcement and the History of Financial Market Manipulation.* . 1 edition ed., Armonk, N.Y:Routledge.

O'Hagan, A. 1991. "Bayes–Hermite quadrature." *Journal of Statistical Planning and Inference*, 29(3): 245–260.

Oprea, Ryan, David Porter, Chris Hibbert, Robin Hanson, and Dorina Tila. 2008. "Can Manipulators Mislead Prediction Market Observers?" *ESI Working Papers*.

Paul, Allen B. 1985. "The Role of Cash Settlement in Futures Contract Specification." 58.

Perdue, Wendy Collins. 1987. "Manipulation of Futures Markets: Redefining the Offense." *Fordham Law Review*, 56: 59.

Poincaré, Henri. 1896. *Calcul des probabilités.* Paris:Georges Carré.

Rasmussen, Carl Edward, and Christopher K. I. Williams. 2005. *Gaussian Processes for Machine Learning*. Cambridge, Mass:The MIT Press.

Reuters Staff. 2014. "London gold-fix banks accused of manipulation in U.S. lawsuit." *Reuters*.

Ridley, Kirstin, and Karen Freifeld. 2015. "Deutsche Bank fined record $2.5 billion over rate rigging." *Reuters*.

Rostek, Marzena, and Marek Weretka. 2015. "Dynamic Thin Markets." *The Review of Financial Studies*, 28(10): 2946–2992.

Tirole, Jean. 1988. *The theory of industrial organization*. MIT press.

Vayanos, Dimitri. 1999. "Strategic Trading and Welfare in a Dynamic Market." *The Review of Economic Studies*, 66(2): 219–254.

Wang, Chaojun. 2016. "Core-Periphery Trading Networks."

Wilson, Robert. 1979. "Auctions of Shares." *The Quarterly Journal of Economics*, 93(4): 675–689.

Wittwer, Milena. 2017. "Connecting Disconnected Markets? An Irrelevance Result." Social Science Research Network SSRN Scholarly Paper ID 2910113, Rochester, NY.

Wittwer, Milena. 2018. "Interconnected Pay-As-Bid Auctions." Social Science Research Network SSRN Scholarly Paper ID 3240230, Rochester, NY.

Zhu, Haoxiang. 2018. "From LIBOR to SOFR: A Multi-Maturity Clock Auction Design."

Appendix

A Supplementary material for section 3

A.1 Proof of proposition 1

Assume residual supply is:

$$z_{RSi}(p,\eta) = d_i(p-\pi)+\eta_i$$

Define $p^*(\eta)$ as the optimal choice of p for any given η_i, that is:

$$p^*(\eta_i) \equiv \arg\max_p U_i(z_{RSi}(p,\eta_i), p; y_{di}, y_{ci})$$

$$= \arg\max_p \pi z_{RSi}(p,\eta_i) + \frac{y_{di} z_{RSi}(p,\eta_i)}{\kappa_i} - \frac{z_{RSi}(p,\eta_i)^2}{2\kappa_i} + y_{ci}p - z_{RSi}(p,\eta_i)\,p$$

Since $z_{RSi}(p,\eta)$ is affine and increasing in p, the objective function concave in p, thus the first-order condition is necessary and sufficient for $p^*(\eta_i)$ to be optimal. Taking the derivative with respect to p and setting to 0, and using that $z'_{RSi}(p,\eta_i) = d_i$, we have:

$$\frac{y_{di}d_i}{\kappa_i} - \frac{z_{RS}(p^*(\eta_i),\eta_i)}{\kappa_i}d_i + y_{ci} - z_{RS}(p^*(\eta_i),\eta_i) - (p^*(\eta_i)-\pi)\,d_i = 0 \tag{65}$$

Hence, any pair $(p^*(\eta_i), z_{RS}(p^*(\eta_i),\eta_i))$ – that is, any point (p,z) which is the agent's optimal choice for some η_i – satisfies (65). Hence, the unique bid curve which passes through the set of all ex-post optimal points is the curve implicitly defined by (65). Solving for $z_{RS}(p^*(\eta_i),\eta_i)$, we get expression (6) of proposition 1.

A.2 Proof of proposition 2

This proof is taken, with minor notational modifications, from Appendix A.4 of Du and Zhu (2012); I reproduce the proof here for completeness.

We seek a vector of demand and residual supply slopes b_i which satisfy the following

equations for all i:

$$d_i = \sum_{j \neq i} b_i = B - b_i \tag{66}$$

$$b_i = \frac{d_i \kappa_i}{\kappa_i + d_i} \tag{67}$$

Rearranging, we have:

$$b_i \kappa_i + b_i d_i = d_i \kappa_i$$

$$d_i (\kappa_i - b_i) = b_i \kappa_i$$

$$d_i = \frac{b_i \kappa_i}{\kappa_i - b_i} \tag{68}$$

Combining (66) and (68), we have:

$$\sum_j b_j - b_i = \frac{b_i \kappa_i}{\kappa_i - b_i}$$

Defining $B \equiv \sum_j b_j$, we have

$$(\kappa_i - b_i)(B - b_i) = b_i \kappa_i$$

This has two solutions:

$$b_i = \frac{B + 2\kappa_i \pm \sqrt{B^2 + 4\kappa_i^2}}{2}$$

In order for $B > b_i$, we must pick:

$$b_i = \frac{2\kappa_i + B - \sqrt{B^2 + 4\kappa_i^2}}{2} \tag{69}$$

This is (11) of proposition 2. B must satisfy:

$$B = \sum_j b_j = \sum_{i=1}^{n} \frac{2\kappa_i + B - \sqrt{B^2 + 4\kappa_i^2}}{2} \tag{70}$$

By multiplying the top and bottom of the RHS by $2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}$ and simplifying,

this becomes:

$$B = \sum_{i=1}^{n} \frac{2\kappa_i B}{2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}}$$

Or,

$$B\left(-1 + \sum_{i=1}^{n} \frac{2\kappa_i}{2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}}\right) = 0 \tag{71}$$

Now, define

$$f(B) = -1 + \sum_{i=1}^{n} \frac{2\kappa_i}{2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}}$$

In order for B to solve (71) when $B > 0$, we need $f(B) = 0$. Now, $f(0) > 0$, $f(B) \to -1$ as $B \to \infty$, and $f'(B) < 0$ for $B > 0$. Hence, $f(B) = 0$ at some unique B, hence there is a unique value of B which solves (71), and thus there is a unique linear equilibrium for any demand slopes $\kappa_1 \ldots \kappa_n$.

From proposition 1, we have that best-response demand functions are:

$$z_{Di}(p; y_{di}, y_{ci}) = \frac{d_i}{\kappa_i + d_i} y_{di} + \frac{\kappa_i}{\kappa_i + d_i} y_{ci} - \frac{\kappa_i d_i}{\kappa_i + d_i}(p - \pi) \tag{72}$$

Substituting $b_i = \frac{\kappa_i d_i}{\kappa_i + d_i}$, we have:

$$z_{Di}(p; y_{di}, y_{ci}) = \frac{b_i}{d_i} y_{di} + \frac{b_i}{\kappa_i} y_{ci} - b_i(p - \pi) \tag{73}$$

Now $d_i = \sum_{j \neq i} b_i$, so (73) is the same as (9).

To find benchmark prices, sum demand and equate to 0:

$$\sum_{i=1}^{n} \left[y_{di} \frac{b_i}{\kappa_i} + y_{ci} \frac{b_i}{\sum_{j \neq i} b_j} - (p_b - \pi) b_i \right] = 0$$

Solving for the benchmark price p_b, we have:

$$p_b - \pi = \frac{1}{\sum_{i=1}^{n} b_i} \left[\sum_{i=1}^{n} \left[\frac{b_i y_{di}}{\kappa_i} + \frac{b_i y_{ci}}{\sum_{j \neq i} b_j} \right] \right]$$

proving (10).

A.3 Proof of proposition 3

Fix $\kappa_1 \dots \kappa_N$, and let

$$s_{max} = \frac{\max_i \kappa_i}{\sum_{i=1}^{n} \kappa_i}$$

In appendix A.4 below, I prove the following claim.

Claim 1. For any demand slopes $\kappa_1 \dots \kappa_N$, for all i, we have:

$$\left(1 - \frac{s_{max}}{1 - s_{max}}\right) \leqslant \frac{b_i}{\kappa_i} \leqslant 1$$

Now, from proposition 1, agents' best-response bid curves are:

$$z_{Di}(p; y_{di}, y_{ci}) = \frac{d_i}{\kappa_i + d_i} y_{di} + \frac{\kappa_i}{\kappa_i + d_i} y_{ci} - \frac{\kappa_i d_i}{\kappa_i + d_i}(p - \pi) \tag{74}$$

We want to bound the difference between this and the approximation from (14):

$$z_{Di}(p; y_{di}, y_{ci}) \approx y_{di} + \frac{\kappa_i}{\sum_{i=1}^{n} \kappa_i} y_{ci} - \kappa_i (p - \pi) \tag{75}$$

Using claim 1, we can bound all three coefficients in the bid curve (74). Recall that $b_i = \frac{\kappa_i d_i}{\kappa_i + d_i}$. Thus, we have:

$$\left(1 - \frac{s_{max}}{1 - s_{max}}\right) \kappa_i \leqslant \frac{\kappa_i d_i}{\kappa_i + d_i} \leqslant \kappa_i$$

$$1 - \frac{s_{max}}{1 - s_{max}} \leqslant \frac{d_i}{\kappa_i + d_i} \leqslant 1$$

Hence, the multiplicative error of the demand slope and the coefficient on y_{di}, comparing the approximation (75) to the exact bid curve (74), is bounded above by

$$1 - \frac{s_{max}}{1 - s_{max}}$$

For the manipulation coefficient, recall that $d_i \equiv \sum_{j \neq i} b_i$. Thus,

$$\frac{\kappa_i}{\kappa_i + d_i} = \frac{\kappa_i}{\kappa_i + \sum_{j \neq i} b_i}$$

Using claim 1, we have:

$$\frac{\kappa_i}{\kappa_i + \sum_{j \neq i} \kappa_j} \geqslant \frac{\kappa_i}{\kappa_i + \sum_{j \neq i} b_i} \geqslant \frac{\kappa_i}{\kappa_i + \left(1 - \frac{s_{max}}{1 - s_{max}}\right) \sum_{j \neq i} \kappa_j} \tag{76}$$

Now, note that:

$$\frac{\kappa_i}{\kappa_i + \left(1 - \frac{s_{max}}{1 - s_{max}}\right) \sum_{j \neq i} \kappa_j} \geqslant \frac{\kappa_i}{\left(1 - \frac{s_{max}}{1 - s_{max}}\right) \left(\kappa_i + \sum_{j \neq i} \kappa_j\right)} = s_i \left(1 + \frac{s_{max}}{1 - 2s_{max}}\right)$$

Hence, (76) becomes:

$$s_i \left(1 + \frac{s_{max}}{1 - 2s_{max}}\right) \leqslant \frac{\kappa_i}{\kappa_i + \sum_{j \neq i} b_i} \leqslant s_i$$

Thus, multiplicative error in the manipulation coefficient is bounded above by $\left(1 + \frac{s_{max}}{1 - 2s_{max}}\right)$.

The approximation bid curve (75) in proposition 3 uses the upper-bound values for the bid slope and the coefficient on y_{di}, and lower-bound values for the coefficient on y_{ci}; hence, the approximation systematically underestimates the manipulation coefficient, although the error is small if s_{max} is small.

A.3.1 Manipulation-induced benchmark price variance

The approximation bounds for bids can also be used to derive an error bound for the approximate expression for manipulation-induced price variance, (18) in subsection 3.6. The exact expression for benchmark prices is:

$$p_b - \pi = \frac{1}{\sum_{i=1}^{n} b_i} \left[\sum_{i=1}^{n} \left[\frac{\kappa_i}{\kappa_i + d_i} y_{ci} + \frac{d_i}{\kappa_i + d_i} y_{di} \right] \right]$$

Thus, the exact contribution of manipulation to price variance is:

$$\left(\frac{1}{\sum_{i=1}^{n} b_i}\right)^2 \left(\frac{\kappa_i}{\kappa_i + d_i}\right)^2 \sigma_{ci}^2 \tag{77}$$

In expression (18) in the main text, we approximate manipulation-induced variance as:

$$\frac{\sum_{i=1}^{n} s_i^2 Var\left(y_{ci}\right)}{\left(\sum_{i=1}^{n} \kappa_i\right)^2} \tag{78}$$

This approximation always underestimates price variance: since we ignore bid shading, the exact manipulation coefficient, $\frac{\kappa_i}{\kappa_i + d_i}$, will always be somewhat larger than s_i, and the exact slope of aggregate demand, $\sum_{i=1}^{n} b_i$, will always be somewhat smaller than $\sum_{i=1}^{n} \kappa_i$. To bound the approximation error, define:

$$\psi \equiv \left(1 - \frac{s_{max}}{1 - s_{max}}\right)$$

Using claim 1, we have:

$$\frac{1}{\sum_{i=1}^{n} b_i} \leqslant \frac{1}{\psi \sum_{i=1}^{n} \kappa_i}$$

$$\frac{\kappa_i}{\kappa_i + d_i} \leqslant \frac{\kappa_i}{\psi \sum_{j=1}^{n} \kappa_j}$$

Thus,

$$\left(\frac{1}{\sum_{i=1}^{n} b_i}\right)^2 \left(\frac{\kappa_i}{\kappa_i + d_i}\right)^2 \sigma_{ci}^2 \leqslant \left(\frac{1}{\psi \sum_{i=1}^{n} \kappa_i}\right)^2 \left(\frac{\kappa_i}{\psi \sum_{j=1}^{n} \kappa_j}\right)^2 \sigma_{ci}^2$$

$$= \frac{1}{\psi^4} \left(\frac{s_i^2 \sigma_{ci}^2}{\left(\sum_{i=1}^{n} \kappa_i\right)^2}\right)$$

Hence, the ratio of the approximate expression for manipulation-induced price variance, expression (78), and its exact value, expression (77), is bounded by ψ^4, that is

$$\left(1 - \frac{s_{max}}{1 - s_{max}}\right)^4$$

A.4 Proof of claim 1

From (11) in proposition 2, we have:

$$b_i = \frac{B + 2\kappa_i - \sqrt{B^2 + 4\kappa_i^2}}{2}$$

This immediately implies that $b_i \leqslant \kappa_i$, as b_i approaches κ_i from below as $B \to \infty$. This proves the upper bound in claim 1,

$$\frac{b_i}{\kappa_i} \leqslant 1$$

Thus, we need only prove the lower bound:

$$\left(1 - \frac{s_{max}}{1 - s_{max}}\right) \leqslant \frac{b_i}{\kappa_i} \tag{79}$$

We will proceed in two stages. Appendix A.4.1 proves claim 2, which states that, given s_{max} and κ_{max}, we can construct an analytical lower bound for B. Appendix A.4.2 then uses the lower bound for B to lower-bound the ratio $\frac{b_i}{\kappa_i}$.

A.4.1 A lower bound for B

This subsection proves the following claim:

Claim 2. Fixing κ_{max}, s_{max}, we have:

$$B \geqslant \frac{2s_{max} - 1}{s_{max}(s_{max} - 1)} \kappa_{max}$$

Proof. From (12) in Proposition 2, given $\kappa_1 \dots \kappa_N$, B satisfies:

$$-1 + \sum_{i=1}^{n} \frac{2\kappa_i}{2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}} = 0 \tag{80}$$

For any $\kappa_1 \dots \kappa_N$, define $B(\kappa_1 \dots \kappa_n)$ as the induced value of B, that is:

$$B(\kappa_1 \dots \kappa_n) = \left\{ B : \ -1 + \sum_{i=1}^{n} \frac{2\kappa_i}{2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}} = 0 \right\}$$

Now, without loss of generality, suppose agent $i = 1$ has the largest demand slope, so that:

$$s_{max} = s_1 = \frac{\kappa_1}{\sum_{i=1}^{n} \kappa_i}$$

Alternatively, we can write this as:

$$\sum_{i=2}^{n} \kappa_i = \left(\frac{1 - s_{max}}{s_{max}} \right) \kappa_1 \tag{81}$$

Consider the following optimization problem:

$$\begin{aligned} &\min_{\kappa_2 \dots \kappa_n} B\left(\kappa_1 \dots \kappa_n\right) \\ &\text{s.t. } \sum_{i=2}^{n} \kappa_i = \left(\frac{1 - s_{max}}{s_{max}} \right) \kappa_1,\ 0 \leqslant \kappa_i \leqslant \kappa_1 \ \forall i \end{aligned} \tag{82}$$

In words, problem (82) states that we choose $\kappa_2 \dots \kappa_n$ to minimize $B\left(\kappa_2 \dots \kappa_n\right)$, fixing κ_1 and s_{max} (which is equivalent to fixing the sum $\sum_{i=2}^{n} \kappa_i$). The minimal value of B from this problem is a lower bound for $B\left(\kappa_1 \dots \kappa_n\right)$ given κ_1 and s_{max}. The following claim characterizes the solution to problem (82).

Claim 3. If $\kappa_2 \dots \kappa_n$ are an optimal solution to (82), then all but 1 element of $\kappa_2 \dots \kappa_n$ must be equal to either κ_1 or 0.

Proof. We use the the implicit function theorem to calculate the derivative $\frac{dB}{d\kappa_i}$. Write expression (80) as:

$$L = -1 + \frac{2\kappa_1}{2\kappa_1 + B + \sqrt{B^2 + 4\kappa_1^2}} + \sum_{i=2}^{n} \frac{2\kappa_i}{2\kappa_i + B + \sqrt{B^2 + 4\kappa_i^2}} = 0 \tag{83}$$

Differentiate (83) with respect to κ_i, to get:

$$\frac{\partial L}{\partial \kappa_i} : \frac{2B\left(B + \sqrt{B^2 + 4\kappa_i^2}\right)}{\sqrt{B^2 + 4\kappa_i^2}\left(B + 2\kappa_i + \sqrt{B^2 + 4\kappa_i^2}\right)^2} \tag{84}$$

which is strictly positive. Differentiate (83) with respect to B to get:

$$\frac{\partial L}{\partial B} : -\sum_{i=1}^{n} \frac{2\kappa_i \left(1 + \frac{B}{\sqrt{B^2 + 4\kappa_i^2}}\right)}{\left(B + 2\kappa_i + \sqrt{B^2 + 4\kappa_i^2}\right)^2}$$

This is strictly negative, so we have:

$$\frac{dB}{d\kappa_i} = -\frac{\frac{\partial L}{\partial \kappa_i}}{\frac{\partial L}{\partial B}} > 0 \tag{85}$$

Hence, B is strictly increasing in κ_i. Differentiating (84) once again in κ_i, we have:

$$\frac{\partial^2 L}{\partial \kappa_i^2} = -\frac{2B}{(B^2 + 4\kappa^2)^{\frac{3}{2}}} < 0$$

Hence, $\frac{\partial L}{\partial \kappa_i}$ is strictly decreasing in κ_i. This implies that, if $\kappa_i > \kappa_j$, then $\frac{\partial L}{\partial \kappa_i} < \frac{\partial L}{\partial \kappa_j}$; from (85), this then implies that $\frac{dB}{d\kappa_i} < \frac{dB}{d\kappa_j}$. Now, suppose we have a vector of demand slopes $\kappa_2 \dots \kappa_n$, such there exists indices i, j such that κ_i, κ_j are strictly between 0 and κ_1, with $\kappa_i \geqslant \kappa_j$. Then, since $\frac{dB}{d\kappa_i} \leqslant \frac{dB}{d\kappa_j}$, the following vector of demand slopes, for sufficiently small δ, lowers the objective value in (82), while maintaining constraint satisfaction:

$$\left(\kappa_2 \dots \kappa_i + \delta, \dots \kappa_j - \delta, \dots \kappa_n\right)$$

In words, this says that, since $\frac{dB}{d\kappa_i}$ is lower for higher κ_i, we can always lower the objective by increasing high values of κ_i and decreasing lower values to keep the sum constant. Thus, if $\kappa_2 \dots \kappa_n$ are optimal for problem (82), there cannot exist two indices i, j such that κ_i and κ_j are both strictly between 0 and κ_1; thus, the solution to (82) must have all but one element of $\kappa_2 \dots \kappa_N$ equal to either κ_1 or 0, proving claim 3. □

The bound from claim 3 is not yet useful, as the optimal value from problem (82) does not admit a simple analytical expression. However, the limit of a sequence of relaxations of problem (82) does yield an analytically tractable solution. I construct a sequence of relaxations of problem (82), parametrized by the integer $h \in \{1, 2, \dots \infty\}$. For any h, let $\kappa_h(x)$ be a function defined on the interval $[2, n+1]$, which is piecewise constant on intervals of length $\frac{1}{h}$. Define $B(\kappa_h(x))$ as:

$$B(\kappa_h(x)) = \left\{ B : -1 + \frac{2\kappa_1}{2\kappa_1 + B + \sqrt{B^2 + 4\kappa_1^2}} + \int_2^n \frac{2\kappa_h(x)}{2\kappa_h(x) + B + \sqrt{B^2 + 4(\kappa_h(x))^2}} dx = 0 \right\} \tag{86}$$

Define the hth minimization problem as:

$$\min_{\kappa_h(x)} B\left(\kappa_h\left(x\right)\right)$$
$$\text{s.t.} \int_2^n \kappa_h\left(x\right) = K\left(1 - s_{max}\right),\ 0 \leqslant \kappa_i \leqslant \kappa_1\ \forall i \tag{87}$$

Effectively, (86) splits each $\kappa_2 \ldots \kappa_n$ value in the original problem (82) into h components, which may have different values of κ_i. To see that problem (87) is a strict relaxation of the original optimization problem (82), note that if we constrain $\kappa_h\left(x\right)$ to be constant on the intervals

$$[2,3), [3,4), \ldots [n, n+1)$$

then both the objective and the constraints in problem (87) reduced to the original problem (82). Thus, any value attainable in the original optimization problem is attainable in the relaxed optimization problem, so the optimal value from problem (87), for any h, is a lower bound for the original problem.

The integral in the relaxed problem (86) is simply a weighted sum over a finite number of values of $\kappa_h\left(x\right)$; thus, as in the original problem, the objective function for the relaxed problem is concave in the value of $\kappa_h\left(x\right)$ on any interval. Thus, claim 3 characterizing the solution to the original problem applies to the relaxed problem for any h: at any optimal solution, $\kappa_h\left(x\right)$ must be equal to either 0 or κ_1 on all intervals but one. Taking the limit as $h \to \infty$, the constraint

$$\int_2^n \kappa_h\left(x\right) = K\left(1 - s_{max}\right)$$

implies that $\kappa_h\left(x\right)$ must be equal to κ_1 on a set with measure arbitrarily close to

$$\frac{1}{s_{max}} - 1$$

and, $\kappa_h\left(x\right) = 0$ otherwise, except for an interval which has measure arbitrarily close to 0. Hence, in the limit as $h \to \infty$, any choice of $\kappa_h\left(x\right)$ which minimizes B has:

$$\int_2^n \frac{2\kappa_h\left(x\right)}{2\kappa_h\left(x\right) + B + \sqrt{B^2 + 4\left(\kappa_h\left(x\right)\right)^2}} dx \to \left(\frac{1}{s_{max}} - 1\right) \frac{2\kappa_1}{2\kappa_1 + B + \sqrt{B^2 + 4\kappa_1^2}}$$

Using (86), in the limit as $h \to \infty$, the minimized value of B thus satisfies:

$$-1 + \frac{2\kappa_1}{2\kappa_1 + B + \sqrt{B^2 + 4\kappa_1^2}} + \left(\frac{1}{s_{max}} - 1\right) \frac{2\kappa_1}{2\kappa_1 + B + \sqrt{B^2 + 4\kappa_1^2}} = 0$$

$$\implies -1 + \frac{1}{s_{max}} \frac{2\kappa_1}{2\kappa_1 + B + \sqrt{B^2 + 4\kappa_1^2}} = 0$$

This can be analytically solved for B, to get:

$$B = \frac{2s_{max} - 1}{s_{max}(s_{max} - 1)} \kappa_1$$

This is thus a lower bound for the optimal value to the original optimization problem (82), hence, this is a lower bound for B given κ_1 and s_{max}. Since we assumed $\kappa_1 = \kappa_{max}$, this proves claim 2. This lower bound is tight whenever $s_{max} = \frac{1}{n}$ for some integer value of n, as it is exactly the equilibrium value of B when there are n agents with identical demand slopes, $\kappa_i = \kappa$. □

A.4.2 Bounding b_i

Now, we use claim 2 to prove the lower bound, (79), of claim 1. We have:

$$b_i = \frac{2\kappa_i + B - \sqrt{B^2 + 4\kappa_i^2}}{2} \tag{88}$$

This shows that b_i is increasing in B. Thus, we can find a lower bound for the ratio $\frac{b_1}{\kappa_1}$ for the largest agent by plugging in the lower bound for B from claim 2 and solving for b_i. This gives:

$$\frac{b_1}{\kappa_1} \geqslant \left(1 - \frac{s_{max}}{1 - s_{max}}\right) \tag{89}$$

To bound the ratio $\frac{b_i}{\kappa_i}$ for all other i, I show that the ratio $\frac{b_i}{\kappa_i}$ is decreasing in κ_i, for B fixed. To see this, differentiate (88), fixing B, to get:

$$\frac{db_i}{d\kappa_i} = 1 - \frac{2\kappa_i}{\sqrt{B^2 + 4\kappa_i^2}}$$

This is decreasing in κ_i, so b_i is a concave function of κ_i fixing B, and $b_i = 0$ when $\kappa_i = 0$. Now,

$$\frac{d}{dx}\frac{b_i(\kappa_i)}{\kappa_i} = \frac{\kappa_i b_i'(\kappa_i) - b_i(\kappa_i)}{\kappa_i^2}$$

This is negative if the numerator is negative; using that $b_i(0) = 0$, write the numerator as:

$$\kappa_i b_i'(\kappa_i) - b_i(\kappa_i) = \kappa_i b_i'(\kappa_i) - \int_0^{\kappa_i} b_i'(\kappa)\, d\kappa$$

$$= \int_0^{\kappa_i} b_i'(\kappa_i) - b_i'(\kappa)\, d\kappa$$

Since $b_i'(\kappa_i)$ is a decreasing function, $b_i'(\kappa_i) \geqslant b_i'(\kappa)$ for all $\kappa \leqslant \kappa_i$, hence the integrand is everywhere nonpositive; hence $\frac{d}{dx}\frac{b_i(\kappa_i)}{\kappa_i} \leqslant 0$ for all $\kappa_i > 0$. Thus, $\frac{b_i}{\kappa_i}$ is a decreasing function of κ_i fixing b_i. This gives us:

$$\frac{b_i}{\kappa_i} \geqslant \frac{b_1}{\kappa_1} \geqslant \left(1 - \frac{s_{max}}{1 - s_{max}}\right)$$

proving the lower bound (79), and thus proving claim 1.

A.5 Proof of proposition 4

Assuming that contract positions y_{ci} and demand shocks y_{di} are both normally distributed, with variances $\sigma_{ci}^2, \sigma_{di}^2$ respectively, the ratio of manipulation-induced variance to total variance, (20), is:

$$\frac{\sum_{i=1}^n s_i^2 \sigma_{ci}^2}{\sum_{i=1}^n \sigma_{di}^2 + \sum_{i=1}^n s_i^2 \sigma_{ci}^2} \tag{90}$$

For notational convenience, I will refer to the ratio in (90) as MVR, short for "manipulation variance ratio." To prove proposition 4, we need to show that:

$$MVR \leqslant \left(\frac{s_{max} E[C]}{E[V]}\right)^2 \tag{91}$$

First, note that contract volume is

$$E[C] = E\left[\sum_{i=1}^n |y_{ci}|\right] = \sum_{i=1}^n \sqrt{\frac{2\sigma_{ci}^2}{\pi}}$$

Appendix A.6 below proves that trade volume can be approximated as follows:

Claim 4. When s_{max} is small, we have:

$$\nu_i \equiv z_{Di}(p_b) \approx y_{di} + s_i y_{ci}$$

Claim 4, together with assumption 2, imply that expected total trade volume in the underlying asset is:

$$E[V] = E\left|\sum_i \nu_i\right| = \sum_{i=1}^{n} \sqrt{\frac{2\left(\sigma_{di}^2 + s_i^2\sigma_{ci}^2\right)}{\pi}}$$

Under assumption 2, we have $\sigma_{di}^2 = \theta^2\sigma_{ci}^2$. Thus,

$$\sum_{i=1}^{n} \sqrt{\sigma_{di}^2 + s_i^2\sigma_{ci}^2} = \sum_{i=1}^{n} \sqrt{\left(\theta^2 + s_i^2\right)\sigma_{ci}^2}$$

Taking the ratio of $E[C]$ and $E[V]$, we get:

$$\frac{E[C]}{E[V]} = \frac{\sum_{i=1}^{n}\sqrt{\sigma_{ci}^2}}{\sum_{i=1}^{n}\sqrt{\left(\theta^2 + s_i^2\right)\sigma_{ci}^2}}$$

Since this is decreasing in s_i, we have:

$$\frac{\sum_{i=1}^{n}\sqrt{\sigma_{ci}^2}}{\sum_{i=1}^{n}\sqrt{\left(\theta^2 + s_i^2\right)\sigma_{ci}^2}} \geqslant \frac{\sum_{i=1}^{n}\sqrt{\sigma_{ci}^2}}{\sum_{i=1}^{n}\sqrt{\left(\theta^2 + s_{max}^2\right)\sigma_{ci}^2}} = \frac{1}{\sqrt{\theta^2 + s_{max}^2}}$$

Hence,

$$\frac{E[C]}{E[V]} \geqslant \frac{1}{\sqrt{\theta^2 + s_{max}^2}}$$

implying that

$$\frac{s_{max}E[C]}{E[V]} \geqslant \frac{s_{max}}{\sqrt{\theta^2 + s_{max}^2}}$$

Or,

$$\left(\frac{s_{max}E[C]}{E[V]}\right)^2 \geqslant \frac{s_{max}^2}{s_{max}^2 + \theta^2} \tag{92}$$

Now, we will show that the RHS of inequality (92) is an upper bound for MVR. Taking

the derivative of the MVR as defined in (90) with respect to s_i, we have:

$$\frac{\partial MVR}{\partial s_i} = \frac{2\sigma_{ci}^2 s_i \left[\left(\sum_{i=1}^n \sigma_{di}^2 + \sum_{i=1}^n s_i^2\sigma_{ci}^2\right) - \left(\sum_{i=1}^n s_i^2\sigma_{ci}^2\right)\right]}{\left(\sum_{i=1}^n \sigma_{di}^2 + \sum_{i=1}^n s_i^2\sigma_{ci}^2\right)^2} > 0$$

Hence,

$$\frac{\sum_{i=1}^n s_i^2\sigma_{ci}^2}{\sum_{i=1}^n \sigma_{di}^2 + \sum_{i=1}^n s_i^2\sigma_{ci}^2} \leqslant \frac{\sum_{i=1}^n s_{max}^2\sigma_{ci}^2}{\sum_{i=1}^n \sigma_{di}^2 + \sum_{i=1}^n s_{max}^2\sigma_{ci}^2}$$

Now, plug in $\sigma_{di}^2 = \theta^2\sigma_{ci}^2$, to get:

$$\frac{\sum_{i=1}^n s_{max}^2\sigma_{ci}^2}{\sum_{i=1}^n \sigma_{di}^2 + \sum_{i=1}^n s_{max}^2\sigma_{ci}^2} = \frac{\sum_{i=1}^n s_{max}^2\sigma_{ci}^2}{\sum_{i=1}^n \theta^2\sigma_{ci}^2 + \sum_{i=1}^n s_{max}^2\sigma_{ci}^2} = \frac{s_{max}^2}{s_{max}^2 + \theta^2}$$

This implies that

$$MVR \leqslant \frac{s_{max}^2}{s_{max}^2 + \theta^2}$$

Combining this with 92, we have:

$$\left(\frac{s_{max}E[C]}{E[V]}\right)^2 \geqslant \frac{s_{max}^2}{s_{max}^2 + \theta^2} \geqslant MVR$$

proving (91), and thus proposition 4.

A.6 Proof of claim 4

To find trade volume, plug equilibrium prices (15) into demand (14), to get:

$$z_{Di}(p_b) = y_{di} + \frac{\kappa_i}{\sum_{j=1}^n \kappa_j} y_{ci} - \kappa_i\left(\frac{\sum_{i=1}^n y_{di}}{\sum_{i=1}^n \kappa_i} + \frac{\sum_{i=1}^n \kappa_i y_{ci}}{\left(\sum_{i=1}^n \kappa_i\right)^2}\right)$$

We can write this as:

$$z_{Di}(p_b) = \left(y_{di} - \frac{\kappa_i}{\sum_{i=1}^n \kappa_i}\sum_{i=1}^n y_{di}\right) + \frac{\kappa_i}{\sum_{j=1}^n \kappa_j}\left(y_{ci} - \frac{\sum_{i=1}^n \kappa_i y_{ci}}{\sum_{i=1}^n \kappa_i}\right)$$

If n is fairly large, then

$$\frac{\kappa_i}{\sum_{i=1}^n \kappa_i}\sum_{i=1}^n y_{di} \to 0$$

and,

$$\frac{\kappa_i \sum_{i=1}^{n} \kappa_i y_{ci}}{\left(\sum_{i=1}^{n} \kappa_i\right)^2} \to 0$$

Thus, we can ignore these terms, and we have approximately:

$$z_D(p_b) \approx y_{di} + \frac{\kappa_i}{\sum_{j=1}^{n} \kappa_j} y_{ci} = y_{di} + s_i y_{ci}$$

The term $s_i y_{ci}$ also approaches 0 as n becomes large, but since we are interested in settings in which y_{ci} may be much larger than y_{di}, I include this term in the volume approximation.

A.7 Symmetric equilibrium

The exact equilibrium of proposition 2 can be solved for analytically in the symmetric case. Assume there are n agents, with identical slopes of demand $\kappa_i = \kappa$. Specializing definition 2 to the symmetric case, we have:

Definition 5. The demand function $z_D(p, y_{di}, y_{ci})$ constitutes a *symmetric linear ex-post equilibrium* if:

1. As in (6), agents' bid curves are ex-post best responses given the slope of residual supply d:
$$z_{Di}(p; y_{di}, y_{ci}) = \frac{d}{\kappa + d} y_{di} + \frac{\kappa}{\kappa + d} y_{ci} - \frac{\kappa d}{\kappa + d}(p - \pi) \tag{93}$$
2. As in (3), the slope of residual supply is the sum of $n-1$ agents' bid curve slopes:
$$d \equiv z'_{RSi}(p) = -\sum_{j \neq i} z'_{Dj}(p; y_{dj}, y_{cj}) = (n-1)\frac{\kappa d}{\kappa + d} \tag{94}$$

(94) and (93) imply that

$$d = (n-1)\frac{\kappa d}{\kappa + d}$$

Solving for d, and plugging in to (93), we find the unique equilibrium bidding functions.

Proposition 7. *In the unique linear ex-post equilibrium, agents with demand shock* y_{di} *and contract position* y_{ci} *submit bid curves:*

$$z_{Di}(p; y_{di}, y_{ci}) = \frac{n-2}{n-1} y_{di} + \frac{1}{n-1} y_{ci} - \frac{n-2}{n-1}\kappa(p - \pi) \tag{95}$$

The slope of residual supply facing any agent is

$$d \equiv z'_{RS}(p) = (n-2)\kappa \tag{96}$$

Summing bid curves and equating to 0, we get the following expression for auction prices:

$$p_b - \pi = \frac{1}{n\kappa}\sum_{i=1}^{n} y_{di} + \frac{1}{n(n-2)\kappa}\sum_{i=1}^{n} y_{ci} \tag{97}$$

Assuming that $y_{di} \sim N\left(0, \sigma_d^2\right)$ and $y_{ci} \sim N\left(0, \sigma_c^2\right)$, we then have:

$$p_b - \pi \sim N\left(0, \frac{\sigma_d^2}{n\kappa^2} + \frac{\sigma_c^2}{n(n-2)^2\kappa^2}\right) \tag{98}$$

B Supplementary material for section 4

B.1 Proof of proposition 5

Agents in submarket m have slopes of demand $\kappa_{m1} \ldots \kappa_{mn_m}$, so this is enough to pin down the equilibrium bid curves by proposition 2; setting $y_{ci} = 0$ when $i > 1$ and $y_{di} = 0$ when $i = 1$, we have (24) and (25). To find prices, sum aggregate demand and equate to 0:

$$\frac{b_i}{\sum_{j\neq i} b_j}\phi_m y_c + \sum_{i=2}^{n_m} y_{dmi}\frac{b_{mi}}{\kappa_{mi}} - \sum_{i=1}^{n_m} b_{mi}(p_m - \pi) = 0$$

$$(p_m - \pi) = \frac{1}{\sum_{i=1}^{n_m} b_{mi}}\left[\frac{b_i}{\sum_{j\neq i} b_j}\phi_m y_c + \sum_{i=2}^{n_m} y_{dmi}\frac{b_{mi}}{\kappa_{mi}}\right]$$

as desired.

B.2 Benchmark price variance algebra

Using (26) of proposition 5, the submarket m price is:

$$(p_m - \pi) = \frac{1}{\sum_{i=1}^{n_m}\kappa_{mi}}\left[\bar{y}_{dm} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m}\kappa_{m1}}\phi_m y_c\right]$$

Rather than restricting to (23), I consider all benchmarks which are affine in submarket prices, which I write as:

$$P_b = C + \sum_{m=1}^{M} \phi_m (p_m - \pi)$$

Given a set of benchmark weights ϕ_m, the risk of the benchmark $P_b = \mu_\pi + \sum_{m=1}^{M} \phi_m (p_m - \pi)$ is:

$$E\left[(P_b - \pi)^2\right] = E\left[\left[C + \sum_{m=1}^{M} \phi_m \left(\frac{1}{\sum_{i=1}^{n_m} \kappa_m}\left[\pi + \bar{y}_{dm} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m} \kappa_{m1}} \phi_m y_c - \mu_\pi\right]\right) - \pi\right]^2\right]$$

We can write this as:

$$E\left[(P_b - \pi)^2\right] = E\left[\left[C - \mu_\pi + \mu_\pi + \sum_{m=1}^{M} \phi_m \left(\frac{1}{\sum_{i=1}^{n_m} \kappa_m}\left[\pi + \bar{y}_{dm} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m} \kappa_{m1}} \phi_m y_c - \mu_\pi\right]\right) - \pi\right]^2\right]$$

Rearranging the inside term, we have:

$$= E\left[\left[(C - \mu_\pi) + (\mu_\pi - \pi)\left(1 - \sum_{m=1}^{M} \phi_m\right) + \sum_{m=1}^{M} \phi_m \left(\frac{\bar{y}_{dm}}{\sum_{i=1}^{n_m} \kappa_m} + \frac{\kappa_{m1}}{\sum_{i=1}^{n_m} \kappa_{m1}} \phi_m y_c\right)\right]^2\right]$$

The terms $\pi - \mu_\pi$, $\bar{y}_{dm}$ and y_c are each independent mean-0 normals, with variances $\sigma_\pi^2, \sigma_{dm}^2, \sigma_c^2$ respectively. Hence this becomes:

$$(C - \mu_\pi)^2 + \left(1 - \sum_m \phi_m\right)^2 \sigma_\pi^2 + \sum_m \phi_m^2 \frac{\bar{\sigma}_d^2}{\left(\sum_{i=1}^{n_m} \kappa_m\right)^2} + \left[\sum_m \phi_m^2 \frac{\kappa_{m1}}{\left(\sum_{i=1}^{n_m} \kappa_m\right)^2}\right]^2 \sigma_c^2 \quad (99)$$

The term $(C - \mu_\pi)^2$ is constant, so it is always optimal to set $C = \mu_\pi$, so all optimal benchmarks take the form (23). Under this restriction, (99) simplifies to (28).

C Supplementary material for section 5

C.1 CFTC Commitments of Traders Data

The COT reports contains four concentration metrics: the percentage of net or gross open interest held by the top 4 and top 8 largest agents. Net concentration differs from gross concentration in that it nets out spreading positions which involve offsetting contract positions at different delivery months; for my purposes, what matters is the fraction of open interest held in any particular delivery month, so gross concentration is a more appropriate metric. Likewise, although top 4 and top 8 concentration are strongly correlated, I use top 4 concentration, as markets with $n = 8$ similarly sized participants are quite competitive in the context of my theory. I filter to data from 2010-01-01 to 2016-12-27.

I divide commodity and financial futures into subgroups based on the CFTC's commodity subgroup codes. The CFTC does not provide a codebook, so I create my own labels, in the "Description" columns of tables 1 and 5. A few commodity categories – in particular, those for oil, gas and electricity – have very large numbers of contracts. There are 92 kinds of oil contracts, 106 kinds of natural gas contracts, and 287 kinds of electricity contracts in my data. However, open interest is very concentrated, and the largest contract markets are much bigger than median contract markets. I must therefore choose a scheme for aggregating over all these contracts. Flat averaging does not reflect differences in size between contract markets, and weighting contract markets by open interest is undesirable because contract units can differ within commodity categories, and also because estimates would generally be dominated by the largest contract markets. Instead, I use only the 10 largest contracts in each commodity category, based on the total open interest summed across all dates where a contract is observed. The numbers in table 1 then reflect flat averages across the 10 largest contract categories. Thus, my estimates reflect only large contract markets in each category.

The COT report data reports open interest for five categories of traders. For commodity futures, these are "Producers/Merchants/Processors/Users" ("Prod+Merch" in table 1), "Swap Dealers", "Managed Money" ("Man. Money"), "Other Reportable", and "Nonreportable". The distinction between "Other Reportable" and "Nonreportable" is that "Other Reportable" represents parties who hold large enough positions to meet reporting requirements, but do not fall into other categories; nonreportable positions are calculated

as a difference between total open interest and total held by reporting parties. For the purposes of my analysis, I combine "Other Reportable" and "Nonreportable" into a single "Other" category. The COT reports also separately reports long, short, and spreading open interest, where spreading open interest represents offsetting long and short contracts with different delivery months. I do not need to separately consider spreading positions for the purposes of my analysis; thus, I add the fraction of spreading positions to both long and short positions. This makes long and short positions summed across participants equal to 100%, up to rounding error, as expected.

"Producers/Merchants/Processors/Users" are participants in markets for the underlying asset – for example, agents who produce, process or store the physical underlying commodity, and are thus exposed to price risk that they hedge in futures markets. Since they directly trade the underlying physical commodity, they are likely to have relatively low costs of taking positions in the asset, so high κ_i values. "Managed money" participants are commodity trading advisors, commodity pool operators, or other funds, who can be thought of as largely financial speculators purchasing risk; they are thus likely to have low κ_i values.[102] The role of "swap dealers" in these markets in less clear. The CFTC's explanatory notes state: "A "swap dealer" is an entity that deals primarily in swaps for a commodity and uses the futures markets to manage or hedge the risk associated with those swaps transactions." A more detailed 2008 CFTC staff report on swap dealers (CFTC, 2008) suggests that swap dealers primarily function to bridge between the OTC and exchange-traded derivatives markets; that is, they enter into customized OTC derivative contracts with customers, and hedge the resultant risk on commodity exchanges. The staff report suggests that swap dealers' customers are both speculators and hedgers. For the purposes of this paper, the swap dealers themselves are largely financial participants, who are likely to have high cost of taking physical positions in the underlying asset, and thus low κ_i values.

C.2 Financial futures

Table 5 reports results, as in table 1, for financial futures. The categories of participants are "Dealer Intermediary", "Asset Manager/Institutional", "Leveraged Funds", "Other Reportable", and "Nonreportable". "Dealers" are sell-side participants – intermediaries that earn commisions from providing liquidity to customers. "Asset managers" are

[102]Further description of the different classes of participants is available here.

largely buy-side institutions such as pension funds, endowments, insurance companies, and mutual funds. "Leveraged funds" are parties such as hedge funds, commodity trading advisors, commodity pool operators and other buy-side parties who are likely to have higher risk tolerance than asset managers.[103] The relative κ_i values of these three classes of participants are less clear, but one might guess that large dealers and hedge funds might find it easier to tie up capital in large speculative positions for underlying assets than the more risk-averse asset managers.

Similarly to table 1, there are large differences in long and short concentration for a number of products – longs are more concentrated than shorts for interest rate products, representing cash-settled contracts such as Eurodollar futures as well as deliverable IR swaps.[104] Gross short and long concentration is relative close for all products other than interest rates. Across participant types, dealers tend to take short positions in equity index and treasury product, and long positions in interest rate products. Asset managers largely take opposing positions, being long equity products and treasuries and short interest rate products. Since dealers are relatively large and may be more willing to sink capital into risky assets than asset managers, contract manipulation may tend to cause equity index prices to be manipulated downwards and interest rates to be manipulated upwards.

[103]Further description of the different classes of participants is available here.

[104]A long position in interest rate products is long fixed, short floating rates, for both deliverable and cash-settled products.

Table 5: Commitments of Traders reports: financial futures

	Top 4		s_{max}	Dealers		Asset Man.		Lev. Money		Other	
Category	Long	Short		Long	Short	Long	Short	Long	Short	Long	Short
FX	42%	37%	20%	32%	25%	14%	15%	29%	31%	25%	28%
Emerging Markets FX	63%	62%	31%	39%	57%	9%	4%	40%	26%	12%	13%
Equity Indices	31%	35%	17%	21%	40%	38%	22%	24%	20%	18%	19%
Treasuries	29%	26%	14%	8%	17%	48%	34%	22%	25%	22%	24%
Interest Rates	61%	45%	27%	55%	39%	8%	15%	27%	30%	10%	16%
Other Indices	52%	62%	28%	24%	38%	25%	6%	32%	45%	19%	11%

Notes. "Top 4" is the fraction of gross long and short open interest held by the 4 largest market participants, defined as q_{l4} and q_{s4} above. s_{max} is the sum of q_{l4} and q_{s4} divided by 4, as in (36). All other columns describe the fraction of long and short open interest held by different classes of market participants. Open interest percentages add to 100%, separately for longs and shorts, across the 4 categories of market participants: "Dealers", "Asset Man.", "Lev. Money", and "Other". Since the CFTC does not provide a codebook for subgroup codes, I manually label subgroups in the "Description" column.

C.3 Discussion of assumption 5

Roughly speaking, assumption 5 can be justified as follows. If we assume that the absolute value of contract position size y_{ci} follows a Zipf's law distribution among agents, with exponent greater than 1, then the top 8 agents have a nonvanishing share of total contract position size as the total number of agents grows large. Observed concentration ratios imply that the share of the 8 largest agents is roughly 20-70% of the total. Using numerical simulations, this is matched by a Zipf's law exponent between 1.02 and 1.51. These exponents imply that the ratio between the top 1 and top 8 shares is between 37.5% and 52.4%. To err on the side of overestimating concentration, assumption 5 sets s_{max} to be equal to 50% of the top 8 share in (36).

An issue with this justification for assumption 5 is that the top 4 long and short agents are not always the 8 largest agents in absolute size; however, this is unlikely to have a large quantitative impact on results.

C.4 Qualitative evidence from recent contract manipulation cases

A core prediction of my theory is that agents with larger κ_i values have larger incentives to manipulate. A sharper statement of this is that agents most likely to manipulate are not necessarily those who are financially largest, or those who hold the largest contract positions, but those agents that have the lowest costs for trading large amounts of the underlying asset. This is qualitatively consistent with a number of manipulation cases in the US: many of the largest cases have involved agents who were active traders in markets for underlying assets.

In 2011, the CFTC brought manipulation charges against Parnon Energy Inc, Arcadia Petroleum Ltd. and Arcadia Energy in 2011, claiming that these parties traded futures and other derivatives based on crude oil prices, then traded large quantities of physical oil to move prices to increase contract payoffs,[105] eventually leading to fines totalling $13 million.[106] Parnon Energy Inc., Arcadia Petroleum Ltd. and Arcadia Energy are all usually active in physical oil markets. In a similar case, the CFTC fined Statoil ASA $4 million for purchasing swap contracts tied to the Argus FEI, an index of propane prices, and then purchasing large quantities of propane in an attempt to increase the Argus FEI.

[105] CFTC Press Release 6041-11
[106] CFTC Press Release 6971-14

Statoil (now Equinor) is an energy company active in oil and gas production and pipeline operation.

In natural gas markets, the CFTC[107] in 2007 and the FERC[108] in 2009 fined Energy Transfer Partners, L.P. (ETP) a total of $40 million for manipulation of natural gas prices. According to the FERC[109], ETP entered into a large number of basis swap contracts, whose value is based on the difference between average natural gas prices at Henry Hub, a liquid trading location, and Houston Ship Channel (HSC), a very illiquid location. ETP dominated gas sales at HSC, often comprising 80% or more of total sales. ETP aggressively sold physical gas for delivery at HSC in order to depress prices, increasing payoffs on basis swap contracts. At the time, ETP owned many large natural gas pipeline systems, and was an active participants in spot markets for natural gas. In a similar case, the CFTC fined Total Gas and Power (TGP) $3.6 million in 2015 for entering into basis swap contracts and then trading large amounts of physical natural gas to move average prices and thus contract payoffs; in some cases, TGP was responsible for over 40% of trades used to calculate index prices.[110]

Other examples of manipulation by agents generally active in markets for the underlying asset are a 2015 CFTC case against Kraft Foods Group, Inc. and Mondelez Global LLC, both active buyers of physical wheat, alleging manipulation of wheat futures prices;[111] and a 1998 CFTC case against Sumitomo Corporation, a company producing copper cathodes, for cornering the copper market by acquiring dominant positions in copper futures and physical copper.[112] In comparison, relatively few high-profile cases of commodity derivatives contract manipulation involve purely financial participants who do not usually participate in markets for the underlying asset.

The distinction between the underlying and contract markets appears to be smaller in the case of purely financial assets; however, many cases of financial benchmark manipulation also seem to involve active participants in underlying markets. The recent FX manipulation scandal involved banks who were actively traded as FX dealers, frequently taking on large positions in both the contract and underlying markets Levine (2014). Similarly, the ISDAFIX manipulation scandal involved dealers who were active in traders of

[107] CFTC Press Release 5471-08
[108] 128 FERC 61,269
[109] 120 FERC 61,086
[110] CFTC Docket No. 16-03
[111] CFTC Press Release 7150-15
[112] In the Matter of Sumitomo Corporation

ISDAFIX derivatives as well as the underlying swaps and swap spreads[113]. It is plausible that different classes of agents have different holding costs for certain kinds of financial assets; for example, dealers are probably in a better position to trade large amounts of repo debt than student loan organizations or mortgage banks, hence it is likely difficult for the latter class of agents to manipulate repo-based interest rate benchmarks even if they hold large interest rate derivative positions.

D Supplementary material for section 6

D.1 Determinants of auction participation

Figure 8 shows a kernel regression of the max number of participants versus date, using the full set of 1,650 auctions. There are significant trends in auction participation over time. In addition, auction participation is correlated with the size of price differences: participation is higher when gold price innovations are larger. This is true when price differences are measured using auction prices, and also using daily gold futures prices from FactSet, suggesting that this is not due to reverse causation – that is, higher auction participation causing gold auction prices to be more variable.

This could confound my results, if trends in auction participation and my measured $\hat{B}_a$ are primarily caused by time trends, or other factors such as gold price differences. To rule out this possibility, the first three volumns of table 6 regress the estimated aggregate demand slope $\hat{B}_a$ on the maximum number of participants, adding controls for price differences and date. To increase sample size, I use the full sample of 1,650 auctions, although auctions that end in a single round cannot be used since I cannot measure the slope of auction demand. I measure participation using the max number of participants across rounds in the auction. Column 1 shows that the coefficient on n is positive, but the coefficient on n^2 is negative; this means that auction aggregate demand slope is an increasing but concave function of the maximum number of participants in the auction. As illustrated in figures 1 and 2 in the main text, this is the core stylized fact that allow me to infer that marginal participants with higher n have lower values of κ_i. This stylized fact survives controls for price difference and absolute price difference, as well as a fourth-order polynomial in auction date.

[113] CFTC Press Release 7180-15

Figure 8: Trends in gold auction participation over time

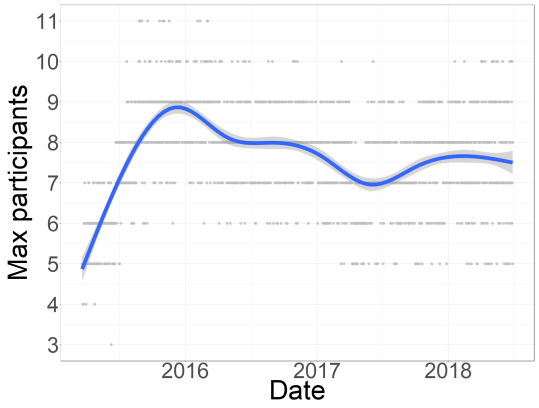

Notes. Kernel regression of the max number of auction participants, $\max_r n_{ar}$, on the auction date. Each data point is a single auction a.

Similarly, columns 4-6 show regressions of auction volume v_a on the max number of participants, price differences, and time. The core stylized fact is that volume is increasing in the number of auction participants. Again, this is robust to controlling for auction price differences, and auction date.

D.2 Dynamic auctions

I model the ICE auction as a continuous-time dynamic auction game. As in the baseline model of section 3, $i \in (1 \dots n)$ agents have types $\kappa_1 \dots \kappa_n$, which are common knowledge, and demand shocks $y_{d1} \dots y_{dn}$ and contract positions $y_{c1} \dots y_{cn}$, which are private information. As before, the utility of agent i for purchasing z units of the asset when the price is p is:

$$\pi z + \frac{y_{di} z}{\kappa} - \frac{z^2}{2\kappa} - pz + py_{ci}$$

Instead of submitting bid curves, agents play a continuous-time auction, which I model as a simple differential game. At time $t = 0$, the auction begins at some deterministic price $p(0)$, which is known to all participants.[114] Participants simultaneously announce

[114]IBA's published auctions specification documents do not describe how the auction starting price is chosen. However, the choice of starting price does not significantly affect agents' optimal strategies, so I assume the price is non-random and commonly known for simplicity.

Table 6: Determinants of gold auction demand slope and volume

	$\hat{B}_a$, 1000oz / (1 USD/oz)			Volume, 1000oz		
	(1)	(2)	(3)	(4)	(5)	(6)
Max n	41.715***	41.073***	31.313***	22.725***	22.145***	31.331***
	(7.979)	(7.965)	(8.599)	(1.479)	(1.479)	(1.707)
Max n squared	−2.665***	−2.609***	−1.905***			
	(0.515)	(0.514)	(0.543)			
Price diff		0.022	−0.015		0.283	0.277
		(0.138)	(0.132)		(0.265)	(0.255)
Abs price diff		−0.546***	−0.177		1.444***	1.319***
		(0.199)	(0.193)		(0.382)	(0.372)
Date polynomial			X			X
N	1,330	1,330	1,330	1,649	1,649	1,649
R^2	0.020	0.026	0.112	0.125	0.134	0.198

Notes. "Max n" is the max number of auction participants across rounds, $\max_r n_{ar}$. "Price diff" is first differences of auction prices, and "Abs price diff" is the absolute value of "price diff". "Date polynomial" is a fourth-order polynomial in calendar date. In the first three columns, the dependent variable is the estimated slope of auction demand $\hat{B}_a$, and in the last three the dependent variable is auction volume v_a.

initial demands $z_D(p(0))$. Thereafter, at any given time t, the price evolves according to the differential equation:

$$\frac{dp}{dt} = \begin{cases} k & z_D(p(t)) > 0 \\ -k & z_D(p(t)) < 0 \end{cases}$$

Agents can update their demand functions $z_{Di}(p(t))$ as the price $p(t)$ changes; agents' demand functions are required to be decreasing in price. The game ends at the first time T when aggregate demand is exactly equal to 0,

$$\sum_{i=1}^{n} z_{Di}(p(T)) = 0$$

at which point each agent purchases $z_{Di}(p(T))$ units of the good for $p(T)$ per unit, and is paid $p(T)$ per unit contract that she holds.

I assume that agents choose the rate at which demand changes with price, $z'_{Di}(p(t)) = \frac{dz_{Di}}{dp}$, rather than the level of demand; this ensures that resultant demand functions are continuous. Since $\frac{dp}{dt}$ is constant throughout any instance of the game, choosing $\frac{dz_{Di}}{dp}$ is equivalent to choosing $\frac{dz_{Di}}{dt}$.

I require agents' reported demand slopes to be finite and bounded away from 0:

$$-M \leqslant z'_{Di}(p(t)) \leqslant -\epsilon < 0 \tag{100}$$

This guarantees that the game will end in finite time. I will show that the dynamic differential auction game admits an equilibrium which coincides with the equilibrium of the static auction in section 3; since agents' bid curves in the static game have slopes which are negative, finite and bounded away from 0, for any $\kappa_1 \ldots \kappa_n$, ϵ can always be chosen small enough in magnitude, and M large enough, that the bounds in (100) are not binding in equilibrium.

Claim 5. The dynamic auction game always ends in finite time.

Proof. Suppose that $\sum_{i=1}^{n} z_{Di}(p(0)) > 0$. Then, we have:

$$\frac{dz_D}{dt} = \sum_{i=1}^{n} \frac{dz_{Di}}{dt} = \sum_{i=1}^{n} \frac{dz_{Di}}{dp}\frac{dp}{dt} = k\sum_{i=1}^{n} \frac{dz_{Di}}{dp} \in [-kM, -k\epsilon]$$

Thus,

$$z_D(p(t)) \in [z_D(p(0)) - kMt,\ z_D(p(0)) - k\epsilon t]$$

Hence, $z_D(p(t)) = 0$ for some $t \in \left[\frac{z_D(0)}{kM},\ \frac{z_D(0)}{k\epsilon}\right]$, hence the game ends in finite time. The proof of the case where $\sum_{i=1}^{n} z_{Di}(p(0)) < 0$ is analogous. □

The setup of this game ignores three features of the ICE auction. First, the price does not evolve smoothly, but jumps in increments, sometimes overshooting the equilibrium price. Second, agents are not required to submit monotone bid functions. Third, the auction does not end when supply and demand are exactly equal, but allows for some volume gap. All of these are difficult to model, since the game becomes discontinuous, and the model would need to take a stance on the bid updating rule, which to my knowledge is not publically documented by ICE. Thus, I adopt the simpler monotone differential game model as an approximation to the ICE auction.

A *history* is a sequence of observed demand slopes $z'_{Di}(p(t))$ of all agents on the interval $t \in [0, T]$. It is equivalent to assume we observe agents' demand functions; thus a history at time t can be described as:

$$h_T = \{(z_{D1}(p(t)) \dots z_{Dn}(p(t))),\ t \in [0, T]\}$$

A *strategy* is a decision $z_{Di}(p(0))$ about what demand to announce at the starting price $p(0)$, and then a choice of $z'_{Di}(p(t))$ for every possible history. Both decisions may also depend on the realization of agents' demand shock y_{di} and contract position y_{ci}. Thus, in full generality, strategies can be very complex. I will restrict attention to a class of *naive linear strategies*, which I define below.

Definition 6. A *naive linear strategy* for agent i is a strategy in which:

$$z_{Di}(p(0)) = f_i(y_{di}, y_{ci}),\ z'_{Di}(p(t)) = b_i$$

that is, agent i's demand function has an intercept which may depend on y_{di} and y_{ci}, and constant slope.

Agents playing naive strategies cannot condition their slopes on the observed behavior of other agents, or their own contract positions and demand shocks – they must commit to a constant slope, independent of the history of the game. Note that (100) implies that

we must have

$$b_i \in [-M, -\epsilon]$$

Naive linear strategies are isomorphic to affine bid curves. For any naive linear strategies described by $f_i(y_{di}, y_{ci}), b_i$, we can construct the equivalent bid curve as:

$$z_{Di}(p; f_i(y_{di}, y_{ci}), b_i) \equiv f_i(y_{di}, y_{ci}) + \int_{p_0}^{p} b_i dp = f_i(y_{di}, y_{ci}) + b_i(p - p_0) \tag{101}$$

Proposition 8. *Naive linear strategies* $f_i(y_{di}, y_{ci}), b_i$ *corresponding to equilibrium bidding strategies* $z_{Di}(p; y_{di}, y_{ci})$ *in the static auction game in proposition 2 also constitute an equilibrium in the dynamic auction game.*

I prove proposition 8 in two steps. Claim 6 shows that, if all agents are restricted to naive linear strategies, the dynamic auction game produces exactly the same outcomes as the static auction game. Claim 7 shows that, if all agents other than i are playing naive linear strategies, agent i can do no better than to play a naive linear strategy. Thus, equilibria in naive linear strategies are equilibria in the broader game. This proves proposition 8, as equilibrium strategies in proposition 2 for the static auction game can be implemented as naive linear strategies in the dynamic auction game.

Claim 6. If all agents play naive linear strategies, the outcomes of the dynamic auction game are exactly the outcomes of a static auction game in which in which agents submit the bid functions specified by (101).

Proof. Suppose agents are playing the naive linear strategies $f_i(y_{di}, y_{ci}), b_i$. Construct the equivalent strictly decreasing bid curves $z_{Di}(p; f_i(y_{di}, y_{ci}), b_i)$. Fix some realization of y_{di}, y_{ci} across agents, and consider the aggregate demand function,

$$\sum_{i=1}^{n} z_{Di}(p; f_i(y_{di}, y_{ci}), b_i)$$

Claim 5 states that the dynamic auction game always ends in finite time; thus, the dynamic auction must end at some time T, at some price $p(T)$, at which $\sum_i z_{Di}(p(T); f_i(y_{di}, y_{ci}), b_i) = 0$. Then each agent purchases $z_{Di}(p; f_i(y_{di}, y_{ci}), g_i(y_{di}, y_{ci}))$ units of the good at price $p(T)$, and is paid $p(T)$ per unit contract she holds. This is exactly the same outcome as agents receive in a static auction game in which agents submit bid curves $z_{Di}(p; f_i(y_{di}, y_{ci}), b_i)$. □

The strategies in the dynamic auction game with naive linear strategies are a subset of their strategies in the static auction, since they are not allowed to condition the slopes of their demand on y_{di}, y_{ci}. However, all equilibria in the static auction game correspond to naive linear strategies, because agents' bid slopes in equilibrium do not depend on y_{di}, y_{ci}, as shown by proposition 2 in the baseline model of section 3. Thus, if agents are restricted to playing naive linear strategies, equilibria in the dynamic auction game correspond to those of the static auction game. The following Claim shows that these strategies also constitute equilibria without the restriction to naive linear strategies:

Claim 7. If all agents other than i are playing naive linear strategies, it is weakly optimal for agent i to play a naive linear strategy.

Proof. Suppose all other agents are playing naive linear strategies, described by $f_j\left(y_{dj}, y_{cj}\right), b_j$. From the perspective of agent i, there is a random residual supply curve:

$$z_{RSi}\left(p, y_{d,-i}, y_{c,-i}\right) \equiv -\sum_{j \neq i} z_{Dj}\left(p, y_{dj}, y_{cj}\right) = -\sum_{j \neq i}\left[f_j\left(y_{dj}, y_{cj}\right) + b_j\left(p - \pi\right)\right]$$

Importantly, the slope of the residual supply curve is constant at $d = \sum_{j \neq i} b_j$. The dynamic auction game concludes when $\sum_{i=1}^{n} z_{Di}\left(p\right) = 0$, that is, when $z_{Di}\left(p\right) = z_{RSi}\left(p, y_{d,-i}, y_{c,-i}\right)$. Hence, the set of all attainable combinations of quantity and price available to the agent, for any realization of $\left(y_{d,-i}, y_{c,-i}\right)$, are described by $z_{RSi}\left(p, y_{d,-i}, y_{c,-i}\right)$. The agent can do no better, using arbitrarily complex strategies, than choosing the optimal point on $z_{RSi}\left(p, y_{d,-i}, y_{c,-i}\right)$ point for every realization of $\left(y_{d,-i}, y_{c,-i}\right)$. Proposition 1 shows that this can be accomplished in the static auction game by submitting the bid curve:

$$z_{Di}\left(p;\, y_{di}, y_{ci}\right) = \frac{d}{\kappa + d} y_{di} + \frac{\kappa}{\kappa + d} y_{ci} - \frac{\kappa d}{\kappa + d}\left(p - \pi\right)$$

Since the slope of z_{Di} does not depend on y_{di} or y_{ci}, agent i can implement this bid curve by playing the naive linear strategy:

$$f_i\left(y_{di}, y_{ci}\right) = \frac{d}{\kappa + d} y_{di} + \frac{\kappa}{\kappa + d} y_{ci} + \frac{\kappa d}{\kappa + d}\pi,\ b_i = -\frac{\kappa d}{\kappa + d}$$

Hence, agent i can do no better than playing a naive linear strategy. □

Claim 7 shows that there exists an equilibrium in the game in which agents play naive linear strategies corresponding to equilibrium strategies in the static auction game, as in

proposition 2. However, this does not imply that naive strategies are the only equilibrium strategies in this game; the strategy space is rich, and it is possible that many collusive equilibria of the kind documented by Wilson (1979) and others may exist. In particular, appendix D.3 provides evidence that agents shade bids more in earlier auction rounds, which is inconsistent with assumption 6; however, appendix D.3 also suggests that bid shading seems not to have large effects on my estimates of the slope of auction demand.

D.3 Round volume

Figure 9 regresses round buy and sell volume separately on round number, using the 509 auctions which constitute my main estimation subsample. Volume is increasing with round number, suggesting that agents are shading bids more in earlier rounds. This is not because of increased participation, as the estimation subsample uses only auctions in which participation is constant throughout, and this trend survives controlling for the number of participants. This observation is inconsistent with assumption 6, which says that agents bid in each round exactly their equilibrium bid curves in a uniform-price double auction. The model is thus not entirely appropriate. If we assume that the final auction outcome is still close to the bid-function submission game, bid shading in earlier rounds likely implies that my estimated slopes of demand are biased downwards; thus, I am likely to overestimate price variance due to manipulation.

As a robustness check, I estimate the slope of auction demand using only auctions which lasted at least 4 rounds, and using only data from rounds 3 and above. If bid shading substantially affects auction aggregate demand slope, demand slopes estimated using only late-round bidding data should be higher than demand slopes estimated using all rounds. The results are shown in figure 10; demand slopes estimated using only late rounds are statistically almost indistinguishable from demand slopes estimated using the full dataset. Thus, while I do not explicitly account for bid shading, these results suggest that bid shading likely does not influence my estimation of auction demand slopes by a quantitatively large amount. Dynamic uniform-price auctions are difficult to model and solve, so I do not address this issue further in this paper.

Figure 9: Buy and sell volume vs gold auction round number

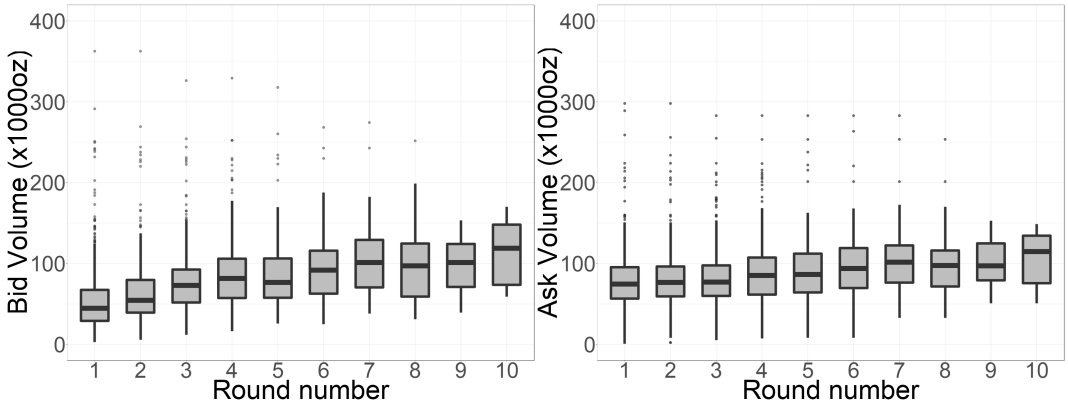

Notes. Round-level average buy volume b_{ar} (left) and sell volume s_{ar} (right) versus round number r.

Figure 10: Gold auction demand slope, all rounds vs late rounds only

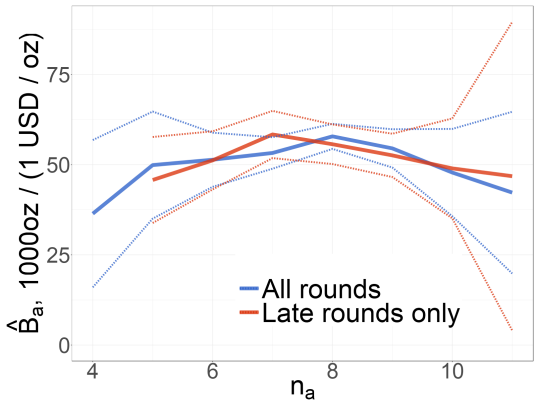

Notes. Slope of auction demand by auction participant number n_a, estimated using data from all auction rounds (blue) and only data from auction rounds 3 and above (red).

D.4 Moment matching: Auction demand slope

As described in (42) the main text, I measure $\hat{B}_a$ for each auction by regressing volume imbalance on prices:

$$\hat{B}_a \equiv -\frac{\sum_{t=1}^{T_a} \left(i_{at} - \bar{i}_a\right)\left(p_{at} - \bar{p}_a\right)}{\sum_{t=1}^{T_a} \left(p_{at} - \bar{p}_a\right)}$$

Given (41), $\hat{B}_a$ thus measured should be equal to $\sum_{i=1}^{n} b_{ai}$.

For any set of parameters $\kappa_1 \ldots \kappa_N$, for auction a with n_a participants, with demand slopes $\kappa_1 \ldots \kappa_{n_a}$, proposition 2 states that there is a unique set of equilibrium bid slopes

$$b_{a1}\left(\kappa_1 \ldots \kappa_{n_a}\right) \ldots b_{an_a}\left(\kappa_1 \ldots \kappa_{n_a}\right)$$

and thus a unique predicted slope of aggregate demand $B_a\left(\kappa_1 \ldots \kappa_{n_a}\right)$. This can be calculated by solving (12) numerically for B and then plugging into (11).

I assume that $\kappa_1 \ldots \kappa_4$ are equal, and I assume that $\log\left(\kappa_i\right)$, for $i \in \{4 \ldots 11\}$ is a second-order polynomial in i, with coefficients described by the parameter vector θ_κ. For a given θ_κ, let

$$B\left(n_a;\ \theta_\kappa\right)$$

represent the predicted slope of aggregate demand given θ_κ. I choose parameter vector θ_κ^* to minimize the distance between predicted and actual slopes of demand:

$$\theta_\kappa^* = \arg\min_{\theta_\kappa} \sum_{a=1}^{A} \left[\hat{B}_a - B\left(n_a;\ \theta_\kappa\right)\right]^2 \tag{102}$$

D.5 Moment matching: Auction volume

We observe the total volume, v_a, in each auction. In the asymmetric model, in equilibrium, all agents trade 0 on average, since their demand shocks have mean 0. In appendix D.6, I show that, given the equilibrium vector of bid slopes b_i, total trade volume for agent i, $z_{Di}\left(p_a\right)$, is mean 0, with variance:

$$Var\left(z_{Di}\left(p_a\right)\right) = \left(\frac{b_{ai}}{\kappa_i}\left(1 + \frac{b_{ai}}{\sum_{j=1}^{n} b_{aj}}\right)\right)^2 \sigma_{di}^2 + b_{ai}^2 \sigma_P^2 - b_{ai}^2 \left(\frac{\frac{b_{ai}}{\kappa_i}}{\sum_{j=1}^{n} b_{aj}}\right)^2 \sigma_{di}^2 \tag{103}$$

The expectation of the absolute value of a mean-0 normal distribution is $\sqrt{\frac{2\sigma^2}{\pi}}$; thus, expected absolute volume for participant i in equilibrium with n agents is:

$$E_n\left[|z_{Di}(p_a)|\right] = \sqrt{\frac{2Var\left(z_{Di}(p_a)\right)}{\pi}}$$

Expected volume for all participants the sum of $E\left[|z_{Di}(p_b)|\right]$ across all participants i. This is a function of n as well as both demand slopes κ_i and the variance of demand shocks σ^2_{di}. As with demand slopes, I assume that $\log\left(\sigma^2_{di}\right)$ for $i \geqslant 4$ is a second-order polynomial in i, with coefficients θ_d. Predicted volume in an auction with n_a participants depends on both θ_κ, through dependence on b_{ai} and κ_i, and θ_d, through σ^2_{di}. I denote predicted volume given parameters as:

$$v\left(n_a;\ \theta_\kappa, \theta_d\right) \equiv \sum_{i=1}^{n} E_n\left[|z_{Di}(p_a)| \mid \theta_\kappa, \theta_d\right]$$

Fixing $\theta_\kappa = \theta^*_\kappa$ at the optimal value from (102), I choose θ^*_d to minimize the distance between predicted and observed volume:

$$\theta^*_d = \arg\min_{\theta_d} \sum_{a=1}^{A} \left[v_a - v\left(n_a;\ \theta^*_\kappa, \theta_d\right)\right]^2$$

D.6 Derivation of expected volume

Using (39), and ignoring contract positions, we have:

$$p_b - \pi = \frac{\sum_{i=1}^{n} \frac{y_{di} b_i}{\kappa_i}}{\sum_{i=1}^{n} b_i}$$

Demand for each agent at the auction clearing price is:

$$z_{Di}(p_b) = y_{di}\frac{b_i}{\kappa_i} - b_i\left(p_b - \pi\right)$$

$$\implies z_{Di}(p_b) = y_{di}\frac{b_i}{\kappa_i} - b_i\frac{\sum_{i=1}^{n} \frac{y_{di} b_i}{\kappa_i}}{\sum_{i=1}^{n} b_i}$$

Variance can be written as:

$$Var\left(z_{Di}\left(p_b\right)\right) = \left(\frac{b_i}{\kappa_i}\left(1+\frac{b_i}{\sum_{i=1}^{n} b_i}\right)\right)^2 \sigma_{di}^2 + b_i^2\sigma_p^2 - b_i^2\left(\frac{\frac{b_i}{\kappa_i}}{\sum_{i=1}^{n} b_i}\right)^2 \sigma_{di}^2$$

So expected volume from agent i is:

$$\sqrt{\frac{2Var\left(z_{Di}\left(p_b\right)\right)}{\pi}}$$

Expected total volume is the sum of expected volume across agents.

D.7 Aggregated benchmark manipulability

Instead of settling gold futures contracts based on a single auction, we could construct an aggregated benchmark based on auctions over a longer period of time;[115] this would decrease the weight of any given auction on the benchmark, and make manipulation more costly. In this subsection, I calculate the manipulability of a benchmark calculated by averaging auction prices over a month. I construct a synthetic month consisting of 22 business days and 44 auctions, with a distribution of number of participants n_a matching the distribution in my sample.[116] I assume that the benchmark administrator can set auction weights $\phi\left(n_a\right)$, which sum to 1, which can vary according to n_a; thus the administrator can assign higher weight to auctions with more participants, which are harder to manipulate. Thus, the aggregated gold price benchmark is:

$$P_b = \sum_{a=1}^{44} \phi\left(n_a\right) p_a, \ \sum_{a=1}^{44} \phi\left(n_a\right) = 1$$

I maintain parameter estimates for κ_i and σ_{di}^2 from previous subsections. I assume that there is a fixed set of agents, $i = 1 \ldots N$, with contract positions y_{ci} which are constant

[115] A drawback is that this benchmark would represent averaged gold prices over a month, rather than gold prices at a single point in time. However, COMEX gold futures contracts are currently settled by physical delivery, and shorts can choose any day within the contract month to make delivery; thus, gold futures contracts are currently effectively only linked to gold prices at monthly resolution. A number of other benchmarks are based on prices averaged over relatively long periods of time; the Platts *Inside FERC* benchmarks are based on average prices of trades during "bid week" once a month (see the S&P Global Platts website), and the ICE Brent Index is calculated based on average oil prices over a month.

[116] Since my sample includes only auctions with a constant number of participants, it does not include any full months of data, so I cannot simply use a fixed month in the data.

throughout all auctions $a \in \{1 \dots 44\}$, and demand shocks y_{dai} which are independently realized for each auction a. I assume that agents' utilities are additive over all auctions that they participate in. Thus, the utility function of agent i for a series of prices p_a, quantities z_{ia} is:

$$\left[\sum_{a:i \leqslant n_a} \pi z_{ai} - \frac{z_{ai}^2}{2\kappa_i} - p_a z_{ai} \right] + P_b y_{ci}$$

The assumption that auction participants' utilities are additively separable across auctions is likely to be violated in practice. Auction participants are likely to treat gold across different auctions as substitutes, so if a manipulator sells a lot of gold in one auction, she will decrease also decrease participants' marginal utility for gold in later auctions; they will lower bids in future auctions, lowering prices and thus the benchmark. Hence, benchmark manipulation is probably somewhat easier than my model suggests. The aggregated benchmark functions as least as well as a single auction; we can think of the results of this section as an optimistic bound, estimating the manipulability of the LBMA gold price benchmark in the best possible case.

In appendix D.8, I calculate the variance of the aggregated benchmark for any given vector of weights $\phi(n_a)$. I choose auction weights $\phi(n_a)$ to minimize total benchmark variance from demand shocks and manipulation. The aggregated benchmark functions substantially better than any individual auction. Maintaining the size of contract positions, $\sqrt{\sigma_{ci}^2}$, at 150,000oz of gold, the square root of manipulation-induced variance is \$0.0328 USD/oz; this is a negligible fraction of monthly innovations to average gold auction prices, which have standard deviation \$38.14 USD/oz.[117] Put another way, under the aggregated auction, the size of agents' contract positions can be greatly increased before manipulation variance is non-negligible. If we set the standard deviation of contract positions to 100,000,000oz, roughly 300 times the size of current contract position limits, the square root of manipulation-induced variance is still only \$21.64 USD/oz, so manipulation variance is only 32.17% of the variance in gold price innovations. Thus, the aggregated benchmark can handle contract positions two orders of magnitude larger than current position limits, without manipulation variance being much larger than the variance in gold price innovations.

[117] Aggregation reduces manipulation variance faster than demand variance, so manipulation also represents a smaller fraction of total benchmark variance under the aggregated benchmark: the square root of demand shock-induced benchmark variance is approximately \$0.180 USD/oz, so manipulation represents only 3.51% of total benchmark variance.

D.8 Aggregated benchmark algebra

The utility function of agent i for a series of prices p_a, quantities z_{ia} is:

$$\left[\sum_{a:i\leqslant n_a} \pi_a z_{ai} - \frac{z_{ai}^2}{2\kappa_i} - p_a z_{ai}\right] + P_b y_{ci}$$

$$= \left[\sum_{a:i\leqslant n_a} \pi_a z_{ai} - \frac{z_{ai}^2}{2\kappa_i} - p_a z_{ai}\right] + y_{ci}\left[\sum_a \phi\left(n_a\right) p_a\right]$$

Where the sum is over all auctions a that agent i participates in, that is, with $i \leqslant n_a$. Since agents cannot influence prices in auctions they do not participate in, we can write the utility of agent i as:

$$= \sum_{a:i\leqslant n_a}\left[\pi_a z_{ia} - \frac{z_{ia}^2}{2\kappa_i} - p_a z_{ia} + \phi\left(n_a\right) p_a y_{ci}\right]$$

Hence, similarly to section 4, if agent i participates in auction a, she bids as if she holds $\phi_a y_{ci}$ contracts tied to the price in auction a. Equilibrium demand functions in auction a are thus:

$$z_{Di}\left(p; y_{ci}, y_{di}\right) = \phi\left(n_a\right) y_{ci}\frac{b_i}{\sum_{j\neq i} b_j} + y_{dai}\frac{b_i}{\kappa_i} - \left(p - \pi\right) b_i$$

Prices are then

$$p_a - \pi_a = \frac{1}{\sum_{i=1}^{n_a} b_i}\left[\sum_{i=1}^{n_a} \phi\left(n_a\right) y_{ci}\frac{b_i}{\sum_{j\neq i} b_j} + y_{dai}\frac{b_i}{\kappa_i}\right]$$

Thus, the aggregated benchmark P_b is:

$$P_b = \sum_{a=1}^{44} \phi\left(n_a\right)\left[\frac{1}{\sum_{i=1}^{n_a} b_i}\left[\sum_{i=1}^{n_a} \phi\left(n_a\right) y_{ci}\frac{b_i}{\sum_{j\neq i} b_j} + y_{dai}\frac{b_i}{\kappa_i}\right]\right]$$

Interchanging the order of summation on the y_{ci} term, we can write:

$$P_b = \sum_{i=1}^{n} \left[\sum_{a=1}^{44} 1\left(i \leqslant n_a\right) \left(\phi\left(n_a\right)\right)^2 \left[\frac{b_i}{\left(\sum_{j \neq i} b_j\right)\left(\sum_{i=1}^{n_a} b_i\right)} \right] y_{ci} \right] + \sum_{a=1}^{44} \phi\left(n_a\right) \left[\frac{b_i}{\kappa_i \sum_{i=1}^{n_a} b_i} \right] y_{dai}$$

Taking the variance, we have:

$$Var\left(P_b\right) = \underbrace{\sum_{i=1}^{n} \left[\left[\sum_{a=1}^{44} 1\left(i \leqslant n_a\right) \left(\phi\left(n_a\right)\right)^2 \left[\frac{b_i}{\left(\sum_{j \neq i} b_j\right)\left(\sum_{i=1}^{n} b_i\right)} \right] \right]^2 \right] \sigma_{ci}^2}_{Manipulation} + \underbrace{\sum_{a=1}^{A} \sum_{i=1}^{n} \left(\phi\left(n_a\right)\right)^2 \left[\frac{b_i}{\kappa_i \sum_{i=1}^{n} b_i} \right]^2 \sigma_{di}^2}_{Demand\ Shocks} \quad (104)$$

I calculate optimal weights $\phi\left(n_a\right)$ to minimize (104).

E Supplementary material for section 7

E.1 The V_t^2 integral

Let K represent option strike prices, and define $O_t\left(K\right)$ as the time t price of the out-of-money option at strike price K. On each date t, we observe H_t strike prices, $K_{t1} \ldots K_{t,H_t}$. Let $C_t\left(K\right)$ and $P_t\left(K\right)$ represent the call and put prices at strike price K respectively, and let r_t represent the 31-day interest rate at date t. From (10) in Carr and Wu (2006), the risk-neutral expected value of return variance is:

$$V_t^2 = \int_0^{\infty} \frac{2}{T} e^{T r_t} \frac{O_t\left(K\right)}{K^2} dK \quad (105)$$

In order to compare across dates with different levels of π_t, I normalize all prices and strikes by the forward index price F_t. With normalized quantities in lowercase,

$$k \equiv \frac{K}{F_t},\ o_t(k) \equiv \frac{O_t(\pi_t k)}{F_t}$$

Hence, (105) becomes:

$$\begin{aligned}
&\int_0^\infty \frac{2}{T} e^{Tr_t} \frac{O_t(K)}{K^2} dK \\
&= \int_0^\infty \frac{2}{T} e^{Tr_t} \frac{F_t o_t(k)}{(F_t k)^2} (F_t dk) \\
&= \int_0^\infty \frac{2}{T} e^{Tr_t} \frac{o_t(k)}{k^2} dk
\end{aligned}$$

Since we have defined $\pi_t(k) = \frac{2}{T} e^{Tr_t} \frac{o_t(k)}{k^2}$, this proves that (46) is equivalent to (105).

E.2 Data

1. **Option Metrics:** I use data for all monthly SPX options daily closing prices prices and trade volume from 2010-05-19 to 2017-12-20. I restrict my sample to monthly SPX options with maturities of 31 days. Using option prices, I estimate option-implied forward prices $\tilde{F}_t$, as described in appendix E.4, and then use only out-of-money options relative to $\tilde{F}_t$, with strike prices between 0.6 and 1.1 times $\tilde{F}_t$. I construct normalized strike prices k_{th} and normalized price vectors $\mathbf{p_t^{H_t}}$ as in (47) and (107), that is,

$$k_{th} \equiv \frac{K_{th}}{\tilde{F}_t},\ \mathbf{o_t^{H_t}}(k_{th}) = \frac{\mathbf{O_t^{H_t}}(K_{th})}{\tilde{F}_t} \tag{106}$$

$$\mathbf{p_t^{H_t}}(k_{th}) = \frac{2}{T} e^{Tr_t} \frac{\mathbf{o_t^{H_t}}(k_{th})}{k_{th}^2} \tag{107}$$

   where time to maturity T is $T = \frac{30.718}{365}$ years, and the interest rate r_t is linearly interpolated based on Option Metrics data on the daily zero coupon yield curve, derived from ICE IBA LIBOR rates. In total, I observe 92 dates with 14,984 option prices.

2. **iVolatility SPX minute-quote data:** I purchased SPX option quote data from iVolatility.com, which I use to estimate demand slopes $\boldsymbol{\kappa_t^{H_t}}(k_{th})$. I use data for all options

31 days to expiry in 2016 and 2017. The data includes minute-level best bid and ask prices and sizes for all monthly SPX options, from 9:31:00 to 16:19:00 each day. I construct option-implied forward prices using data each minute, as in appendix E.4, and I use these, together with the interest rate series r_t, to normalize prices each minute as in (106) and (107). I aggregate data for each day, taking the average bid plus ask size over all minutes for a given strike price, and the bid-plus-ask-size weighted average bid-ask spread, throughout the day.

3. **VIX settlement values:** To make figure 11, I use a time series of monthly VIX futures settlement values downloaded from the CBOE website.[118]

4. **VIX futures open interest at settlement:** Data on VIX futures open interest at settlement, from 2016 to 2017, used for estimating the size of contract positions σ_c^2, is from the CBOE website.[119]

Figure 11 shows my calculated VIX values, using the CBOE Riemann sum weights, plotted against the official monthly VIX settlement series. The series should not exactly coincide, as my VIX values are calculated using exchange closing prices on days preceding VIX settlement dates. However, the series track each other closely, indicating that my data cleaning steps do not significantly distort the level or time-series movements of VIX.

E.3 Equilibrium prices $\mathbf{p_t^{H_t}}$

Regular traders' utilities are given in (53) and the manipulator's utility is (54). As in section 4, the manipulator's utility is thus separable across submarkets, and can be written as:

$$\sum_{h=1}^{H_t}\left[\pi_t\left(k_{th}\right)z\left(k_{th}\right)+\frac{y_{td1}\left(k_{th}\right)z\left(k_{th}\right)}{\kappa\left(k_{th}\right)}-\frac{\left(z\left(k_{th}\right)\right)^2}{2\kappa\left(k_{th}\right)}-\right.$$
$$\left.\mathbf{p_t^{H_t}}\left(k_{th}\right)z\left(k_{th}\right)+\boldsymbol{\phi}_{\mathbf{t}}^{\mathbf{H_t}}\left(k_{th}\right)\mathbf{p_t^{H_t}}\left(k_{th}\right)y_{tc}\right]+Cy_{tc}$$

The manipulator acts as if she holds $\phi_t y_{tc}$ contracts in submarket k_{th}. Equilibrium in each submarket k_{th} is thus exactly as described in proposition 5 of section 4, specialized to the

[118] CBOE Final Settlement Prices.
[119] CBOE Historical Data

Figure 11: Calculated vs official VIX series

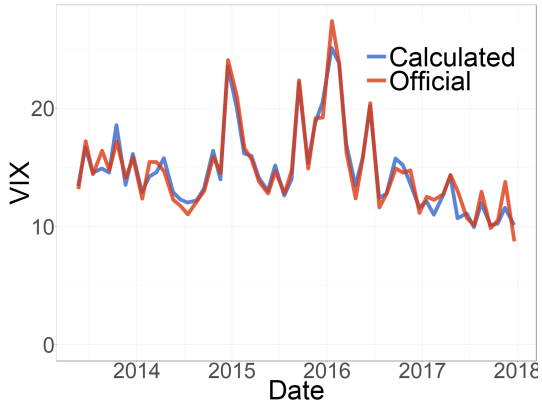

Notes. The blue line shows VIX calculated using Riemann sums, as CBOE does, using my observed vectors of option prices, $\mathbf{p}_t^{H_t}$, on afternoons 31 days before option expiration. The red line shows the official VIX settlement series, based on option prices in the morning 30 days before option expiration.

case in which each submarket has n symmetric agents. Using expressions for symmetric equilibrium outcomes from appendix A.7, regular agents' equilibrium demand functions in submarket k_{th} are:

$$z_{tDi}\left(\mathbf{p}_t^{H_t}(k_{th}), y_{tdi}(k_{th}), k_{th}\right) = \frac{n-2}{n-1} y_{tdi}(k_{th}) - \left(\frac{n-2}{n-1}\right)\left(\mathbf{p}_t^{H_t}(k_{th}) - \pi_t(k_{th})\right)\boldsymbol{\kappa}_t^{H_t}(k_{th}) \quad (108)$$

The manipulator's equilibrium demand function in submarket k_{th} is:

$$z_{tD1}\left(\mathbf{p}_t^{H_t}(k_{th}), y_{tdi}(k_{th}), y_c, k_{th}\right) = \frac{n-2}{n-1} y_{td1}(k_{th}) + \frac{1}{n-1}\boldsymbol{\phi}_t^{H_t}(k_{th}) y_{tc} - \left(\frac{n-2}{n-1}\right)\left(\mathbf{p}_t^{H_t}(k_{th}) - \pi_t(k_{th})\right)\boldsymbol{\kappa}_t^{H_t}(k_{th}) \quad (109)$$

The price in submarket k_{th} is thus:

$$\mathbf{p}_t^{H_t}(k_{th}) = \pi_t(k_{th}) + \frac{\bar{\mathbf{y}}_{dt}^{H_t}(k_{th})}{n\boldsymbol{\kappa}_t^{H_t}(k_{th})} + \frac{\boldsymbol{\phi}_t^{H_t}(k_{th})}{n(n-2)\boldsymbol{\kappa}_t^{H_t}(k_{th})} y_{tc},\ h \in \{1 \ldots H_t\}$$

as desired.

E.4 Measurement of forward price $\tilde{F}_t$

At each date, we observe vectors of put and call option prices $\mathbf{C_t^{H_t}}(K_{th})$ and $\mathbf{P_t^{H_t}}(K_{th})$. I measure the forward index level π_t following the methodology of Andersen, Bondarenko and Gonzalez-Perez (2015). Put-call parity implies that, for any K_h, we should have:

$$F_t = K_{th} + e^{Tr_t}\left(\mathbf{C_t^{H_t}}(K_{th}) - \mathbf{P_t^{H_t}}(K_{th})\right) \tag{110}$$

Given r_t and T, we can thus infer F_t from any single pair of put and call prices at the same strike K_{th}. Following Andersen, Bondarenko and Gonzalez-Perez (2015), I measure $\tilde{F}_t$ by taking the average value of F_t implied by (110), using bid prices for call options and ask prices for put options, for all strike prices where put and call prices differ by at most \$25; this encompasses approximately 10 option pairs for each date t.

E.5 Gaussian process parametrization

I parametrize the functions $\mu(k), \sigma_\pi(k), \sigma_d^2(\cdot)$ using splines. From figure 4, the derivative of $\mu(k)$ is discontinuous at the forward index level $k = 1$; thus, I model $\mu(\cdot)$ separately for $k > 1$ and $k < 1$, so:

$$\mu(k) = \begin{cases} \mu_-(k) & 0.6 \leqslant k \leqslant 1 \\ \mu_+(k) & 1 \leqslant k \leqslant 1.1 \end{cases}$$

I model each of $\mu_-(k)$ and $\mu_+(k)$ as an i-spline, with 4 interior knots each, with uniform quantile spacing between knots. I-splines are simply a convenient set of basis functions for capturing the shape of the $\mu(k)$ function. As each set of splines has 7 degrees of freedom, $\mu(k)$ is thus modelled as a sum of 14 basis functions; the i-spline basis ensures that $\mu(k)$ is equal to 0 at the boundaries $k = 0.6$ and $k = 1.1$. Let θ_μ represent the spline parameters determining $\mu(k)$.

In principle, it is economically reasonable to impose that $\mu_-(1) = \mu_+(1)$, and that $\mu(k) \geqslant 0\ \forall k$. However, leaving $\mu(k)$ unconstrained allows me to analytically optimize for the $\mu(k)$ parameters, fixing any values of other parameters $\sigma_\pi(k), \sigma_d^2(\cdot), l$; this greatly improves numerical performance. The economic restrictions $\mu_-(1) = \mu_+(1)$ and $\mu(k) \geqslant 0\ \forall k$ almost holds in my final estimates. Since VIX is an integral over $\pi_t(k)$, discontinuities in $\pi_t(k)$ are smoothened out by the integral operator, so some discontinuities in $\pi_t(k)$ are unlikely to substantially affect final results.

I model $\sigma_\pi(k)$ and $\sigma_d^2(k)$ as b-splines, with 6 interior knots uniformly distributed in k-quantile space, and thus 9 degrees of freedom. Both $\sigma_\pi(k)$ and $\sigma_d^2(k)$ are required to be positive; as the b-spline basis consists of strictly positive functions, I need only impose that all spline coefficients are positive. Let θ_π, θ_d respectively represent the spline parameters determining $\sigma_\pi(k)$ and $\sigma_d^2(k)$.

The final model for the Gaussian process for $\pi_t(k)$ and $\mathbf{p_t^{H_t}}$ thus has 33 parameters: 14 $\mu(k)$ parameters, 9 $\sigma_\pi(k)$ parameters , 9 $\sigma_d^2(k_t)$ parameters, and the length-scale parameter l.

E.6 Log likelihood

The covariance matrix of prices on observed dates is:

$$\mathbf{\Sigma_{p,t}^{H_t}} = \mathbf{\Sigma_{\pi,t}^{H_t}} + \mathbf{\Sigma_{demand,t}^{H_t}}$$

From the definitions of $\mathbf{\Sigma_{\pi,t}^{H_t}}$ and $\mathbf{\Sigma_{demand,t}^{H_t}}$, given θ_π, θ_d, we then have:

$$\mathbf{\Sigma_{p,t}^{H_t}}(h, h' \mid \mathbf{\Theta}) = \sigma_\pi(k_{th}; \theta_\pi)\,\sigma_\pi(k_{th'}; \theta_\pi)\exp\left(-\frac{(k_{th}-k_{th'})^2}{2l^2}\right) + \delta_{hh'}\sigma_d^2(k_{th}; \theta_d) \quad (111)$$

Thus, fixing parameters $\mathbf{\Theta} = \theta_\mu, \theta_\pi, \theta_d, l$, the collection of observed vectors $\mathbf{p_t^{H_t}} \equiv (p_t(k_{t1}), \ldots p_t(k_{tH_t}))$ are independent observations (over dates t) of draws from a multivariate Gaussian distribution, with mean $\boldsymbol{\mu}^{\mathbf{H_t}} \equiv (\mu(k_{t1}; \theta_\mu), \ldots \mu(k_{tH_t}; \theta_\mu))$ and covariance matrix described by (111). Log likelihood for a given date t, taking the grid of points $\{k_{t1} \ldots k_{tH_t}\}$ as exogeneous, is the multivariate normal log likelihood function, which is:

$$\log p\left(\mathbf{p_t^{H_t}} \mid \mathbf{\Theta}\right) = -\frac{1}{2}\left(\mathbf{p_t^{H_t}} - \boldsymbol{\mu}^{\mathbf{H_t}}\right)^\top \left(\mathbf{\Sigma_{p,t}^{H_t}}\right)^{-1}\left(\mathbf{p_t^{H_t}} - \boldsymbol{\mu}^{\mathbf{H_t}}\right) - \frac{1}{2}\log\left(\det\left(\mathbf{\Sigma_{p,t}^{H_t}}\right)\right)$$

The log likelihood of the entire sequence of $\mathbf{p_t}$ values over all dates is the sum over log likelihood of each date t:

$$\sum_t \log p\left(\mathbf{p_t^{H_t}} \mid \mathbf{\Theta}\right) = \sum_t \left[-\frac{1}{2}\left(\mathbf{p_t^{H_t}} - \boldsymbol{\mu}^{\mathbf{H_t}}\right)^\top \left(\mathbf{\Sigma_{p,t}^{H_t}}\right)^{-1} \left(\mathbf{p_t^{H_t}} - \boldsymbol{\mu}^{\mathbf{H_t}}\right) - \frac{1}{2}\log\left(\det\left(\mathbf{\Sigma_{p,t}^{H_t}}\right)\right)\right] \quad (112)$$

I choose values of parameters $\mathbf{\Theta}$ by maximizing log likelihood (112).

E.7 Log likelihood maximization: Numerical details

Log likelihood is given in (112). For numerical stability, I compute the log determinants, $\log\left(\det\left(\mathbf{\Sigma_{p,t}^{H_t}}\right)\right)$, using the Cholesky decomposition of $\mathbf{\Sigma_{p,t}^{H_t}}$:

$$\log\left|\mathbf{\Sigma_{p,t}^{H_t}}\right| = 2sum\left(\log\left(diag\left(chol\left(\mathbf{\Sigma_{p,t}^{H_t}}\right)\right)\right)\right)$$

I then maximize log likelihood using a three-layer nested optimization routine.

1. θ_μ: In the innermost layer, I find the optimal value of θ_μ given θ_π, θ_d, l. Maximizing likelihood is equivalent to minimizing the first term:

$$\sum_t -\frac{1}{2}\left(\mathbf{p_t^{H_t}} - \mathbf{X_{\mu,t}}\theta_\mu\right)^\top \left(\mathbf{\Sigma_{p,t}^{H_t}}\right)^{-1} \left(\mathbf{p_t^{H_t}} - \mathbf{X_{\mu,t}}\theta_\mu\right)$$

   where the mean vector $\boldsymbol{\mu}^{\mathbf{H_t}}$ is expressed as the spline matrix $\mathbf{X_{\mu,t}}$,[120] multiplied by spline parameters θ_μ. This is a least-squares problem, with analytical solution:

$$\theta_\mu = \left(\left(\mathbf{X_\mu}\right)^\top \left(\mathbf{\Omega}\right)^{-1} \mathbf{X_\mu}\right)^{-1} \left(\mathbf{X_\mu}\right)^\top \left(\mathbf{\Omega}\right)^{-1} \mathbf{p} \quad (113)$$

   Where $\mathbf{X_\mu}, \mathbf{p}, \mathbf{\Omega}$ represent respectively the spline matrix, the price vector, and the covariance matrix of errors respectively, across all dates and strike prices. Since I assume prices are independent across dates, the covariance matrix $\mathbf{\Omega}$ is sparse, with all covariances of elements at distinct dates t equal to 0. Thus, the terms in (113)

[120] Note that $\mathbf{X_{\mu,t}}$ is a matrix, not a vector, violating somewhat my notation principles.

simplify to:

$$(X_\mu)^\top (\Omega)^{-1} p = \sum_t (X_{\mu,t})^\top \left(\Sigma_{p,t}^{H_t}\right)^{-1} p_t^{H_t}$$

$$(X_\mu)^\top (\Omega)^{-1} X_\mu = \sum_t (X_{\mu,t})^\top \left(\Sigma_{p,t}^{H_t}\right)^{-1} X_{\mu,t}$$

Thus, using (113), I can analytically solve for θ_μ, given any choices of θ_π, θ_d, l. Plugging these θ_μ values into (112) gives the log likelihood for any choices of θ_π, θ_d, l.

2. θ_π, θ_d: In the second layer of the optimization, fixing any l, I use a gradient descent optimization procedure over the 18 parameters θ_π, θ_d, solving for θ_μ at each step using the innermost layer of the optimization loop. Plugging these $\theta_\mu, \theta_\pi, \theta_d$, values into (112) gives log likelihood for any l.

3. l: In the outermost layer of the optimization, I run matlab's fminbnd function maximize the likelihood calculated in the second optimization layer with respect to l. This is a one-dimensional minimization algorithm which is based on golden-section search and parabolic interpolation, hence is robust to nonconvexity. This is needed because the likelihood function with respect to l, after optimizing for all other parameters, appears to have a unique maximizer in l, but is nonconvex.

E.8 Estimating demand slopes, $\kappa_t^{H_t}$

Under these assumptions, let $a_{tk_{th}}$ and $b_{tk_{th}}$ be the normalized ask and bid prices of options of strike price k_{th} at date t. Let $z_{atk_{th}}$ and $z_{btk_{th}}$ represent ask and bid volume respectively. I assume that:

$$\sum_{i=1}^{n} z_D\left(a_{tk_{th}}\right) = -z_{atk_{th}}$$

$$\sum_{i=1}^{n} z_D\left(b_{tk_{th}}\right) = z_{btk_{th}}$$

In words, similarly to section 6 I assume that the top-of-book bid and ask values accurately reflect observations of aggregate demand. This would be true if all agents submitted bid and ask quotes throughout the course of the day representing their true supply and demand functions in a static auction, and the buy and sell volume at the best bid and ask

quotes respectively reflect the sum of all market participants' demand functions at bid and ask prices. This is unlikely to be true; Hasbrouck (2007, ch. 9) suggests that this is likely to substantially underestimate demand slopes.

Given these assumptions, from (108) and (109), aggregate demand is:

$$\sum_{i=1}^{n_k} z_D\left(\mathbf{p}_t^{H_t}(k_{th})\right) = \frac{n-2}{n-1}\bar{\mathbf{y}}_{dt}^{H_t}(k_{th}) + \frac{1}{n-1}\boldsymbol{\phi}_t^{H_t}(k_{th})\, y_c - n\left(\frac{n-2}{n-1}\right)\left(\mathbf{p}_t^{H_t}(k_{th}) - \pi(k_{th})\right)\boldsymbol{\kappa}_t^{H_t}(k_h)$$

Thus,

$$\frac{z_{atk_{th}} + z_{btk_{th}}}{a_{tk_{th}} - b_{tk_{th}}} = n\left(\frac{n-2}{n-1}\right)\boldsymbol{\kappa}_t^{H_t}(k_h)$$

Hence, we can estimate $\boldsymbol{\kappa}_t^{H_t}(k_h)$ as:

$$\boldsymbol{\kappa}_t^{H_t}(k_h) = \left(\frac{n-1}{n(n-2)}\right)\frac{z_{atk_{th}} + z_{btk_{th}}}{a_{tk_{th}} - b_{tk_{th}}} \tag{114}$$

To implement (114) in data, I calculate the average bid and ask volume $z_{atk_{th}} + z_{btk_{th}}$ over all minutes in the trading day, and average bid-ask spread $a_{tk_{th}} - b_{tk_{th}}$, weighted by bid plus ask volume, over all minutes; I use the ratio of these averages to calculate $\boldsymbol{\kappa}_t^{H_t}(k_h)$ using 114.

Figure 12 shows $\boldsymbol{\kappa}_t^{H_t}$, and the two components that it is based on, the bid and ask volume $z_{atk_{th}} + z_{btk_{th}}$ and the bid-ask spread $a_{tk_{th}} - b_{tk_{th}}$, for an example date, 2017-09-19. We can make a few observations. First, all series are very unsmooth; this is partially explained by the fact that, as noted in figure 12, options at round-numbered strike prices have lower bid-ask spreads and larger size than options at other strikes. In particular, besides the round-numbered strikes of 1750 and 2000, options at strike prices which are multiples of 25 have on average higher size and lower bid-ask spread than options not multiples of 25. This is because long-dated options are only traded at strike prices which are multiples of 25, and options at strikes which are multiples of 5 are only introduced around three months before expiration; plausibly, this causes market participants to cluster at options with strikes that are multiples of 25 for liquidity reasons.

Second, the estimated slope of demand is in fact lower for $k \approx 1$. This is surprising, given that close-to-the-money options are traded so often. Looking at the red and blue

Figure 12: Quote size and bid-ask spread vs strike price

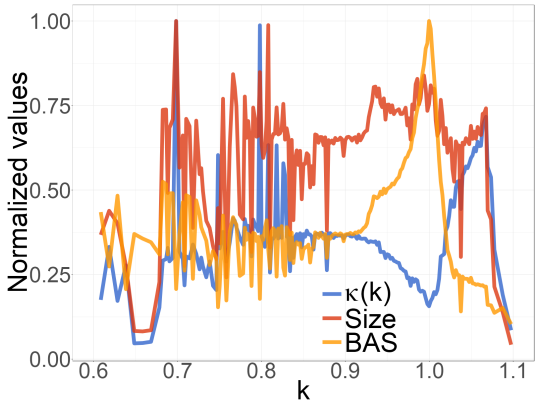

Notes. The blue line shows the estimated slope of demand $\boldsymbol{\kappa}_t^{H_t}(k_h)$, the red line shows bid plus ask volume $z_{atk_{th}} + z_{btk_{th}}$, and the yellow line shows the bid-ask spread $a_{tk_{th}} - b_{tk_{th}}$, for the date 2017-09-19. All values are normalized so the maximum value of each series is 1.

lines in figure 12, size is approximately constant for values of k close to 1, but the bid-ask spread is approximately 4 times higher for $k \approx 1$ than k further away from 1. Delving further into this, figure 13 shows bid-ask spread (yellow), average option price (green), and bid-ask spread as percentage of price (purple), as a function of strike price. The yellow line is much smaller near 1; bid-ask spread as a percentage of option price is lowest for close-to-the-money options with $k \approx 1$, which is intuitive. However, as in figure 4 in the main text, the average option price (green) is much higher for $k \approx 1$. Thus, the bid-ask spread, in units of $\mathbf{p}_t^{H_t}(k_{th})$, is in fact higher for close-to-the-money options. Intuitively, since V_t^2 integrates over option prices, manipulators care about absolute, not relative, price impact, and this is actually higher for options with $k \approx 1$ because both average option prices and bid-ask spreads are higher.

It is possible that my estimates of size in figure 12 are flawed – that markets are in fact much deeper than volume at the best bid and offer suggest. However, if markets are deeper for k close to 1, so that $\boldsymbol{\kappa}_t^{H_t}(k_h)$ is in fact larger than my estimates, this would only increase optimal submarket weights for $k \approx 1$, strengthening the main conclusion of this section that manipulation-sensitive VIX weights should put higher weight on close-to-the-money options with $k \approx 1$.

Figure 13: Bid-ask spread, average option price, and percentage bid-ask spread

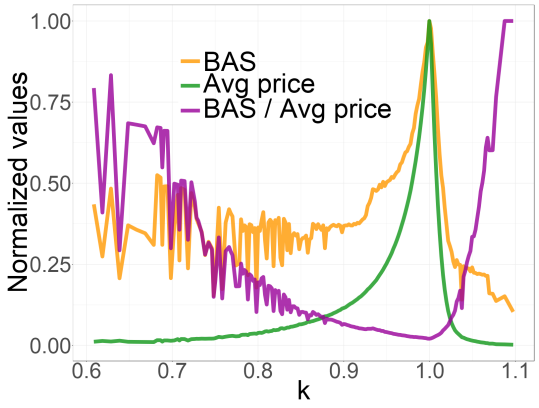

Notes. The yellow line shows the bid-ask spread $a_{tk_{th}} - b_{tk_{th}}$, the green line shows the average option price $\mathbf{p}_t^{H_t}(k_{th})$, and the purple line shows the bid-ask spread divided by the average price, for the date 2017-09-19. All values are normalized so the maximum value of each series is 1.

E.9 Estimating VIX contract position size, σ_c^2

Using the CBOE website's data on VIX futures open interest at settlement, on average from 2016-2017, there are 48,552 open VIX futures contracts on settlement days. I assume that there is a single manipulator holds a position of size roughly 10% of open interest, so I set the size of VIX contracts to $\sigma_{VIX}^2 = (4855.2)^2$. Since my benchmark represents V_t^2, I must convert the VIX futures position into an equivalent V_t^2 position. I do this using linear approximation. VIX corresponds to 100 times the the square root of V_t^2:

$$VIX = 100\sqrt{V_t^2}$$

Differentiating, we have:

$$\frac{dVIX}{dVIX^2} = \frac{100}{2\sqrt{V_t^2}}$$

The mean value of $\sqrt{V_t^2}$ (estimated using the CBOE formula) over the time period covered by my data is 0.168, so I set:

$$\frac{dVIX}{dV_t^2} \approx \frac{100}{2\,(0.168)}$$

Thus, a single VIX contract is worth approximately 298 V_t^2 contracts. Also, the contract multiplier for VIX futures is \$1000,[121] compared to \$100 for SPX options,[122] so we need to multiply V_t^2 contract positions by another factor of 10. Thus, I set the standard deviation of the size of contract positions to:

$$\sqrt{\sigma_c^2} \approx \left(\frac{100}{2\,(0.168)}\right)(10)\sqrt{\sigma_{VIX}^2} = 14,450,248$$

E.10 Bayesian quadrature history

The procedure that I use – estimating an integral by assuming that the function is drawn from a Gaussian process prior, and then calculating the expectation of the integral over the posterior – is referred in the literature to alternatively as *Bayesian quadrature*, or Bayesian Monte Carlo, or Bayes-Hermite quadrature. This procedure has recently received some attention in the machine learning literature, from researchers studying Gaussian processes.[123] The method has a long history; O'Hagan (1991) derives analytical formulas and calculates simulation results, and Diaconis (1988) surveys the literature, tracing the origins of the method back to Poincaré (1896).

Despite its long history, to my knowledge, there are relatively few applications of Bayesian quadrature, and most existing applications focus on the problem of generating quadrature rules for exactly observed functions. In my setting, I apply Bayesian quadrature to approximate the integral of a function which is observed with error; I am not aware of other settings in which Bayesian quadrature is used for purpose. Qualitatively, the Bayesian quadrature formula, (60), lowers weights on function observations with higher errors, and increases weights on nearby function observations to compensate, in a manner which is controlled by the Gaussian process length-scale or smoothness parameter l.

E.11 Bayesian quadrature algebra

We wish to approximate the integral over the posterior mean function on the RHS of (58), that is,

$$\int_{0.6}^{1.1} E\left[\pi_t(k) \mid \mathbf{p}_t^{\mathbf{H}_t}\right] dk \tag{115}$$

[121] VIX Contract Specifications
[122] SPX Options Product Specifications
[123] See, for example, (Ghahramani and Rasmussen, 2003)

We can numerically approximate (115) using a Riemann sum (57) as follows. Fix a uniformly spaced grid of evaluation points $\{k_1 \dots k_G\}$ on the interval $[0.6, 1.1]$, and define the sum vector:

$$\mathbf{S}^{\mathbf{G}} \equiv \left(\frac{1.1 - 0.6}{G}\right) \mathbf{1}^{\mathbf{G}}$$

Where $\mathbf{1}^{\mathbf{G}}$ is a length-G vector with all elements equal to 1. In my calculations I use a grid with $G = 2000$, which is on average over 10 times as dense as observed option strike prices. Let $\boldsymbol{\pi}_t^{\mathbf{G}}$ represent the vector of values of $\pi_t(k)$ on the grid $\{k_1 \dots k_G\}$. We do not observe $\boldsymbol{\pi}_t^{\mathbf{G}}$, but we can calculate its posterior expectation given the observed vector of option prices, $\mathbf{p}_t^{\mathbf{H}_t}$, that is, $E\left[\boldsymbol{\pi}_t^{\mathbf{G}} \mid \mathbf{p}_t^{\mathbf{H}_t}\right]$. The Riemann sum approximation to (115) is then:

$$\int_{0.6}^{1.1} E\left[\pi_t(k) \mid \mathbf{p}_t^{\mathbf{H}_t}\right] dk \approx \left(\mathbf{S}^{\mathbf{G}}\right)^{\top} E\left[\boldsymbol{\pi}_t^{\mathbf{G}} \mid \mathbf{p}_t^{\mathbf{H}_t}\right] \tag{116}$$

Under the Gaussian process prior, $\mathbf{p}_t^{\mathbf{H}_t}$ and $\boldsymbol{\pi}_t^{\mathbf{G}}$ are jointly normally distributed, hence the expectation $E\left[\boldsymbol{\pi}_t^{\mathbf{G}} \mid \mathbf{p}_t^{\mathbf{H}_t}\right]$ can be calculated using Gaussian projection formulas. Specifically, (51) and (56) imply that the joint distribution of $\mathbf{p}_t^{\mathbf{H}_t}, \boldsymbol{\pi}_t^{\mathbf{G}}$ is:

$$\begin{pmatrix} \mathbf{p}_t^{\mathbf{H}_t} - \boldsymbol{\mu}^{\mathbf{H}_t} \\ \boldsymbol{\pi}_t^{\mathbf{G}} - \boldsymbol{\mu}^{\mathbf{G}} \end{pmatrix} \sim N\left(0, \begin{bmatrix} \boldsymbol{\Sigma}_{p,t}^{\mathbf{H}_t} & \boldsymbol{\Sigma}_{p\pi,t}^{\mathbf{H}_t,\mathbf{G}} \\ \boldsymbol{\Sigma}_{\pi p,t}^{\mathbf{G},\mathbf{H}_t} & \boldsymbol{\Sigma}_{\pi,t}^{\mathbf{G}} \end{bmatrix}\right) \tag{117}$$

$\boldsymbol{\mu}^{\mathbf{G}}$ is the vector of values of the mean function $\mu(k)$ evaluated on the grid G, that is, $\boldsymbol{\mu}^{\mathbf{G}}(k_g) = \mu(k_g)$. $\boldsymbol{\Sigma}_{p,t}^{\mathbf{H}_t}$ and $\boldsymbol{\Sigma}_{\pi,t}^{\mathbf{G}}$ are as defined in (51) and (56), that is,

$$\boldsymbol{\Sigma}_{p,t}^{\mathbf{H}_t}(h, h') = \sigma_\pi(k_{th})\,\sigma_\pi(k_{th'}) \exp\left(-\frac{(k_{th} - k_{th'})^2}{2l^2}\right) + \delta_{hh'}\sigma_d^2(k_{th}) + \left(\frac{\phi_{k_{th}}}{n(n-2)\,\boldsymbol{\kappa}_t^{\mathbf{H}_t}(k_{th})}\right)\left(\frac{\phi_{k_{th'}}}{n(n-2)\,\boldsymbol{\kappa}_t^{\mathbf{H}_t}(k_{th'})}\right)\sigma_c^2$$

$$\boldsymbol{\Sigma}_{\pi,t}^{\mathbf{G}}(g, g') = \sigma_\pi(k_g)\,\sigma_\pi(k_{g'}) \exp\left(-\frac{(k_g - k_{g'})^2}{2l^2}\right)$$

The rectangular matrices $\boldsymbol{\Sigma}_{p\pi,t}^{\mathbf{H}_t,\mathbf{G}}, \boldsymbol{\Sigma}_{\pi p,t}^{\mathbf{G},\mathbf{H}_t}$ have elements defined by:

$$\boldsymbol{\Sigma}_{p\pi,t}^{\mathbf{H}_t,\mathbf{G}}(h, g) = \boldsymbol{\Sigma}_{\pi p,t}^{\mathbf{G},\mathbf{H}_t}(g, h) = \sigma_\pi(k_{th})\,\sigma_\pi(k_g) \exp\left(-\frac{(k_{th} - k_g)^2}{2l^2}\right)$$

The posterior distribution of the vector $\boldsymbol{\pi}_t^G - \boldsymbol{\mu}^G$, conditional on observing the vector $\mathbf{p}_t^{H_t} - \boldsymbol{\mu}^{H_t}$, is then normal, with mean and variance :

$$\mathsf{E}\left[\left(\boldsymbol{\pi}_t^G - \boldsymbol{\mu}^G\right) \mid \mathbf{p}_t^{H_t}\right] = \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\left(\mathbf{p}_t^{H_t} - \boldsymbol{\mu}^{H_t}\right) \tag{118}$$

$$\mathsf{Var}\left[\left(\boldsymbol{\pi}_t^G - \boldsymbol{\mu}^G\right) \mid \mathbf{p}_t^{H_t}\right] = \boldsymbol{\Sigma}_{\pi,t}^{G} - \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\boldsymbol{\Sigma}_{p\pi,t}^{H_t,G} \tag{119}$$

Rearranging (118), we can write:

$$\mathsf{E}\left[\boldsymbol{\pi}_t^G \mid \mathbf{p}_t^{H_t}\right] = \boldsymbol{\mu}^G + \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\left(\mathbf{p}_t^{H_t} - \boldsymbol{\mu}^{H_t}\right) \tag{120}$$

Finally, plugging into (116), we have:

$$\int_{0.6}^{1.1} \mathsf{E}\left[\pi_t(k) \mid \mathbf{p}_t^{H_t}\right] dk \approx \left(\mathbf{S}^G\right)^\top \left(\boldsymbol{\mu}^G - \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\boldsymbol{\mu}^{H_t} + \boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{p,t}^{H_t}\right)^{-1}\mathbf{p}_t^{H_t}\right) \tag{121}$$

E.12 Posterior inference algebra and computation

Define the vector function $\boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)$,[124] as a vector with elements:

$$\boldsymbol{\eta}_t^{H_t}(k_h) = \frac{\boldsymbol{\phi}_t^{H_t}(k_h)}{n(n-2)\,\boldsymbol{\kappa}_t^{H_t}(k_h)}$$

We can write the fixed point problem as:

$$\left(\boldsymbol{\phi}_t^{H_t}\right)^\top = \left(\mathbf{S}^G\right)^\top \left(\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t} + \sigma_c^2 \boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\left(\boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\right)^\top\right)^{-1}\right) \tag{122}$$

Define an operator $\mathcal{T}(\boldsymbol{\phi})$ as the RHS of (122):

$$\mathcal{T}\left(\boldsymbol{\phi}_t^{H_t}\right) = \left(\mathbf{S}^G\right)^\top \left(\boldsymbol{\Sigma}_{\pi p,t}^{G,H_t}\left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t} + \sigma_{ct}^2 \boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\left(\boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\right)^\top\right)^{-1}\right)$$

[124] This violates my vector notation principles, but is the only instance in the text where I need to explicitly use a vector-valued function.

We can speed up the computation of $\mathcal{T}\left(\boldsymbol{\phi}_t^{H_t}\right)$ by using the Sherman-Morrison formula for matrix inverses, which states that:

$$\left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t} + \sigma_{ct}^2 \boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\left(\boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\right)^T\right)^{-1}$$

$$= \left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t}\right)^{-1} - \frac{\left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t}\right)^{-1} \boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\left(\boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\right)^T \left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t}\right)^{-1} \sigma_{ct}^2}{1 + \left(\boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right)\right)^T \left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t}\right)^{-1} \boldsymbol{\eta}_t^{H_t}\left(\boldsymbol{\phi}_t^{H_t}\right) \sigma_{ct}^2}$$

The matrix $\left(\boldsymbol{\Sigma}_{\pi,t}^{H_t} + \boldsymbol{\Sigma}_{demand,t}^{H_t}\right)^{-1}$ does not depend on $\boldsymbol{\phi}_t^{H_t}$, and can be precomputed.

In order for $\boldsymbol{\phi}_t^{H_t}$ to constitute manipulation-consistent Bayesian quadrature weights, as in definition 4, it is equivalent that $\boldsymbol{\phi}_t^{H_t}$ satisfies:

$$\boldsymbol{\phi}_t^{H_t} = \mathcal{T}\left(\boldsymbol{\phi}_t^{H_t}\right) \tag{123}$$

I was unable to solve (123) using the standard minimum-distance gradient descent approach, so I implemented a step-down iteration scheme. Pseudocode for the scheme is:

```
function fixedpointsolver(phistart):
t = 0.1
philist(1) = phistart
while absdiff > tol:

    for(i in 2:N):

        philist(i) = (1-t)*philist(i)+ t * T(philist(i-1))

    if(stuckcheck(philist)):

        t=t/2

    absdiff = sum(abs(T(philist(1)) - philist(1)))
    philist(1) = philist(N)

return philist(1)
```

Intuitively, the `fixedpointsolver()` function slowly adjusts a guess vector $\boldsymbol{\phi}_t^{H_t}$ towards $\mathcal{T}\left(\boldsymbol{\phi}_t^{H_t}\right)$, according to the rule:

$$\boldsymbol{\phi}_{t,new}^{H_t} = (1-t)\,\boldsymbol{\phi}_{t,old}^{H_t} + t\mathcal{T}\left(\boldsymbol{\phi}_{t,old}^{H_t}\right)$$

After N iterations, a function `stuckcheck(philist)` is run on the list of the previous N iterates, `philist`. This function, whose pseudocode I do not include, runs a regression of certain fixed elements of `philist` on the iteration index `i`, and calculates the R^2 value; if R^2 is high, the iteration is likely proceeding smoothly, whereas if R^2 is low the iteration is likely stuck in a cycle. Hence, if R^2 is low, the step size `t` is lowered and the iteration continues. The function concludes when `absdiff`, the L1 difference between $\boldsymbol{\phi}_t^{H_t}$ and $\mathcal{T}\left(\boldsymbol{\phi}_t^{H_t}\right)$, is below `tol`.