Derived Pricing: Fragmentation, Efficiency, and Manipulation

Craig Pirrong
Bauer College of Business
University of Houston

January 14, 2019

1 Introduction

The use of the price in one transaction to determine the price of another–"derived pricing"–is extremely common in financial markets. This article explores two particular types of derived pricing: "trade at settlement" contracts, and off-exchange trading through crossing networks where transactions prices are based on subsequent prices in "lit" exchange auction markets. Adaptations of well-known microstructure models to these particular types of trading produce several interesting results.

First, when (a) private information is short-lived, and (b) some uninformed traders can utilize these forms of trading, but others cannot, derived pricing causes market fragmentation: the uninformed traders who can utilize them do so but the privately informed do not. Thus, derived pricing leads to a partial separating equilibrium between informed and uninformed traders.

Second, this "cream skimming" fragmentation lowers the cost of the uninformed traders who can utilize derived pricing but raises the liquidity costs of those who cannot. However, total adverse selection costs for uninformed traders decline.

Third, derived pricing creates the potential for trade-based manipulation strategies. Since the market using derived pricing serves only uninformed traders, whereas the exchange market serves both informed and uninformed traders, the price impact of trades is smaller in the derived pricing instrument than on exchange. This disparity in price impact allows a manipulator to profit by accumulating a position in the derived pricing market, and then affect the value of that position by trading in the exchange market.

The remainder of this article is organized as follows. Second 2 presents a model of trading in trade-at-settlement ("TAS") contracts in the absence of manipulation. Section 3 shows how the existence of TAS markets create manipulative opportunities, and how such manipulations affect prices. Section 4 presents an example of manipulation of TAS contracts by the trading firm Optiver in 2008. Section 5 adapts the model to crossing networks. Section discusses the implications of the models for the regulation of markets in the United States. Section 7 summarizes.

2 Trading-at-Settlement Contracts

Trading-at-Settlement ("TAS") contracts are common. For example, the CME Group offers TAS trading on 46 different futures contracts. The trader enters into a TAS contract prior to the determination of the futures settlement price (which, for instance, is based on the volume-weighted-average price between 2:28 and 2:29:59 for actively traded crude oil futures contracts).

In CME Group TAS contracts, the buyer (seller) pays (receives) a price equal to the official settlement price, plus or minus four ticks: the differential to the settlement price is negotiated between the TAS buyer and seller in the Globex continuous auction market.[1] The possibility of a differential means that there can be a bid/ask spread in the TAS market, and that TAS trades can have price impact in the TAS market, as well as in the market for the underlying futures.

To model the TAS market, I adapt the well-known Admati-Pfleiderer (1988) ("AP88") model of discretionary trade timing. As in AP88, there are multiple trading periods t in a day, $t = 1, 2, \ldots, T$. In my adaptation of the model, the last trading period T in a day is the settlement period, and the equilibrium price from this last round of trading is the settlement price used to determine the price of TAS contracts. The value of the asset at T is:

$$\tilde{V} = \mu + \sum_{t=1}^{T} \tilde{\delta}_t \tag{1}$$

where $\tilde{\delta}_t$ are independent random variables. For simplicity, the $\tilde{\delta}_t$ are normal, with mean zero and variance σ^2 for all t.

Immediately prior to each t, N informed traders receive signals about $\tilde{\delta}_t$. Informed trader i receives a signal $s_{it} = \tilde{\delta}_t + \epsilon_{it}$, where the ϵ_{it} are i.i.d. For simplicity, the ϵ_{it} is a normal variate, with variance ϕ for all i and t. Prior to the next round of trading at $t+1$, $\tilde{\delta}_t$ is publicly revealed. Thus, informed traders' information is short-lived.[2] At time t, informed trader i trades x_{it} contracts to maximize his profits.

[1] CME Rulebook, Rule 524.

[2] As I discuss below, similar results would obtain if informed traders' information is long-lived, but there are many informed traders so competition between them is intense. See Holderness and Subrahmanyam (1992).

In addition to informed traders, there is a continuum of uninformed liquidity traders located along the line segment $[0, 1]$. There are two types of uninformed traders. In each period, there are non-discretionary liquidity (noise) traders who must trade a given quantity in that period. Non-discretionary trader j must trade Y_{jt} units at time t, where Y_{jt} is a normal random variable with mean zero . Furthermore, there are discretionary traders who can choose to trade in the market at t and pay/receive the equilibrium price at t, or to buy/sell a TAS contract. If j trades in the TAS market at t, she trades Y_{jt} in that market, where Y_{jt} is a random variable. The purchase/sales price for the TAS contract equals the market price at time T, plus/minus a differential F that is determined in the TAS market at t.

For each t, the variance of $\int_0^1 Y_{jt} dj = S$. Furthermore, for each t, the fraction of discretionary traders is ρ.

Finally, there is a risk neutral, perfectly competitive market maker. The market maker observes the total net order flow ω_t in the market for the instrument, which is the sum of the informed trader orders and the liquidity traders' orders. The market maker also observes net order flow ν_t in the TAS market. At t the market maker chooses a linear pricing schedule in the "lit" exchange market for the instrument:

$$P_t = \mu + \sum_{\tau=1}^{t} \tilde{\delta}_t + \lambda_t \omega_t \tag{2}$$

The market maker also chooses a linear pricing schedule in the TAS market:

$$U_t = \gamma_t \nu_t \tag{3}$$

Assume initially that there is no TAS market, so that neither discretionary liquidity traders nor informed traders can trade in the TAS market. Then, the results of AM88 imply that the equilibrium involves:

$$\lambda_t = \frac{\sigma^2}{N+1}\sqrt{\frac{N}{S(\sigma^2+\phi)}} \tag{4}$$

Furthermore,

$$x_{it} = \beta s_{it} \tag{5}$$

with

$$\beta = \sqrt{\frac{S}{N(\sigma^2+\phi)}} \tag{6}$$

Now permit trading on the TAS market. First note that no informed trader will trade there for any $t < T$ because her information will become public before the settlement price is determined, meaning that the settlement price will incorporate this information. Therefore, it is not possible to trade profitably on a $t < T$ signal in the TAS market.

This immediately implies that the market maker is willing to set $\gamma_t = 0$. This is true because the expected profit for the market maker of trading the TAS contract is zero. At t, the expected price at T is:

$$E_t(P_T) = E_t[\mu + \sum_{\tau=1}^{T}\tilde{\delta}_t + \lambda_t\omega_t] = \mu + \sum_{\tau=1}^{t-1}\tilde{\delta}_\tau \tag{7}$$

If the market maker buys a TAS, he expects to pay $E_t(P_T)$, and expects to receive $E(\tilde{V}) = \mu + \sum_{\tau=1}^{t-1}\tilde{\delta}_\tau$.

Now consider discretionary liquidity trader j. If he trades in the exchange market, he expects to pay a liquidity cost of $\lambda_t Y_j^2$ (the product of the expected price $\lambda_t Y_j$ and the quantity Y_j). However, as just shown, the liquidity cost in the TAS market is zero because $\gamma_t = 0$. Therefore, in equilibrium, all discretionary traders choose the TAS market.

Since no informed traders are present in the TAS market in the exchange market, N remains unchanged. Furthermore, since all discretionary traders choose the TAS market, the variance of liquidity trader order flow in the exchange market is $S(1-\rho)$. Thus, in the new equilibrium:

$$\lambda_t^{TAS} = \frac{\sigma^2}{N+1}\sqrt{\frac{N}{S(1-\rho)(\sigma^2+\phi)}} \tag{8}$$

$$\beta = \sqrt{\frac{S(1-\rho)}{N(\sigma^2+\phi)}} \tag{9}$$

Thus, TAS trading causes depth in the exchange market to decline (i.e., price impact of order flow increases), and the intensity of informed trading to fall as well. The decline in the depth of the exchange market hurts the non-discretionary liquidity traders who must trade there. However, total liquidity costs decline due to the introduction of TAS trading. Without TAS, total liquidity cost is:

$$C^0 = \lambda_t S = \frac{\sigma^2}{N+1}\sqrt{\frac{NS}{\sigma^2+\phi}} \tag{10}$$

With TAS, total liquidity cost is:

$$C^{TAS} = \rho \cdot 0 + \lambda_t^{TAS}(1-\rho)S = \frac{\sigma^2}{N+1}\sqrt{\frac{NS(1-\rho)}{\sigma^2+\phi}} < C^0 \tag{11}$$

Thus, TAS trading has distributive effects. It makes discretionary liquidity traders better off, and non-discretionary traders worse off. On net, the liquidity traders are better off: this benefit to liquidity traders as a whole comes at the expense of informed traders, whose profitability declines because they have less uninformed order flow to trade against–and profit from.

TAS therefore fragments markets. Although the term fragmentation is often used pejoratively in securities and derivatives markets, here it is beneficial for liquidity traders as a whole because it reduces adverse selection.

The introduction reduces the informativeness of prices because it reduces the intensity of informed trading. However, given that by assumption private information becomes public in short order, this is of little relevance because real investment decisions, or other decisions that are made conditional on securities prices are little affected, or totally unaffected, by the slight delay in the revelation of information. Indeed, given that very short-lived private information has little social value (in terms of improving real decisions), but can be costly to collect, this type of privately informed trading is a form of rent seeking. By reducing the profitability of such trading, TAS markets reduce the incentive to spend real resources in order to gain such a fleeting, and socially useless, information advantage.

Now consider the case of long-lived information (i.e., information that is not made public prior to T). Here the role of TAS will depend on the amount of competition between informed traders. Consider the case where there is a single informed trader, i.e., $N = 1$. Here the equilibrium derived above will not hold. If γ_t=0, the informed trader could profit by trading TAS because the settlement price would not reflect all her information: indeed, she can determine how much of the information is revealed in the price at T. With $N = 1$, introduction of TAS cannot lead to a partial separating equilibrium in which informed traders eschew TAS trading. Indeed, TAS trading offers no benefit to discretionary uninformed traders because in equilibrium $\gamma_t = \lambda_t$. If this were not the case, the single informed trader would shift her trading to the market with the greater depth.

Next consider the case where information is long-lived, but N is large. As shown by Holderness and Subrahmanyam (1992), with large N, competition between informed traders causes information to be incorporated into prices almost immediately even if that information is long-lived. Thus, competition between a large number of informed traders restores the semi-separating TAS equilibrium for t sufficiently below T. Due to this competition, even long-lived information produced sufficiently long before T will be incorporated into prices by T. Thus, informed traders have no incentive to trade TAS at such a $t < T$ even if γ_t is zero (or positive, but less than λ_t) because the price they pay will fully reflect their information.

Thus, TAS leads to a semi-separating equilibrium in which informed traders do not use this order type when information is rapidly reflected in prices, either because it is disclosed rapidly to the public, or because there is sufficient competition between informed traders. TAS facilitates "cream skimming" of a portion of uninformed order flow, to the detriment of the remainder of uninformed order flow and informed traders. On balance, however, liquidity traders are better off: discretionary traders gain more than non-discretionary ones lose.

This means that the survival of TAS trading in a particular market sheds light on the nature of private information. Specifically, a functioning TAS market indicates that private information for the underlying instrument is short-lived.

There are several other implications of this analysis. First, derived pricing such as TAS is most likely to occur in large volumes when information is short lived or there is substantial competition between informed traders: it is not viable when a small number of traders possess long-lived information. Second, the volume of TAS activity should depend on the time before the settlement price is determined: the potential for adverse selection in TAS, and hence the viability of TAS trading as a mechanism for inducing self-selection based on information, increases as the time until settlement declines. Thus, TAS trading volume should decline through the day. Third, the impact of TAS trades on the price of the underlying should be smaller than the price impact of regular trades that take place in an exchange auction because TAS traders are uninformed, whereas some traders on exchange are informed.

3 Manipulation With Derived Pricing

Trade-based manipulation (i.e., manipulative strategies that exploit the impact of trades on prices) has been the subject of much controversy and some academic research. As Jarrow (1992) pointed out, trade-based manipulation must overcome a basic problem: even if a manipulator succeeds in driving up prices through purchases, how can he profit if sales drive down prices by as much or more?[3] Thus, profitable trade-based manipulations require purchases and sales have asymmetric price responses.

A variety of articles propose different sources of asymmetry. In one example, Allen and Gorton (1992) posit that purchases are more likely to be driven by private information than sales, so purchases will have a larger price impact than sales: they show that this asymmetry can result in profitable trade-based manipulation strategies. In another example, Kumar and Seppi (1992) present a model in which a cash-settled derivatives contract trades

[3]If market makers are risk averse, or incur trade processing costs, if the permanent price impacts of purchases and sales is identical, trade-based manipulative strategies will lose money because the manipulator must compensate the market maker for trade processing or inventory/risk costs.

before private information is revealed, and the underlying instrument whose price is used to determine the settlement price of the derivative is traded when an individual or individuals has private information. With this asymmetry, a manipulator can make a randomized purchase or sale in the derivatives market with little price impact (because there is no private information then), and then trade in the same direction in underlying market when the derivative is settled: due to the presence of private information at this time, this transaction moves the settlement price of the derivative in a way that profits (on average) the manipulator.

The analysis in Section 2 demonstrates that TAS contracts create trading opportunities with asymmetric price impacts. This suggests that TAS may therefore also create opportunities for profitable trade-based manipulation, and this is indeed the csae.

To show this formally, consider a simplification of the model in Section 2. A single trader receives an informative signal at two times, t_0 and T. The information revealed (possibly noisily) to the informed trader at t_0 is revealed publicly at t_1, $t_0 < t_1 < T$. As before, there is a continuum of liquidity traders, and at t_0 fraction ρ of these traders can trade TAS: the remainder must trade on the exchange at t_0. There is a continuum of noise traders at T. As in Section 2, the variance of noise trader order flow is S at both dates; the variance of the fundamental information is σ^2 at both dates, and the variance of signals is ϕ at both dates: for simplicity, I assume that the signals are perfectly informative so $\phi = 0$. The TAS contract settles based on the market clearing price at T.

In addition to the liquidity traders and informed trader, and a risk neutral market maker, there is a manipulator. As in Kumar-Seppi (1992) this trader is wealth constrained and subject to margin requirements which limits the size of a trade he can make. The manipulator's wealth is $|W|$. The market maker knows the distribution of W, which is $N(0, \sigma_W^2)$.

Due to the public revelation of t_0 information prior to T, the informed traders do not trade TAS at t_0. Further, as before, all discretionary liquidity traders utilize the TAS market, meaning that there will be order flow in the TAS market. The manipulator can trade in the TAS market at t_0, and then trade on the exchange market at T: trades on the lit market impact the price then, and therefore affect the price the manipulator pays (for purchases) or receives (for sales) on his TAS contracts.

Assume the manipulator trades Δ TAS contracts at t_0. In addition, the net order discretionary liquidity traders in the TAS market is e. Assume the price in the TAS market is $F(\Delta + e)$, and that the price in the exchange market at T is $P_T(x + u + z, \Delta + e)$, where x is the order submitted by the informed trader; u is the net order flow of the liquidity traders, and z is the order the manipulator submits to the exchange at T. The price is a function of the TAS market order flow because the market maker can use this order flow to make inferences about z.

As in Kumar-Seppi, I consider linear equilibria. At T, the market maker implements the following pricing rule:

$$P_T(x + u + z, \Delta + e) = \mu + \tilde{\delta}_{t_0} + \lambda_T[z + u + x - E(z + u + x|\Delta + e)] \quad (12)$$

Given this pricing rule, the informed trader chooses x to maximize his profit:

$$E_{u,z}\{x[\tilde{V} - P_T(x + u + z, \Delta + e)]|\tilde{V}, \Delta + e\} \quad (13)$$

Given Δ and the pricing rule, the manipulator chooses z to maximize:

$$E_{x,u}[\tilde{V}(z+\Delta)-P_T(x+u+z,\Delta+e)(z+\Delta)|e] = E_{x,u}[(\tilde{V}-P(x+u+z))(z+\Delta)|e] \tag{14}$$

This expression is derived as follows. After trading at T, has a position in the underlying of $z+\Delta$, which is worth $\tilde{V}(z+\Delta)$. The manipulator pays (receives) $P(x+u+z)$ for each TAS contract bought (sold), and pays (receives) the same price for each unit bought (sold) on the exchange at T. The manipulator conditions expectations on the discretionary trader TAS net order flow e because (a) he can infer this from total TAS order flow and his own TAS trade, and (b) as will be seen, the market maker conditions his price at T on TAS net order flow (because it provides information about the nature of time T order flow).

Given the pricing rule, the first order conditions for the manipulator's problem at T imply:

$$z = \frac{-\Delta + E(z|\Delta+e)}{2} \tag{15}$$

This is similar to the manipulator's choice in Kumar-Seppi, with the exception that the manipulator's trade at T is the opposite sign of the TAS trade at t_0, whereas in Kumar-Seppi the time T trade is in the same direction as the trade in the cash settled derivative. This difference arises because in Kumar-Seppi the buyer of the cash-settled derivative effectively sells it at T at the settlement price, and hence wants to drive up that price, whereas the buyer of a TAS contract buys at P_T, and hence wants to drive down that price. The crucial similarity is that in both cases, the manipulator takes into account the fact that the market maker incorprates his forecast of z based on order flow in the t_0 market into his pricing function even though it conveys no information about fundamentals.

Adapting the proof in Kumar-Seppi to take into account the difference between (15) and (A4) in their article produces the following equilibrium:[4]

$$z = \frac{(\kappa - 1)\Delta + \kappa e}{2} \tag{16}$$

$$\kappa = \frac{\sigma_W^2}{\sigma_W^2 + \rho S} \tag{17}$$

$$\lambda_T = \frac{\sigma}{2[S + \sigma^2(z|\Delta + e)]^{.5}} \tag{18}$$

$$\sigma^2(z|\Delta + e) = \frac{\kappa \rho S}{4} \tag{19}$$

$$E(z|\Delta + e) = \kappa(\Delta + e) \tag{20}$$

$$F(\Delta + e) = 0 \tag{21}$$

Thus, manipulation occurs in equilibrium. Furthermore, a comparison of (18) with (10) (with $N = 1$ and $\phi = 0$) shows that manipulation increases liquidity in the T exchange market because it adds noise to the order flow then. As a result, (a) informed traders gain, (b) T liquidity traders gain, (c) the manipulator gains, and (d) the costs of manipulation are borne by the t_0 discretionary traders who trade in the TAS market. The discretionary traders continue to use TAS nonetheless because price impact is still smaller in this market.

When employing this strategy, the manipulator breaks even on the transactions he undertakes during the auction that determines the settlement price at T. The units bought (sold) during this auction are matched against

[4]The equilibrium at t_0 is identical to the one derived in Section 2.

an equivalent number of TAS contracts sold (bought) prior to T: since the price received (paid) at the auction is the same as the paid (recieved) on the matched, the trades at T are perfectly hedged. The manipulator trades fewer units at the auction than the size of his TAS position, and he makes a profit (on average) on those units: *via* the TAS, he buys (sells) at price that on average is below (above) $\tilde{V}$.

This analysis brings home a basic lesson: order types or trading mechanisms that are designed to be unattractive to informed traders, and which thereby permit some of the uninformed to reduce their trading costs, also create opportunities for manipulation. By design, the price impact of these order types and trading mechanisms (such as TAS) is smaller than in the exchange market because they facilitate self-selection by trader type, which reduces adverse selection. But as noted by Jarrow (1992) such differential price impact can create manipulative opportunities. Manipulative trading is parasitic, and diminishes the value of the segmenting order type/trading mechanism to the clientele that it is intended to serve, but if not deterred, it is an inevitable consequence of the creation of this self-selection mechanism.

Exchanges recognize the potential for manipulation inherent in TAS. For example, CME Rules state:

> All market participants are reminded that any trading activity that is intended to disrupt orderly trading or to manipulate or attempt to manipulate a settlement price to benefit a TAS position will subject the member and/or the market participant to disciplinary action for any of a number of rule violations, including, but not limited to:

- price manipulation or attempted price manipulation
- wash trading
- conduct detrimental to the interest or welfare of the Exchange or conduct which tends to impair the dignity or good name of the Exchange
- engaging in conduct inconsistent with just and equitable principles of trade

Thus, the CME considers the trading strategy analyzed in this section to be manipulative. I now examine a particular alleged TAS manipulation that occurred on CME markets, specifically, the markets for gasoline, heating oil, and crude oil futures.

4 A Case Study: The Optiver Case

In July, 2008, the Commodity Futures Trading Commission ("CFTC") filed a complaint against the Dutch trading firm Optiver, alleging that it had manipulated New York Mercantile Exchange ("NYMEX") gasoline, heating oil, and crude oil futures contracts on various dates in March, 2007.[5] The CFTC identified the dates on which these manipulations occurred, and whether Optiver was a buyer or seller on each date. Therefore, this episode provides an informative case study of manipulation facilitated by TAS.

[5] United States Commodity Futures Trading Commission v. Optiver US *et al* 08 CIV 6560 (S.D.N.Y., 2008). Optiver settled these allegations by paying $14 million to the CFTC under a consent order entered on April 19, 2012. In addition to the monetary penalty, three Optiver traders were banned from trading commodities for periods ranging from four to eight years. Optiver settled class action litigation relating to the same conduct for $16.7 million in June, 2015. I was an expert for the Plaintiffs in that ligitation.

The CFTC alleged that Optiver engaged in a strategy to "bully the market." The trader devised a computer program that it called The Hammer that it used to implement this strategy.

The strategy was broadly similar to that modeled in Section 3 above. NYMEX offers a TAS order type on its Globex system. The buyer (seller) pays (receive) the official settlement price, and receives a long (short) position in the underlying futures contract. The settlement price equals to the volume weighted average price ("VWAP") during the last two minutes of trading (2:28:00 to 2:29:59) on the Globex system, plus or minus a differential (of up to ten ticks) negotiated between the buyer and seller at the initiation of the TAS trade. TAS trading occurs on the Globex trading system, and begins at 6:00PM on the day prior to the relevant settlement date, and ends at 2:30PM on the settlement date.

Early in the trading day, Optiver established a position in the TAS market (a buy of April 2007 gasoline futures TAS, say) by submitting limit orders. One of the features of The Hammer was an algorithm intended to optimize Optiver's TAS order submission strategy in order to increase the likelihood that these orders would be executed. Optiver would usually accumulate the bulk of its TAS position during the morning of the settlement day. If it obtained a sufficiently large TAS position, starting about 2:25 (three minutes prior to the beginning of the settlement window) Optiver began to trade heavily in the underlying (e.g., April 2007 gasoline futures) in the opposite direction as its TAS position through market orders. Optiver's volume during the 2:25-2:30 period was roughly equal to the size of the TAS position it had accumulated. Optiver traded so that approximately 25 percent of its volume was executed from 2:25-2:27:59, and the remainder during the settlement period. It intended that its volume during the 2:28-2:29:59 settlement window would represent such a large fraction of total volume during this period that its average execution price during this period would approximately equal the VWAP, and thus the settlement price.

In broad strokes, Optiver's strategy was similar to that derived in the model in Section 3. The firm established a large TAS position well prior to settlement, and traded in large volume in the auction market during the settlement period. Moreover, the volume traded during the settlement period was a fraction of the size of the TAS position Optiver accumulated prior to the auction.

The differences between the model of Section 3 and Optiver's strategy reflect simplifications to make the model more tractable and transparent.

First, unlike in the model, there was no single closing call auction on NYMEX to determine the settlement price: instead, the settlement price was the average of prices over a two-minute interval in a continuous auction. However, Optiver's strategy was tailored to try to replicate a feature of the model, namely the fact that futures trades during the settlement period were hedged by a portion of the TAS position.

Second, in the actual market it was of course not the case that the "true" value of the underlying was revealed after the closing auction, so that Optiver's expected profit was not based on the difference between this true value and the settlement price. Instead, Optiver's profit was determined by the difference between the price on its futures trades during the 2:25-2:27:59 pre-settlement period, and the final settlement price. In the case when Optiver was selling futures because it had bought TAS, it anticipated that although its sales during the pre-settlement period would drive down prices, its sales during the settlement window would drive down prices further, and cause the settlement price to be less than its average transaction price during the pre-settlement period.

A tedious yet straightforward modification of the model of Section 3 can accommodate this. Specifically, the model can be modified by adding a trading period $T^- < T$. In the most straightforward modification, the manipulator chooses to trade z^- units at T^-, and then is constrained to trade $-\Delta - z^-$ at T: that is, its position after the close of trading at T must be flat. Optiver in fact did this, and it can be rationalized as reflecting the manipulator's aversion to the risk of carrying an open position overnight (as is the case for most market making firms, like Optiver). This simplifies the analysis by reducing the number of choice variables (to Δ and z^-).

The main complication arising from this modification is that the manipulator's trading at T^- reveals information about the time-T trade $-\Delta - z^-$ that the market maker can use to adjust his pricing function. Although this complication makes the analysis more involved, the same basic result follows: manipulation is profitable.

An evaluation of price movements on the days that the CFTC alleges that Optiver demonstrates that its trading often did move prices by a substantial amount in the direction of its end-of-day trading. Figure 1 depicts the second-by-second average of the bid-ask midpoint in NYMEX May gasoline ("RBOB") futures on the nine days the CFTC alleged it "bullied" this market. The figure starts at 2:20:00, and continues through 2:45:59: trading continued on Globex even after the settlement window ended. Figure 2 depicts the bid-ask midpoints by second for May heating oil futures on the four alleged manipulation dates for that contract. Figure 3 presents the bid-ask midpoints by second for May crude oil futures on the five alleged manipulation dates.

These figures demonstrate several points. First, prices moved in the direction of Optiver's futures trading (and hence in the opposite direction of its TAS trading) on most of the alleged manipulation days. The formal model implies that since net order flows from other traders (informed traders and liquidity traders) (a) is sometimes opposite the manipulator's, and (b) can be larger in absolute value than the manipulator's, there will be days on which prices move in the opposite direction of the manipulator's trading. According to the model, the manipulator makes money on average across manipulative attempts, not on each attempt.

Second, the price movements are largely permanent.[6] This conforms with the model's prediction. The manipulator's trade in the settlement period auction cannot be distinguished from those of noise traders or informed traders. Moreover, since trades have a price impact only because of the presence of informed traders who cause adverse selection, and since information has permanent price impacts, the manipulative trades have a permanent price impact.

[6] In the microstructure literature, the "permanent" price impact of a trade is typically measured by looking at the change in the bid-ask midpoint from the time of a trade, to some period thereafter. Depending on the study, the ending price may be measured one, two, five, ten or even thirty minutes after the trade in question. The graphs measure the price change 15 minutes after the end of Optiver's trading, and hence the 2:45 price is a measure of the permanent price impact, according to the literature.

The fact that the manipulator's trades cannot be distinguished from informed trades has other implications that are borne out in the Optiver case. Specifically, information that is relevant for the price of one futures maturity (e.g., April crude oil) is relevant for the price of prices for other maturities. Thus, trades in one futures maturity impact the prices of other maturities. And indeed, in March, 2007, Optiver's trades in the front month futures contracts resulted in price movements similar to those observed in Figures 1-3, with the magnitude of these impacts generally declining with maturity (as predicted by the theory of storage–see Pirrong, 2011).

Moreover, in the case of crude oil and refined product futures, information that is relevant for one (e.g., crude oil) can be relevant for the others. And indeed this was the case in March, 2007.

These points can be demonstrated in a regression framework. I have obtained data on Optiver's trading in the front month (April) gasoline, heating oil, and crude oil futures during the 2:25-2:30 period for each day in March, 2007. Furthermore, based on transaction data and bid-ask data from NYMEX, using a tick test I have estimated net aggressive order flow for the 2:25-2:45 time interval: I subtracted Optiver's net orders from this net aggressive order flow to determine net order flow from traders other than Optiver. I then estimate regressions of the following form:

$$\Delta F_{imt} = \alpha_i + \sum_{j=1}^{3} \beta_{jm} Q_{jt} + \sum_{j=1}^{3} \gamma_{jm} I_{jt} + \epsilon_{it} \tag{22}$$

In this expression, ΔF_{it} is the change in maturity m futures price i from 2:25-2:45 on day t. Index i equals 1 for gasoline futures, 2 for heating oil futures, and 3 for crude oil futures. I examine four different maturities m:

April, May, June, and July The variable Q_{jt} is Optiver's net order in futures contract j: days on which Optiver did not trade are included in the sample, so the values of Q_{jt} are equal to zero for these days. The variable I_{jt} is the net non-Optiver aggressive order flow for the front month of futures contract j on day t. Note that the regression coefficients are indexed by m, meaning that price impacts of trading in April futures contracts can differ across maturities and contracts.

The results of this analysis are presented in Table 1. The number of coefficients is large, so the table presents only the coefficients on Optiver's trading. The results are readily summarized. First, Optiver's trades in the front month of a given contract (e.g., gasoline) are associated with permanent movements in the prices of at least the first four maturities of that contract in the direction of the trade, i.e., buys (sells) are associated with higher (lower prices). Second, Optiver's trades in front month gasoline futures are associated with permanent price movements in the direction of the trade in each of the first four maturity heating oil futures contracts, and Optiver's trades in front month heating oil futures are associated with permanent price movements in the direction of the trade in each of the first four maturity gasoline futures contracts. The coefficients on Optiver's trades in a given commodity (e.g., gasoline) in the regression using the price change of that commodity are all significant at the one percent level. Coefficients on Optiver's trades in gasoline (heating oil) in the heating oil (gasoline) regressions are all significant at the 10 percent level, and sometimes significant at the 1 percent level.[7]

[7]The signs of the coefficients on net aggressive order flow are typically of the right sign,

It should also be noted that Optiver's trades from 2:25-2:30 were almost perfectly negatively correlated with the size of its TAS position, and that the coefficients of regressions of its trades during this period against its TAS position are statistically indistinguishable from -1.00 for heating oil and crude oil. The coefficient for gasoline, -.88, is statistically different from -1.00 at the 1 percent level.

Overall, these results demonstrate that Optiver's actions in March, 2007, and the effects of those actions, are consistent with the theory of TAS manipulation. The firm accumulated large TAS positions, and then traded in the opposite direction of those positions shortly before and during the settlement period. On average, these trades caused the price of the future to move in the direction of the trade: indeed, as expected given the informational linkages between different maturities, and between related commodities, trades in one maturity of one commodity often affected other maturities and commodities. Moreover, these impacts were permanent.

5 Crossing Networks and Other Dark Pools

Dark pools are a ubiquitous part of the equity trading landscape. There are many different types of dark pools that employ different pricing models. Some are dark pools are essentially auction markets that allow the entry of priced orders, and may or may not include a market maker (usually the broker-dealer who operates the platform). Some dark pools are crossing networks ("CN") that cross offsetting orders at a price derived from another market. Even here there is diversity. Some cross contra orders continuously,

but not always, and are not statistically significant.

usually at the bid-ask midpoint derived from exchanges. Others cross orders periodically, and use a price from an exchange market from some time after the cross. For example, POSIT uses the average of the bid and offer from seven minutes after the cross. Instinet Last Daily Cross and Barclays Internal CN utilize the official closing price of the stock.[8]

This last type of dark pool bears the closest similarity to the models of Sections 2 and 3, and these models can be modified to demonstrate that this type of dark pool serves the same economic function as TAS, and is subject to the same kind of manipulation. The main modification is necessitated by the execution risk in CNs. (See Zhu, 2014 for an issue of execution risk in dark pools.) That is, if the number of buy orders in the cross is greater than (less than) the number of sell orders, only a fraction of the buy orders (sell orders) will be executed. Thus, in contrast to the TAS market, and the model of the TAS market presented above, there is no market maker to absorb order imbalances in the CN.

The fraction of an order submitted to a CN that is executed depends on the allocation rule. A common allocation rule is *pro rata.* For simplicity, I will assume that rule is employed: the analysis is easily modified to incorporate alternative rules, including those that give priority to volume.

Consider the modeling framework of Section 2, with a CN replacing the TAS market: for now, there are no manipulators. To take into account the possibility of incomplete order fills in the CN, I assume that all unmatched orders are submitted to the exchange for execution at the market price at $T^+ \geq T$: the orders not matched in the CN are therefore subject to adverse

[8] See Degryse, Van Achter, and Wuyts (2006) for a description of crossing networks.

selection.[9] Further, I assume that the orders are matched in the CN at the market clearing price at T.

In that framework, discretionary uninformed traders prefer to submit orders to the CN even though there is a probability that less than 100 percent of their order will be crossed. The intuition is straightforward: the discretionary uninformed incur liquidity (price impact) costs on all orders executed on the exchange at T. Conversely, they incur no liquidity/price impact cost on orders matched on the CN. Therefore, discretionary liquidity trader j wants to maximize the fraction of his order Y_j that is executed on the CN. He does so by submitting his entire order there. This is true for all j. As before, in the framework of Section 2, where private information is short-lived because it is revealed publicly prior to T or because of competition between informed traders, the informed shun the CN.

This implies that the CN serves the same function as TAS orders. It screens out informed traders, thereby reducing the exposure of a subset of uninformed traders to adverse selection.

And as in the TAS model, the existence of this pool of uninformed orders that incur no liquidity cost creates the asymmetric price response necessary for a manipulative strategy to work. To incorporate the lack of a market maker, and hence the possibility that some orders submitted to the CN will not be executed, the model must be modified slightly. Specifically, it is insufficient to specify a distribution of net discretionary uninformed order flow. Instead, it is necessary to specify a distribution for total uninformed buy

[9] In the manipulation model discussed below, the analysis is somewhat simpler if $T^+ > T$. In the model without manipulation, the primary conclusion is readily demonstrated if $T^+ = T$.

orders and a distribution for total uninformed sell orders. These distributions must have positive support, such as a rectified Gaussian, inverse Gaussian, or lognormal. Denote the distribution for the discretionary sell orders $g_s(y_s)$ and a distribution of discretionary buy orders $g_b(y_b)$.

Assume the manipulator submits an order $\Delta > 0$ to the CN: a symmetric analysis holds for a sell order. Under a *pro rata* rationing rule, the number of the manipulator's units that are matched is:

$$\alpha(y_s, y_b, \Delta)\Delta = \min[\Delta, \frac{\Delta y_s}{\Delta + y_b}] = \Delta \min[1, \frac{y_s}{\Delta + y_b}] \tag{23}$$

Therefore, the fraction of the order executed is:

$$\alpha(y_s, y_b, \Delta) = \min[1, \frac{y_s}{\Delta + y_b}] \tag{24}$$

Consider the manipulator's decision at T, when the crossing price is determined at the exchange auction. Assume that the market maker on the exchange chooses a linear pricing rule. This pricing rule reflects the dark nature of the crossing network: the exchange market maker cannot condition on the order flow in the CN. Therefore:

$$P_T^{CN}(x + u + z) = P_{t_0} + \lambda_T^{CN}(x + u + z) \tag{25}$$

where as before x is the informed trader's order; u is the net non-discretionary liquidity order flow; and z is the manipulator's order at T.[10] Note that due to the darkness of the CN, the market maker cannot condition his pricing rule on the order flow in the CN.

[10]This version of the model assumes that the unmatched discretionary liquidity trader order flow is executed on the exchange at $T^+ > T$. All of the conclusions hold when the liquidity orders that are not matched in the CN are executed at T.

The informed trader's decision is, as usual:

$$x = \frac{V - P_{t_0}}{2\lambda_T^{CN}} \tag{26}$$

Given Δ, the manipulator maximizes:

$$E_{V,u,y_b,y_s}[(V - P_T^{CN})(\alpha(y_b, y_s, \Delta)\Delta + z)] \tag{27}$$

This simplifies to:

$$-\lambda_T^{CN} E_{y_b,y_s}[z^2 + z\alpha(y_b, y_s, \Delta)\Delta] \tag{28}$$

The relevant first order condition is:

$$z = \frac{-\Delta E_{y_b,y_s}\alpha(y_b, y_s, \Delta)}{2} \tag{29}$$

This is similar to the manipulator's choice in the TAS setting, with the exceptions that (a) Δ is multiplied by the expected fraction of the manipulator's CN order that is matched, and (b) the manipulator takes into account the expected discretionary order flow that is executed on the exchange because of incomplete matching. This latter factor is irrelevant in the TAS case, because all discretionary order flow is executed in the TAS market. Furthermore, the expectation depends on Δ, because the manipulator's order affects the cross.

Note that z depends on the execution fraction α. Since this fraction depends on the rationing rule, the intensity of manipulative trading (given Δ) will as well. Rationing rules that give priority to volume result in a larger α than the *pro rata* rule, and therefore induce more aggressive manipulative trading on the exchange when the crossing price is determined.

It is possible to show that as in Section 3, and as in Kumar-Seppi, the manipulator's profit is strictly increasing in $|\Delta|$. Therefore, as before the manipulator takes the largest position that his wealth W allows, and randomizes between buying and selling in the CN. Given the distribution of W, and the densities $g_b(.)$ and $g_s(.)$, it is possible to estimate the distribution of z. This, in turn allows determination of the variance of order flow z, which I denote as $\sigma^2_{z,CN}$, implying that the variance of the order flow is $S_f = \sigma^2 + (1-\rho)S + \sigma^2_{z,CN}$. The linear pricing rule is determined by regressing V on f, which produces:

$$\lambda_T^{CN} = \frac{cov(V, \frac{V}{2})}{\lambda_T^{CN} S_f} \quad (30)$$

Thus:

$$\lambda_T^{CN} = \sqrt{\frac{\sigma^2}{2S_f}} \quad (31)$$

The variable z will not be Gaussian, and its distribution will depend on the densities of the discretionary buy and sell order flows (as well as on the specific allocation rule).[11] It is therefore not possible to solve for S_f in closed form, but it can be done numerically.

Thus, CNs utilizing derived pricing are vulnerable to manipulation.[12] The intuition is virtually identical to that for TAS. Crossing networks are

[11] z is not Gaussian even though Δ is Gaussian because in (29) Δ is multiplied by a non-linear function of Δ.

[12] Degryse, Van Achter, and Wuyts (2006) argue that CNs reduce the profitability of this form of manipulation by delaying fixing of the crossing price to 5-7 minutes after the cross. This would mitigate the ability of the manipulator to exploit temporary price impacts arising from market maker risk aversion or trade processing costs. However, the price impacts related to adverse selection are permanent, and even with delayed pricing, the manipulator could exploit these effects. The models studied herein assume risk neutral market makers, and no trade processing costs. Therefore, price impacts are due to adverse selection alone, and CNs that use delayed pricing are subject to the kind of manipulation modeled here.

unattractive to traders whose private information is short lived due to its impending announcement, or robust competition between the informed. Thus, by design, CNs cater to uninformed traders, and trades executed on CNs do not impact prices. This creates an asymmetry between the CN and the main exchange where the informed and some uninformed trades, and a manipulator can exploit this asymmetry by establishing a position in the dark CN, and trading on exchange in order to affect the price at which the CN order is executed.

The formal model assumes a time lag between the crossing of orders and the pricing of the crossed orders. Some crossing networks, including the largest (Liquidnet) cross continuously at the market midpoint. This raises the issue of whether these continuous crossing networks ("CCN") are subject to manipulation similar to that modeled above.

A manipulator could submit an order to the continuously crossing CN, and simultaneously trade on the contra side in the market used to set the crossing price, and profit in a similar way as a trader in TAS or a delayed pricing CN. However, for this strategy to work, there would have to be some other reason for informed traders to eschew trading in the CCN.

The mechanism for this separation posited above is the mandated pricing delay, which screens out those with short-lived private information. CCNs have no *mandated* pricing delay, but for operational reasons, including factors created by CCN rules, there can be significant delays between order submission and execution, low execution probabilities, and incomplete executions. Zhu (2014) shows that all of these factors make a trading venue relatively unattractive to informed traders, especially those with short-lived information. Thus, to be viable (by being attractive to discretionary liquidity traders because of lower adverse selection risk), CCNs must have a mechanism for screening out informed traders. But this is precisely what makes them vulnerable to manipulative strategies like those modeled here.

One last issue deserves comment. In the formal model the manipulator's trade on exchange is based on the *expected* execution ratio on the CN. The manipulator's profits cannot be reduced, and could be increased, if the on-exchange trade could be conditioned on information about y_s and y_b, which permits more precise estimates of the execution ratio $\alpha(y_b, y_s, \Delta)$. Thus, the manipulator would like to pierce the dark.

CCNs are likely to be more vulnerable to this than CNs using delayed pricing. A trader can send a small order to a CCN (sometimes called "pinging the pool"): if that order is executed, the trader learns that there is a potential cross of a larger order. Given this information, the trader can (a) determine whether a manipulation opportunity exists, and (b) condition his manipulative strategy (direction and size in both the CCN and the exchange market) on this information.

Of course informed traders could also utilize this "pinging" strategy to pick off liquidity-trader orders in the CCN. But if the CCN is too vulnerable to such predatory informed trading, it will not attract liquidity trader order flow.

This demonstrates yet again the *yin* and *yang* of the trading mechanisms studied here: the very features that make a trading mechanism attractive to liquidity traders because they facilitate the screening of informed traders make the mechanism a useful part of a manipulative trading strategy.

6 Regulation of Manipulation of Derived Pricing Mechanisms Under US Law

Manipulation is illegal in futures markets under the Commodity Exchange Act, and in securities markets under the Securities Exchange Act. The foregoing analysis sheds light on how the burden of proof could be met in a case involving manipulation of a derived pricing mechanism.

Commodities law in the United States imposes a four part test to prove manipulation: (a) the existence of an artificial price, (b) the ability to cause an artificial price, (c) the accused in fact caused the artificial price, and (d) the accused specifically intended to cause the artificial price. Under securities law, it suffices to show that the accused traded with the intent of distorting prices. Since the required elements of proof for securities manipulation are a subset of those for commodities, I will discuss the issue in terms of the latter.

First consider the existence of an artificial price. As with much manipulation law, there is some dispute about the definition of this concept, but in a nutshell it means a price that diverges from the competitive price that reflects supply and demand fundamentals: put differently, it is the price that would prevail but for a manipulative act.

In the context of a trade-based manipulation, the manipulative act is a trade or trades intended to move the price that is used to determine the payoff to a position in a contract or instrument held by the manipulator, e.g., the settlement price used to determine the price of a TAS. If it is determined that an alleged manipulator traded during and perhaps before the settlement window[13], an event study methodology can be used to quantify the impact of his trading. In the case of a single manipulative trade, the price movement during the period of his trading can be calculated, and compared to a measure of variability during a similar time period on non-manipulation days to determine whether this price movement is of the correct sign (e.g., the price rises during the period of a buy) and is statistically significant. In the case of multiple alleged manipulative episodes, the average trade-direction adjusted price movement during a window of time encompassing the manipulator's trading can be calculated, and the variability of this average can be determined from price movements on non-manipulation days.

The Optiver case indicates that there is potential manipulative spillover across different contract months, or across related commodities, or both. Thus, event studies can be undertaken for other contract months and related commodities.

Given data on the alleged manipulator's trades, it is also possible to estimate regressions like those presented in Section 5 to determine whether the alleged manipulator's trades were associated with price changes during the trading window, and whether the relation is statistically significant.

Since manipulative trades would not have taken place but for the manipulation, any price impact resulting from these trades injects artificiality into prices. In essence, market participants assign some probability to the possibility that an order that is manipulative, and hence is not informative,

[13] Trading before the settlement window could, as in Optiver, be the means by which the trader realized a profit. Furthermore, in the presence of adverse selection, trading before the window would have a persistent impact on prices, and therefore could impact the settlement price.

was submitted by an informed trader. Thus, the price impact of the accused manipulator's trades is a measure of price artificiality.

If a trader engages in manipulative trading on multiple occasions, the power of the event studies will be higher, and the evidence of price artificiality more robust.

Insofar as ability to cause is concerned, trades are likely to impact prices in any anonymous market. Temporary price impacts almost certainly exist (due to market maker risk aversion and trade processing costs): adverse selection costs are also likely to exist due to privately informed trading. The existence of price impacts in the market in question in comparable periods (e.g., around the settlement window) on non-manipulation days is sufficient to show the existence of such price effects, and hence the ability to cause. Conventional methods employed in microstructure research can be used to document these price impacts.

With respect to causation, the event study analysis described above is highly probative. Given well-known principles of market microstructure theory and empirics, a finding that an alleged manipulator's trades bear a statistically significant association with price moves (based on the event study, a regression of price movements on trades, or both) provides strong evidence that the trading caused the observed price movements.

Intent is often considered the most difficult element to prove in a manipulation case. In the Optiver case, the Respondents helpfully described their trading strategy and motivations to one another in various communications that were obtained by the CFTC. However, such damning communications may not always be available, and even if they exist, communications can often be ambiguous. The formal analysis provides some indicators that can be used to prove intent.

Specifically, in a derived pricing manipulation, the manipulator trades in opposite directions in the market for the derived pricing instrument (e.g., TAS or the CN) and in the underlying market that is used to establish the price for the derived pricing contract.[14] Moreover, the manipulator's trading in the underlying market is concentrated during the settlement period.

A trader (e.g., a discretionary liquidity trader) uses a derived pricing market to establish a position in the underlying in lieu of using the underlying market itself in order to minimize liquidity costs. Therefore, a discretionary liquidity trader would be less likely than a manipulator to trade in the opposite direction, especially during the settlement period. Indeed, to the extent that the trader cannot execute her entire desired quantity in the derived pricing mechanism (as often happens in a CN, for instance), she is likely to trade in the same direction in the derived pricing mechanism and the exchange market. In contrast, a manipulator *always* engages in a trade opposite to the derived pricing instrument trade.

A manipulator could argue that he is attempting to provide liquidity to the derived pricing market. However, manipulative trading strategies differ from typical market making. For example, a trader scalping a TAS contract could be expected to buy the TAS at the bid and sell it at the offer, rather than systematically establishing a TAS position, and then offsetting it in the underlying market: that is, a market maker would absorb temporary

[14]This is in contrast to manipulation of a cash-settled derivatives contract, *a la* Kumar-Seppi: in that case, the manipulator trades in the same direction in the derivative and underlying markets.

supply and demand imbalances in the TAS market, and then reverse his position when those imbalances reverse. Moreover, and crucially, the formal model demonstrates that the manipulator takes an offsetting position in the underlying market during the settlement period on only a fraction of his derived pricing instrument position. In the Optiver case, for example, the firm repeatedly and intentionally traded volumes representing only about 75 percent of its maximum TAS position during the settlement period. A market making strategy would more reasonably involve liquidating the entire remaining derived pricing position during the settlement period.

Furthermore, a market maker would prefer to trade by limit order to capture the half-spread in the underlying market, rather than by market order (and therefore pay the half-spread). For instance, liquidating a TAS position through market orders in the underlying market during the settlement period would be expected at best to be a scratch trade: the trader would receive the TAS market half-spread to initiate the position, and pay the half-spread in the underlying market to cover the TAS position by market order. A trading strategy (like Optiver's) which seldom (if ever) involved capturing the spread in the TAS market, and which involved crossing the spread in the underlying market, is highly atypical for a market maker. This is especially true given that the spread in the underlying market is likely to be greater due to the greater adverse selection risk in that market. It is, however, exactly what a manipulator would do (as the model demonstrates).

7 Summary and Conclusions

Derived pricing–conditioning the price of a transaction on a price from another market–can facilitate the screening of traders with private information, and thereby reduce the vulnerability of some uninformed liquidity traders to adverse selection. Mechanisms like Trade at Settlement contracts, or crossing networks with delayed pricing, can function as such a screening mechanism. This is especially true when private information is short-lived because it will be disclosed publicly soon after a trader obtains it, or when there is intense competition between informed traders.

These mechanisms fragment markets: only a fraction of transactions are executed on exchanges. This tends to reduce liquidity on exchanges, and increase the adverse selection costs incurred by the liquidity traders who must trade there because it is too costly for them to utilize the trading mechanism that screens out the informed. However, I demonstrate herein that it is possible that adverse selection costs are lower overall because the gains to those who can trade on the off-exchange market that screens exceed the losses to those who cannot. Thus, this fragmentation can be efficiency enhancing.

This gain does not come without cost: the very features that make off-exchange trading venues attractive to liquidity traders create opportunities for profitable manipulation. To be profitable, trade-based manipulative strategies require that purchases and sales have different impact on prices: the screening of informed traders *via* derived pricing makes price impact of trades on mechanisms that employ derived pricing lower than the price impact of trades in a market (like an exchange) that does not screen. This differential impact makes possible manipulative trading strategies.

Given the incentive of liquidity traders to avoid adverse selection, there is an incentive to create screening mechanisms that facilitate such avoidance, and indeed. This article demonstrates that Trade at Settle contracts and crossing networks, and derived pricing more generally, can do so. Thus, the theory predicts that derived pricing should be ubiquitous. But the theory predicts that this fragmentation of markets based on screening of informed traders will inevitably make trade-based manipulation possible and profitable. These phenomena are inextricably linked.