On Commodity Price Limits*

Rajkumar Janardanan
SummerHaven Investment Management

Xiao Qiao
SummerHaven Investment Management

K. Geert Rouwenhorst
Yale School of Management

February, 2017

Abstract

Price limits on commodity futures appear to restrict price discovery but not stop speculation. Consistent with delayed price discovery, returns continue in the same direction after limit days and do not reverse after one week, whereas returns are small after large price moves that do not hit limits. High volatility immediately before limits can be attributed to low inventories and increased stockout risk rather than excess speculation. Market participants trade similarly around limit days and non-limit days; speculators do not reduce their positions in response to limits. The options-implied futures price on the limit day predicts the next-day futures price, which suggests price discovery partially moves from the futures market to the options market.

* We are grateful to Geetesh Bhardwaj, Kurt Nelson, Ashraf Rizvi, and Kevin Sheehan for helpful comments. The views expressed are those of the individual authors and do not necessarily reflect official positions of SummerHaven Investment, LLC. All errors are our own.

1. Introduction

Price limits are a simple tool exchanges use to curb large movements in daily prices. Studying price limits is of broad interest, because they have deep implications for market microstructure, trade execution, and risk management. Their importance became most notable following the October 1987 stock market crash, which lead to the growth in interest on price limits and circuit breakers in the stock market. However, there is a relatively small body of work on commodity markets. We shed light on this topic by studying price limits in commodity markets over 25 years.

We analyze a large data set from 01/07/1991 to 05/23/2016 of nine commodities: soybean oil, corn, cotton, feeder cattle, live cattle, lean hogs, soybean, and wheat. We document prices tend to continue in the same direction as the limit day move. This effect is strong across all nine commodities. After limit up days, positive returns are 7-10% more likely compared to the unconditional probability. After limit down days, the probability of negative returns is 9-12% higher compared to the unconditional probability.

Our finding is consistent with the hypothesis that price discovery is hampered by price limits. For example, if the price limit for lean hogs is $3.00 per hundred weight (cwt), the current price is $30, and the true informationally efficient price should be $34, the price only adjusts by $3.00 to $33 on the limit day. On the following day when trading resumes, the price moves up by $1.00 to $34, what it would have been on the limit day had the limits not been in place.

Average returns around limit days are also consistent with the price discovery hypothesis. We adopt the event study methodology, taking limit hit days as "events" and tracing out average returns before and after. We find economically and statistically large return continuation on the day after the limit day. For limit up days, the average return on the following day is 40 to 62 basis points, and for limit down days, the average return on the following day is -38 to -63 basis points. When we compare our findings to the day after large price moves between 90% and 100% of the limit size, we do not find any return continuation. Average returns after 90% limit days are virtually zero for both positive and negative moves. These large average returns after limits are not driven by the large price moves, but rather are associated with limits themselves. Limits appear to restrict price discovery, which leads to return continuation.

If market participants can easily switch between commodity futures and other financial instruments, it is possible the incomplete price discovery on the futures markets can lead to price discovery in other markets. In particular, the options markets are closely related to the futures markets but do not have price limits.

We find evidence that price discovery partially moves to the options markets. Returning to our earlier example, because the options markets do not have price limits, on the limit day when the informationally efficient price change should be $4.00 for lean hogs, the options-implied price should be $34 rather than the limit price of $33. Next day when trading resumes on the futures market, the price is able to move to the options-implied $34. Based on this idea, we calculate options-implied futures prices on limit days from traded options, and test their forecasting power for actual futures prices on the following day.

Options-implied prices forecast the next day opening futures prices. To a lesser extent, they also forecast the next day closing futures prices. In forecasting regressions, a 1% increase in the return calculated from limit day close to options-implied prices is associated with a 0.76% increase in the close-to-open futures returns, and 0.15% increase in the close-to-close futures returns. Implied prices forecast next day opening prices better than closing prices, perhaps because next day closing prices also reflect new information over the course of the day.

If limits help to stop speculation, we would expect to see speculators reduce their long positions or add to their short positions after limit ups, and vice versa for limit downs. Using positions data of market participants from the Commodity Futures Trading Commission (CFTC), we find direct evidence against this hypothesis. Speculators add to their long positions after limit ups and hedgers add to their short positions. After limit downs, speculators add to their short positions and hedgers add to their long positions. This behavior is similar to that observed for the 90% moves. Limits do not appear to alter the trading behavior of market participants.

Price limits are more likely to occur when futures price volatility is elevated. High volatility could be symptomatic of high speculative activity, or could reflect changing fundamentals and increased risk level. Note these two effects may not be mutually exclusive, as speculation could change as a function of fundamentals. We provide evidence that the elevated volatility does not appear to be associated with higher speculation, given the trading behavior of speculators is similar for limits or non-limits. Instead, high volatility appears to be related to low commodity inventories.

We offer an alternative perspective on price limits through providing a fundamentals-based explanation of limit occurrences. Just before limit events, the futures curve around the contract which eventually hits the limit is in steep backwardation. Furthermore, inventory is relatively low and implied volatility is high. These three pieces of evidence paint a coherent picture of low inventory before limit events. Low inventories are associated with stockout risk (Deaton and Laroque, 1992), leads to high futures volatility, which in turn increases the probability that price limits are hit. To the extent price limits provide a bound for futures volatility, it is fundamental volatility rather than excessive speculation that price limits bound. Our inventory-based explanation of limits is consistent with price limits stopping price discovery and not curbing speculation.

We make several contributions to the price limits literature. Our main contribution is to provide a fundamentals-based explanation of limits. Literature on price limits has mostly focused on the trading aspect, relating limits to microstructure effects. We show low inventory is related to the occurrence of price limits, so they may be restricting fundamental volatility rather than excessive volatility from speculation.

Our second contribution is to provide direct evidence against the hypothesis that limits stop speculation. To address the question whether price limits curb speculation, the literature has mostly relied on indirect evidence using returns data (Ma et all, 1989, Kim and Rhee, 1997). Rather than providing indirect evidence using returns, our analysis using CFTC positions data provides direct evidence that price limits do not curb speculation.

Our third contribution is to treat limit days as "events" and trace out average returns around these events. In doing so, we establish a link between the limits literature and the event study literature. Previous work on limits mostly focused on analyzing the returns immediate after limit days without accounting for longer-term effects. For example, it is possible that the price continuation immediately after limit days could eventually reverse. We calculate average returns from three days before to five days after the limit day and do not find return reversal.

Compared to existing work, our paper is most closely related to Reiffen et al. (2006), who also investigate the price discovery hypothesis. Our paper provides a more comprehensive and deeper analysis of limits, and a fundamentals-based explanation of the occurrence of limits. Our paper is related to a broader strand of literature on price limits. Ma et al. (1989) analyze price limits on silver, corn, soybeans, and Treasury bond markets, but they do not compare options-implied prices with actual prices. Lehmann (1989) and Miller (1989) point out that Ma et al. (1989) build in a bias by excluding consecutive limit days. We get around this problem by considering both all limit days and only non-consecutive limit days. Evans and Mahoney (1996) study options-implied prices and limit days for cotton, whereas we do so for nine commodities to greatly improve statistical power and generalize their findings.

The paper proceeds as follows. Section 2 discusses the history of commodity price limits and takes a first look at the data. Section 3 examines the return behavior around limit days. Section 4 relates the futures and options markets around limit days. Section 5 provides an explanation of limits based on commodity fundamentals. Section 6 concludes.

2. Background and Data

2.1 Background and History

Price limits restrict prices from rising above or falling below pre-specified levels. For example, the price limit on soybeans is \$0.70 per bushel. If yesterday's closing price is \$8.00 per bushel, then today's trading must be between \$7.30 and \$8.70 per bushel. Otherwise, trading is temporarily halted until the market starts trading within those prices again. If today's closing price is outside of the price limit bounds, investors are not able to transact at this price and we experience a "limit day".

Price limits aren't the same as circuit breakers. Circuit breakers halts trading for a fixed period of time, whereas when price limits are reached trading is still allowed between the upper and lower limits. Price limits are also not price controls, or absolute price limits, which restricts the level of prices rather than changes in prices. Several reasons have been cited by exchanges and economists for the existence of price limits. To understand these arguments, it is useful to review the history of price limits.

Price limits can be traced back to at least the rice futures on the Dojima exchange in Japan in the early 18th century to curb excess volatility (Moser, 1990) during a period when prices were falling. The first price limit rule in the United States was imposed during World War I. The New York Cotton Exchange placed a price limit of three cents per pound on the cotton futures contract on August 27, 1917 (Howell, 1934). No justification was provided except that German submarine action had been affecting transatlantic cotton prices. Soon after in December, limits were imposed in hog markets. Grain markets had absolute price limits during this period rather than a limit on price changes.

In 1922, the U.S Supreme Court supported two arguments in favor of price limits: 1) Speculators cause price fluctuations, which cannot be reflected in spot markets quickly enough, leading to the farmers taking a lower price in the store. 2) Prices could be occasionally cornered or driven away from the true demand-supply price, and price limits would bring prices back to their fundamental levels. Both of these statements allege the victim is the producer and the offender is the speculator. This sentiment is still sometimes expressed today. Both arguments imply that if limits were effective, they would reduce speculation, dampen return volatility, or lead to price reversals.

In early 1925, wheat prices rallied to over two dollars per bushel, but by the spring prices had collapsed below \$1.20 (Hoffman, 1941). On May 25, 1925, Time Magazine published an article blaming the collapse of the prices on a bear raid by the famed speculator Jesse Livermore. In response to a senate investigation, the Secretary of Agriculture W.M. Jardine instructed the commodities exchanges to institute limits on daily price fluctuations (Grains Futures Administration. Fluctuations in Wheat Future: Letter from the Secretary of Agriculture Transmitting, In Response to Senate Resolution No. 222, of June 9, 1926). The Chicago Board of Trade formally adopted price limit rules in 1925.

Two additional arguments are occasionally made about the benefits of price limits. First, a trading halt may provide an opportunity for brokers to contact their clients and obtain instructions in times of turbulence (Telser, 1981). Second, Brennan (1986) argues "The existence of price limits…is explained by demonstrating that price limits may act as a partial substitute for margin requirements in ensuring contract performance." Neither of these arguments requires that prices or volatility to be correct, nor do they assert that price limits should somehow correct them.

Note that price limits are not unique to commodities markets. During the 1987 stock market crash, stock markets hit circuit breakers whereas the equity futures markets continued to trade. The subsequent investigation, popularly known as the Brady Commission, recommended and subsequently imposed price limits on futures markets. Demand for liquidity and order imbalances were cited as reasons for limits.

2.2 Theories and Hypotheses

We organize our analysis around two key questions. First, do price limits stop speculation? Stopping speculation is the motivation the U.S. Supreme Court upheld in 1922 in favor of price limits. If limits curb excessive speculation, some testable implications after limit events include the following:

1. Return reversal, because the reason the limits were hit is due to large volatility, not because of fundamental reasons (Ma et al., 1989)
2. Decreased volatility, as excessive speculation subsides (Ma et al., 1989)
3. Speculators to reduce their long position and/or increase their short position after limit up, and increase their long position and/or decrease their short position after limit down

We test these implications to shed light on the plausibility of the stopping speculation hypothesis.

Second, do price limits hinder price discovery? This is a main competing hypothesis for the argument that limits stop speculation offered after the 1987 stock market crash. If limits disrupt the price discovery process on the futures markets, then after limit events we would expect:

1. Return continuation, because the price could not fully adjust to its appropriate level in the presence of limits (Kim and Rhee, 1997)
2. No change in volatility before and after price limits
3. Speculators do not change their trading behavior relative to non-limit days

Related to these two key questions, we investigate if there is a fundamentals-based explanation for the occurrence of limits. The existence of such an explanation would support the hypothesis that price limits hinder price discovery.

Options markets do not have price limits. We hypothesize that the futures market participants switch to trading options on limit days. Price discovery may move to the options markets, such that the informationally efficient price is reflected there. This hypothesis implies the options-implied futures prices on the limit days would be good predictors of the next day futures prices when they are no longer constrained. We test the following:

1. Options-implied futures prices forecast next day opening futures prices
2. Options-implied futures prices forecast next day closing futures prices

We do not address Telser's (1981) point that a trading halt provides an opportunity for brokers to contact their clients or Brennan's (1986) point about price limits being a partial substitute for margins.

2.3 Data

Price limits are mostly in place for agricultural and livestock commodities; they are uncommon in energy and metals markets. We consider nine commodities with price limits on the futures contracts: soybean oil (BO), corn (C), cotton (CT), feeder cattle (FC), live cattle (LC), lean hogs (LH), soybean (S), soybean meal (SM), and soft red winter wheat (W).

Daily opening and closing prices for the above commodities are available from Bloomberg. Our sample is from 01/07/1991 to 05/23/2016. Lean hog futures first started trading in its February 1997 contract. Prior to that contract, live hog futures were traded until the December 1996 contract. Since producers and end users of both hog contracts are likely to be the same, we merge the two commodities into one and refer to it as lean hogs.

Price limits are from the CME Group. We identify closing limit days through comparing closing-to-closing price changes with the appropriate price limits. Price limits may change over time due to regulation changes. Limits could also be expandable such that following a limit hit, the limits get wider (e.g. 3 cents per bushel on the first limit day, the limit for the following day is 4.5 cents). We take these considerations into account, and search if the absolute price change exceeds the imposed price limits. We mark it as either a "limit up" or a "limit down" day depending on direction of change. Cotton price limits data were collected by Judith Ganes, a renowned consultant on softs markets (www.jganesconsluting.com).

Options data are from the Commodity Research Bureau (CRB). For each futures contract, we select the closest at-the-money put and call relative to the limit-day futures close price. We use the Citigroup 3-month Treasury Bill index from Bloomberg to construct the riskfree rate. Implied volatility data is collected directly from CRB who uses the Black and Scholes (1973) model to generate the values.

For market participants' trading positions, we obtain the Commitment of Traders (COT) reports published weekly by the U.S. Commodity Futures Trading Commission (CFTC) through Bloomberg. This report includes the long, short, and spread positions of commodity futures and options market participants. The participants are categorized as commercials, non-commercials, and non-reportables. Following the existing literature (e.g. Bessembinder, 1992; De Roon et al., 2000), we classify commercials to be hedgers and non-commercials to be speculators.

Inventory data is from United States Department of Agriculture (USDA) Livestock and Seed Division, USDA Economic Research Services, and New York Board of Trade. We use the monthly time series of inventory for each commodity. Inventory data are not always available until the end of our sample. Data for BO ends in July 2011, August 2014 for C, July 2014 for S, September 2011 for SM, and August 2014 for W.

2.4 Summary Statistics

Table 1 presents the summary statistics for the limit events. There may be multiple contracts for the same commodity hitting limits on the same day, and we count these occurrences separately. Livestocks (FC, LC, and LH) tend to hit limits more often compared to agricultural commodities (BO, C, S, SM, and W), but CT also experience many limit days. There are a total of 2063 limit ups and 2393 limit downs in our sample. About 10% of these limit hits are consecutive events: two limit ups or two limit downs in a row. Specifically, there were 209 consecutive limit ups and 247 limit downs. Not surprisingly, it is rarer to see more than two consecutive days. Four or more consecutive limit hits occur about 1% of the time.

Live cattle had a notable string of limit events: the October 2003 live cattle contract hit limit up for seven consecutive days from 10/06/2003 to 10/14/2003. In the summer of 2003, mad cow disease was discovered in Canada, and many countries proceeded to ban Canadian beef. The U.S. suppliers had to step in to fill the supply shortage. In October 2003, exports were up 39% in the U.S. Combined with holiday season demand, this strong export growth pushed the physical market prices from 160 to 190 cents per pound. This supply shock is likely responsible for the consecutive limit up days in October 2003.

Table 1: Price Limits Summary Statistics, 01/07/1991 – 05/23/2016.

	BO	C	CT	FC	LC	LH	S	SM	W	TOTAL
Number of Limits	340	364	877	1036	535	515	302	266	221	4456
Limit Up	146	165	480	414	236	259	107	112	144	2063
Limit Down	194	199	397	622	299	256	195	154	77	2393
2 Consecutive Limits	29	13	142	125	53	46	13	4	31	456
2 Cons Limit Ups	2	5	96	30	27	18	4	3	24	209
2 Cons Limit Downs	27	8	46	95	26	28	9	1	7	247
3 Consecutive Limits	1	2	45	40	13	9	0	0	5	115
3 Cons Limit Ups	0	1	37	2	6	3	0	0	5	54
3 Cons Limit Downs	1	1	8	38	7	6	0	0	0	61
4 Consecutive Limits	0	0	20	17	7	2	0	0	3	49
4 Cons Limit Ups	0	0	16	0	4	0	0	0	3	23
4 Cons Limit Downs	0	0	4	17	3	2	0	0	0	26
5 Consecutive Limits	0	0	12	8	4	0	0	0	1	25
5 Cons Limit Ups	0	0	10	0	3	0	0	0	1	14
5 Cons Limit Downs	0	0	2	8	1	0	0	0	0	11
6 Consecutive Limits	0	0	5	0	2	0	0	0	0	7
6 Cons Limit Ups	0	0	5	0	2	0	0	0	0	7
6 Cons Limit Downs	0	0	0	0	0	0	0	0	0	0

Note: This table presents the number of limit events for soybean oil (BO), corn (C), cotton (CT), feeder cattle (FC), live cattle (LC), lean hog (LH), soybean (S), soybean meal (SM), soft red winter wheat (W). Different contracts for the same commodities are counted separately. Consecutive limit days are doubly counted (e.g. three consecutive limit up days for one contract is also counted as two occurrences of two consecutive limit ups).

Cotton is another commodity which experienced a long string of limit events. Cotton futures incurred consecutive limit up days in April and May 1995. At that time, there was very little cotton deliverable against the New York exchanges. Open interest on the exchanges was also large. The circumstances forced short covering and led to explosive price increases.

It is also possible for different contracts of the same commodity to hit both limit up and limit down on the same day. There were two occurrences in our sample:

1. 10/09/2003: the October 2003 LC contract hit limit up, while the April 2004, June 2004, August 2004, and October 2004 contracts all hit limit down
2. 10/14/2003: the October 2003 LC contract hit limit up, while the February 2004, April 2004, June 2004, August 2004 hit limit down

The above events can be explained given market participants' expectations of the live cattle market following the mad cow scare. Following the spike in beef prices, futures participants aggressively bought all different contracts across maturities. As participants gathered more information, they recognized that the price spike in live cattle is likely to be a short-term event and that over time; both the export and domestic demand would adjust to prices. As a result, although the near-term contract hit limit up, the further out contracts were sold off and hit limit down. In short, mad cow disease caused excess volatility in the live cattle market which in turn led to unusual limit hits along the curve.

3. Market Behavior around Limit Days

We document commodity futures behavior before and after limit days and find prices tend to continue in the same direction as the limit event. If a contract experiences a limit up, the price next day is higher on average. If a contract experiences a limit down, the price next day is lower on average. This finding is consistent with delayed price discovery due to the restriction imposed by limits.

Exchanges with existing price limits argue that limits prevent day-to-day wild swings in stock prices. (e.g. Kim and Rhee, 1997) The Supreme Court also found in favor of this argument in 1922. Under this hypothesis, speculators sometimes drive up volatility with their frenzied trading. Price limits temporarily suspend trading so that market participants can reevaluate where the price should be. In effect, price limits serve to reduce speculation. If a limit up is reached one day, we expect prices to reverse and there to be a negative return when trading resumes. If several limit ups occur in a row, we expect eventually to observe negative returns when prices reverse. Similarly, if prices hit limit down, we expect to observe positive returns as prices reverse. Another implication for this hypothesis is that volatility of the futures contract should be lower after limits compared to before, which we test in a later section.

A competing hypothesis is that price limits may hinder price discovery (Phylaktis et al., 1999). If market participants' expectation for a new price level, perhaps due to new information released, is beyond the bounds implied by the price limits, the price can only adjust by the price limit before halting. In this scenario, there is no excess volatility from speculators. We would expect to see prices move in the same direction as the limit move as price discovery continues when market resumes function. If a limit up is reached one day, we expect prices to continue and there to be a positive return when trading resumes. We expect to see negative returns following limit down hits.

We find economically large return continuations around limit days. For this exercise, we adopt the event study methodology. The limit day is represented by t=0; the day after the limit day is t=1, the day prior to the limit day is t=-1, etc. Figure 1 reports the average returns and cumulative close-to-close returns around the limit day to illustrate the economic impact of price limits.

Figure 1: Average Returns and Cumulative Returns around Limit Day.

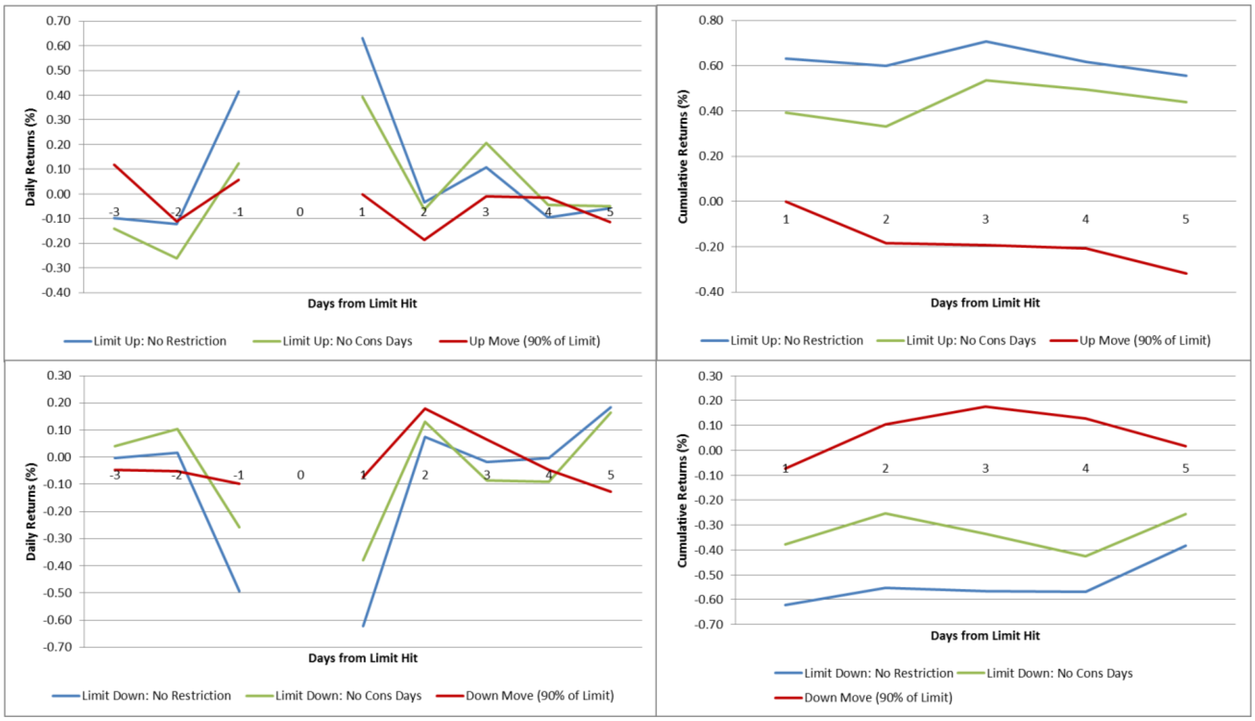

Note: This table reports the close-to-close returns and cumulative returns before and after a limit day event. t=0 marks the limit day. The blue line averages across all limit up (top) or limit down (bottom) events. The green line drops consecutive limit days and only uses single limit days as events for conservative estimates. As a comparison, we find days in which price movements were between 90% and 100% of the limit but do not hit the limit. The red line documents average returns before and after the 90% limit moves. We omit the limit day price moves as they are a magnitude larger.

The top left panel presents the result for average daily returns around limit up days. On the day after a limit up day, on average across the nine commodities the next day return is 63 basis points if we include all limit days and 39 basis points if we only include non-consecutive limit days. Omitting consecutive limit days provides a conservative estimate for the degree of price continuation, so that the 39 basis points can be viewed as a lower-bound value.

To understand the magnitude of the numbers in Table 2, we compare the behavior around limit events with those around large price moves that do not actually hit the limit. To do so, we first find days in which the price move was between 90% and 100% of the price limit. For example, if the limit is $3, we find all of the days with price moves between $2.70 and $3.00 and label them as "90% Move Up", and price moves between -$3.00 and -$2.70 as "90% Move Down". If next day returns are a function of large price moves rather than limits, then limit and non-limit days should exhibit similar next day returns. However, if there is something special about limit days, we expect different patterns for limit versus non-limit days. Kim and Rhee (1997) use a similar approach to evaluate limit hits on the Tokyo Stock Exchange.

It is possible that the price continuation is due to the time-series properties of these returns. If returns were autocorrelated, we may observe these patterns. To address this concern, we compare these values to returns around large price increases between 90% and 100% of the limits. 90% moves provide us with a baseline of how much of the price continuation is due to the time-series behavior of returns rather than because of limits. We find the average next day return is zero after 90% move up, which suggests limits are not merely large price movements but also alter the next-day price distribution.

Interestingly, on the day before the limit up day, the average returns are also positive. This may be a symptom of price pressure building up in the same direction as the limit day. On the top right panel, it is clear the large positive returns after the limit up day do not reverse over the next week. For the 90% price moves, the cumulative returns are closer to zero compared to those after limit days. The slight downward drift may be due to the mean reversion in commodities returns. (Bessembinder et al., 1995)

The bottom panels present the results for limit down days. These plots are almost mirror images of those for limit up days. The average returns on the day after a limit down day is -62 basis points for all limit days and -38 basis points if we only include non-consecutive limit days. Cumulative returns show slight reversal in the week after the limit down day, mostly driven by the positive average return five days after the limit day. For the 90% moves, both daily average returns and cumulative returns are close to zero. There is a small upward drift after 90% down moves, again perhaps due to mean reversion in returns. These results for limit down and large negative price moves support the idea that price continuation after limits is not driven by autocorrelation in returns. As a robustness check, Appendix A provides nonparametric estimates of price continuation using the conditional probability of positive and negative returns. Appendix B illustrates the impact of price limits on contracts along the curve that are not directly hit.

Figure 1 illustrates the economic magnitudes of the average returns before and after limit days, but does not report their statistical significance. Table 2 complements Figure 1 with pooled regressions of daily returns on limits or 90% move indicator variables:

$$r_{i,t+1} = a + b_{up} I_{Up} + b_{down} I_{Down} + \gamma FE_i + e_{i,t+1} \qquad (1)$$

Where $r_{i,t+1}$ is the one-day return on for contract i. I_{Up} and I_{Down} are indicator variables equal to one if the previous day was a limit up or limit down day, or a day with 90% move up or down. We include contract fixed effects, FE_i, to remove any latent contract-specific heterogeneities in average returns. Daily returns are calculated based on the same contract and do not include rolling from one contract to another.

Table 2: The Effect of Limit Hits on Next Day Returns.

	BO	C	CT	FC	LC	LH	S	SM	W	TOTAL
	All Limit Days									
b_{up}	0.58%	1.14%	0.72%	0.35%	0.57%	0.44%	0.73%	0.35%	1.06%	0.62%
	(1.10)	(3.77)	(3.19)	(2.82)	(4.07)	(2.53)	(1.75)	(0.65)	(1.95)	(5.95)
b_{down}	-0.72%	-0.64%	-1.04%	-0.30%	-0.41%	-0.88%	-0.63%	-0.48%	-1.37%	-0.63%
	(-1.61)	(-1.75)	(-4.75)	(-2.26)	(-1.99)	(-4.24)	(-1.43)	(-0.93)	(-1.98)	(-4.61)
Fixed Effects	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract
	No Consecutive Limit Days									
b_{up}	0.54%	1.15%	0.24%	0.25%	0.36%	0.20%	0.52%	0.16%	0.47%	0.38%
	(1.05)	(3.71)	(0.96)	(2.03)	(2.41)	(1.23)	(1.22)	(0.30)	(0.82)	(3.44)
b_{down}	-0.27%	-0.48%	-0.80%	-0.03%	-0.08%	-0.53%	-0.50%	-0.43%	-1.44%	-0.39%
	(-0.53)	(-1.28)	(-3.50)	(-0.22)	(-0.50)	(-3.23)	(-1.07)	(-0.83)	(-2.36)	(-2.75)
Fixed Effects	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract
	90% Limit Move									
b_{up}	-0.29%	-0.14%	-0.07%	0.20%	0.20%	-0.03%	-0.24%	0.15%	-0.25%	-0.01%
	(-0.91)	(-0.44)	(-0.32)	(2.44)	(2.23)	(-0.12)	(-0.58)	(0.39)	(-0.41)	(-0.13)
b_{down}	-0.35%	0.75%	-0.22%	-0.08%	0.14%	-0.11%	-0.28%	-0.77%	0.37%	-0.09%
	(-0.77)	(1.74)	(-1.15)	(-0.73)	(1.11)	(-0.46)	(-0.87)	(-1.31)	(0.94)	(-0.76)
Fixed Effects	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract

Note: We examine the effect of limit days on subsequent one-day returns. The regression specification is the following:

$$r_{i,t+1} = a + b_{up} I_{Up} + b_{down} I_{Down} + \gamma FE_i + e_{i,t+1}$$

Where $r_{i,t+1}$ is the close-to-close futures returns for contract i. I_{Up} (I_{Down}) is an indicator variable equal to one if the previous day is a limit up (limit down) day, or if the previous day was a large price move between 90% and 100% of the limit (-90% and -100% of the limit). FE_i are contract fixed effects to account for the latent heterogeneities among different contracts. T-statistics in parentheses are clustered by time.

Consistent with the results in Figure 1, average returns are positive after limit up days and negative after limit down days for every commodity. The rightmost column pools all nine commodities, such that the estimate coefficients are weighted-averages of the individual commodities coefficients. On the day immediately following limit up days, on average commodities gain 62 basis points. After limit down days, on average commodities lose 63 basis points. The small differences between these numbers and those in Figure 1 may be due to contract fixed effects in the regressions. In the middle panel, conservative estimates made by excluding consecutive limit days also show statistically and economically large price continuation as the top panel.

The bottom panel of Table 2 conducts the same analysis as the top two panels but replaces limit days with 90% limit move days. For these events, I_{Up} and I_{Down} correspond to 90% up and down moves. There does not appear to be a consistent pattern for the follow day's returns. After positive 90% moves, the next day average returns range from -29 basis points to 20 basis points, and after down moves, the next day average returns range from -77 basis points to 75 basis points. The heterogeneity in point estimates appear random, and they wash out when we pool all of the commodities together into one large regression in the rightmost column. After up or down 90% moves, the following day returns are nearly zero. Price continuation after limit days is not due to the large price move experienced on the limit day, but rather that the price move was not large enough.

4. Price Discovery in the Options Market

In the previous section, we showed prices tend to continue to move in the same direction as the limit hits. This finding suggests price discovery is hindered by price limits. If these prices limits did not exist, we would expect to see price discovery take place fully on the day of the limit hits.

Options contracts are traded in which the underlying is the commodity futures contracts. Unlike futures contracts, options are not subject to price limits. If the futures and options markets are not segmented – participants can easily access both markets – price discovery on limit days could move from the futures market to the options market (Reiffen et al., 2006). On limit days when the futures market stops trading, information cannot be impounded into prices in a timely manner. If market participants choose to express their views on the options market instead, the options-implied futures price should reflect current information.

Commodity options and futures prices are closely related through the put-call parity (Black, 1976). Consider a European call c_t and put p_t with strike price K, and their underlying futures contract F_t. By the Law of One Price, the following relationship holds:

$$c_t - p_t = (F_t - K)B(t,T) \qquad (2)$$

Where $B(t,T)$ is the present value of a zero-coupon bond that matures to \$1 at time T. $B(t,T)$ brings time T cash flows F_t and K to time t. We can rearrange (2) to isolate F_t:

$$F_t = K + \frac{c_t - p_t}{B(t,T)} \qquad (3)$$

To operationalize (3), we use the Citigroup 3-month Treasury Bill index to construct the interest rate used for $B(t,T)$, and in turn construct F_t. Specifically, on a limit hit day, we use the put and call associated with the closest strike price to the closing futures price and the prevailing interest rate to compute the implied futures price.

The above hypothesis predicts that the options-implied futures prices would forecast the unrestricted futures prices after limit hits. We test this hypothesis in the following sections by comparing the implied futures prices with the next day futures prices.

4.1 Forecasting Futures Prices

If price discovery moves from the futures market to the options market, options-implied futures prices are good forecasts for the next day futures prices when trading resumes. We can test this hypothesis by running forecasting regressions of the next day prices on the options-implied prices. Although our economic hypothesis calls for regressing actual prices on implied prices, one potential concern is that prices are integrated processes. To circumvent this problem, we examine the forecasting power of options-implied prices in return space. $F^o_{i,t+1}$ is the opening price of t+1 and $F^c_{i,t+1}$ is the closing price of t+1. $F^{implied}_{i,t}$ is the options-implied futures price at time t. Our specification is as follows:

$$r_{i,t+1}^{c,o} = a + br_{i,t}^{implied} + \eta_{i,t+1}^{c,o} \qquad (4)$$

$$r_{i,t+1}^{c,c} = a + br_{i,t}^{implied} + \eta_{i,t+1}^{c,c} \qquad (5)$$

Where $r_{i,t}^{implied} = \frac{F_{i,t}^{implied}}{F_{i,t}^{c}} - 1$ is the options-implied futures return calculated from the closing price on the limit day. $r_{i,t+1}^{c,o} = \frac{F_{i,t+1}^{o}}{F_{i,t}^{c}} - 1$ is the futures return from the limit day close to the next day open (close-to-open returns). $r_{i,t+1}^{c,c} = \frac{F_{i,t+1}^{c}}{F_{i,t}^{c}} - 1$ is the futures return from the limit day close to the next day close (close-to-close returns). . If $r_{i,t}^{implied}$ is a good prediction for the next day returns, we expect to see a=0, b=1, and R^2=100%. (Mincer and Zarnowitz, 1969) These values form our null hypothesis. Alternatively, b > 0 suggests implied prices are informative for next day actual prices

Table 3: Forecasts of Next Day Futures Returns.

	BO	C	CT	FC	LC	LH	S	SM	W	ALL
	LHV = Close-to-Open Returns									
a	-0.14%	0.00%	-0.05%	0.09%	0.04%	-0.25%	-0.18%	-0.05%	0.31%	-0.01%
	(-0.62)	(0.00)	(-0.55)	(1.02)	(0.50)	(-1.66)	(-0.78)	(-0.16)	(1.99)	(-0.16)
b	0.79	0.63	0.65	0.39	0.39	0.96	0.15	0.74	0.36	0.76
	(-1.07)	(-1.36)	(-4.54)	(-7.19)	(-6.88)	(-1.64)	(-5.99)	(-1.09)	(-3.50)	(-2.85)
F Test	0.27	0.23	0.00	0.00	0.00	0.07	0.00	0.28	0.00	0.00
N	114	198	390	318	179	221	126	93	178	1817
R^2	22.6%	40.2%	51.2%	38.8%	49.2%	94.1%	8.1%	18.2%	23.6%	68.6%
	LHV = Close-to-Close Returns									
a	-0.33%	0.01%	-0.05%	0.11%	0.18%	-0.08%	-0.02%	0.17%	-0.12%	0.02%
	(-0.56)	(0.04)	(-0.26)	(0.67)	(1.29)	(-0.39)	(-0.05)	(0.36)	(-0.29)	(0.15)
b	0.53	0.50	0.62	0.34	0.34	0.01	0.09	0.24	0.39	0.15
	(-1.25)	(-1.82)	(-4.28)	(-9.23)	(-7.17)	(-46.39)	(-9.61)	(-2.14)	(-3.26)	(-20.37)
F Test	0.23	0.30	0.00	0.00	0.00	0.00	0.00	0.02	0.00	0.00
N	118	200	403	327	179	221	131	93	184	1856
R^2	2.8%	13.7%	25.9%	15.2%	19.9%	0.0%	1.4%	0.8%	6.8%	3.7%

Note: We compute options-implied futures prices on limit days using put-call parity and forecast next day returns. We run the following regressions:

$$r_{i,t+1}^{c,o} = a + br_{i,t}^{implied} + \eta_{i,t+1}^{c,o}$$

$$r_{i,t+1}^{c,c} = a + br_{i,t}^{implied} + \eta_{i,t+1}^{c,c}$$

Where $r_{i,t+1}^{c,o}$ is the close-to-open return from the limit day closing price to the next day opening price for contract i. $r_{i,t+1}^{c,c}$ is the return from close of limit day to close of the following day. $r_{i,t}^{implied}$ is the return computed from the close of the limit day and the options-implied limit price on the limit day. "F Test" provides the p-value for the joint test that a=0 and b=1. T-statistics clustered by time are shown in parentheses. For a, t-stats are calculated from zero. For b, t-stats are calculated from one.

Table 3 presents the return forecasting results. For close-to-open returns, the intercepts a are close to zero. Although the intercept for wheat (W) is marginally significant, the pooled coefficient for the rightmost column for all commodities is zero. All of the estimates for b are statistically significant and positive. The t-statistics for b are calculated from one. The pooled estimate for b is 0.76, indicating that a 1% increase in the return calculated from the implied futures price and the limit closing price is associated with a 0.76% increase in the close-to-open return after the limit day. We cannot reject the null hypothesis that implied returns forecasts close-to-open returns at the 5% level for four of the nine commodities, but for the pooled regression we reject the null.

The bottom panel of Table 3 contains the return-forecasting regressions for close-to-close returns. Implied returns are unbiased forecasts for the close-to-close returns, as shown by the economically and statistically small estimates for a. The forecasting power is not as strong as for next day opening prices, however, as the estimates for b are generally smaller. Pooled estimates show a zero intercept (0.02%), and the forecasting coefficient b equals 0.15: on average, a 1% increase in the options-implied return implies a statistically significant 0.15% increase in the close-to-close return on the day after the limit day. Options-implied prices are informative about next day futures prices (open and close), but they are not perfect predictors for them.

Our results suggest that the options-implied prices are better forecasts of opening compared to closing prices. This finding may be due to the additional information released over the course of the trading day. Whereas the opening price is a good proxy for the price that would have prevailed on the limit day, had the limit not existed, the closing price is opening price plus the price adjustment to new information on the day.

All of the commodities have forecasting coefficient b of less than one, which suggests the options-implied price on average overshoots their targets. One possible explanation is that when the a futures market hits a price limit, trading stops and the volume moves over to the corresponding options market. For limit ups, the buying pressure combined with the sudden increase in demand for liquidity pushes the call prices too high relative to the puts, because options market makers have limited risk absorption capacity. The options-implied futures price will then be too high relative to the unconstrained futures price. Similarly, for limit downs, the selling pressure combined with the increase in liquidity demand pushes the put prices too high relative to the calls, and the implied price will be too low relative to the unconstrained price. In both cases, the options-implied futures price overshoots the price that would have prevailed in the absence of limits.

Suppose the forecasting coefficient b were zero for our tests, which would mean the options-implied prices do not contain any information about the next day prices. Table 3 illustrates that although options-implied futures prices may not perfectly forecast the next day futures prices, they do contain useful information. Positive predictive coefficients indicate options-implied prices tell us the direction as well as the approximate magnitude of the next day price moves. Our results are consistent with price discovery moving from the futures to the associated options markets on limit days, but limited risk capacity of market makers combined with unbalanced demand for puts and calls push options-implied futures prices too far in the same direction as limit events.

Our findings are consistent with the results in Reiffen et al. (2006). These authors compare options-implied futures prices and next day open prices for live cattle, lean hogs, and pork bellies. They use the difference between the options-implied prices and the limit day closing price to forecast the price change from limit day close to next day open, and find that a $1 increase in the predictor is associated with $0.50 to $1 increase the close-to-open price change. Their regression R^2 of 4% to 30% is also in line with our findings in Table 3.

What happens to price discovery for consecutive limit days? One possibility is that the options-implied price on the first limit day will forecast the eventual futures price after the string of consecutive limit days. Another possible outcome is for the options-implied price to only be weakly informative about the eventual futures price.

We repeat our analysis as in Table 3 for consecutive limit occurrences, and find that the first limit day options-implied prices provide weak forecasts for the futures price after all of the limit day have passed. The forecasts are unbiased, such that the intercept a as in Table 3 is close to zero, but b is 0.05 and statistically insignificant – a 1% increase in the options-implied returns forecasts a 0.05% increase in actual futures returns calculated from the first limit day to the futures price after the consecutive limits. Although we still get a positive association between options-implied returns and actual futures returns, the magnitude is too small to be statistically distinguished from noise.

Our finding for consecutive limit days could be due to the limited information market participants have on the first limit hit in a string of limit days. New information may be released after the initial limit day on which the options-implied futures price is taken, so this implied price becomes stale when we get to the day after the string of limits.

4.2 Trading at Settlement (TAS) and Price Discovery

CME and CBOT implemented TAS, or Trading at Settlement, for the commodities we study starting 06/08/2015. Throughout the trading day, market participants can choose to transact around the unknown settlement price. In particular, trades can be up to four ticks higher or lower than the settlement price. TAS covers all nine commodities we consider. An interesting feature of TAS potentially allows participants to trade after price limits are hit. If a particular live cattle hits limit up, the futures market ceases to trade unless the transaction is below the limit price. In contrast, TAS trades placed above the settlement price may still trade, such that transactions may occur at the settlement price plus up to four ticks.

At a glance, it may appear that TAS loosens the strict price limits and helps with price discovery. Indeed, theoretically it is possible for the unrestricted price to be within four ticks of the settlement price on a limit day. For example, if a live cattle contract hits limit up, and the unrestricted price is the limit price plus two ticks, the TAS market would facilitate price discovery by allowing transactions at that price. However, empirically the above thought experiment rarely occurs. Price limits are typically significantly larger than four ticks: for live cattle, the price limit is 12 ticks. TAS almost always trades at zero relative to the settlement price on limit days, because if a limit is hit, the unrestricted price is unlikely contained in the range of plus or minus four ticks. Furthermore, since the options market does not have price limits, those investors who want to transact at the unrestricted price can switch from futures to options and synthesize their desired positions. For all of these reasons, TAS does not facilitate price discovery on the futures market.

5. A Fundamental Explanation of Limit Events

Price continuation after limit days is consistent with the interpretation that price limits do not curb speculation and hinders price discovery. In this section we provide direct evidence that price limits do not stop speculation and market participants continue to trade the same way after limits. We also provide a fundamental explanation for the occurrence of limits that renders limits not particularly useful for curbing speculation. The basis, inventory, and implied volatility around limit days paint a coherent picture that price limits occur because inventory is low and prices are especially volatile, which leads to a higher probability of limit events.

5.1 Changes in Trader Positions

We use the Commitments of Traders (COT) reports from the U.S. Commodity Futures Trading Commission (CFTC) to uncover the trading behavior of market participants around limit days. We find that price limits do not alter the trading behavior of speculators and hedgers, thereby providing direct evidence that price limits do not prevent excessive speculation.

Every week, the CFTC publishes the COT report which includes the long and short positions of commodity futures and options market participants. The participants are categorized as commercials, non-commercials, and non-reportables. Following the existing literature (e.g. Bessembinder, 1992; De Roon et al., 2000), we classify commercials to be hedgers and non-commercials to be speculators.

Our measure of the trading behavior of speculators is the change in hedging pressure. Hedging pressure (HP) is defined as the net short position of hedgers divided by the open interest (Bessembinder, 1992). The change in hedging pressure, ΔHP, is the difference in hedging pressure between the nearest report after a limit day and the nearest report prior to the same limit day. Our measure of the trading behavior of speculators is the change in speculative position. We define speculator position (SP) as the net long position of speculators divided by the open interest. The change in speculator position (ΔSP) is the difference in speculator position before and after a limit day. While the levels of hedging pressure and speculator positions are informative about the positions of market participants, the change in these quantities provides information about how hedgers and speculators trade.

Table 4 contains the results for the trading measures. The top panel presents the changes in hedging pressure. On average, the change in hedging pressure from week to week is virtually zero. However, after limit up days, the change in hedging pressure is on average positive, whereas after limit down days, the change in hedging pressure is on average negative. Hedgers apparently tend to take more short positions after limit up and more long positions after limit down. Hedgers act as contrarian traders around limit days.

The bottom panel presents results for the change in speculative position. On average, the unconditional change in speculative position is zero. After limit up days, the change in speculative position is on average positive, indicating an increased net long position for speculators. Similarly, after limit down days speculators appear to be more net short on average. Speculators take the opposite side of hedgers and act as momentum traders around limit days. T-statistics presented are for the difference between the unconditional average positions changes and position changes around limit days. It is evident for both change in hedging pressure and change in speculative position, the differences are statistically large.

Table 4: Changes in Net Hedger and Speculator Positions around Limit Days.

	BO	C	CT	FC	LC	LH	S	SM	W	ALL
	ΔHP									
Avg	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Limit Up	4.6%	1.3%	1.2%	0.0%	-0.2%	1.1%	1.7%	2.0%	0.7%	1.0%
	(2.8)	(2.7)	(2.9)	(0.0)	-(0.7)	(3.3)	(3.0)	(2.0)	(2.0)	(5.7)
Limit Up - 90% Up	-0.7%	0.1%	-1.3%	-0.4%	0.1%	0.0%	0.4%	-0.8%	-1.4%	-0.6%
	-(0.5)	(0.2)	-(3.2)	-(1.7)	(0.4)	(0.1)	(0.8)	-(0.8)	-(4.2)	-(3.7)
Limit Down	-2.4%	-1.7%	-1.3%	-1.2%	-0.7%	-1.8%	-2.2%	-3.5%	-1.1%	-1.6%
	-(8.9)	-(6.6)	-(5.7)	-(6.8)	-(3.5)	-(5.0)	-(6.3)	-(6.9)	-(2.4)	-(16.1)
Limit Down - 90% Down	-0.3%	-0.3%	0.7%	0.2%	0.1%	-1.0%	0.5%	1.1%	-0.3%	0.2%
	-(1.2)	-(1.3)	(2.8)	(1.0)	(0.7)	-(2.5)	(1.3)	(2.2)	-(0.6)	(2.1)
	ΔSP									
Avg	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	0.0%	0.0%	0.0%
Limit Up	3.0%	0.4%	0.9%	0.0%	-0.6%	1.5%	1.1%	1.3%	0.7%	0.7%
	(2.5)	(1.3)	(2.7)	(0.2)	-(1.5)	(4.3)	(2.7)	(2.1)	(1.8)	(4.5)
Limit Up - 90% Up	-0.9%	-0.4%	-1.2%	-0.4%	-0.1%	0.4%	0.5%	-0.7%	-1.1%	-0.5%
	-(0.8)	-(1.2)	-(3.3)	-(1.1)	-(0.2)	(1.3)	(1.3)	-(1.2)	-(2.9)	-(3.4)
Limit Down	-1.7%	-1.1%	-1.0%	-1.0%	-0.6%	-1.4%	-1.4%	-2.2%	-1.2%	-1.2%
	-(7.2)	-(5.6)	-(5.3)	-(4.1)	-(2.8)	-(3.9)	-(6.1)	-(5.9)	-(4.5)	-(12.1)
Limit Down - 90% Down	-0.4%	-0.2%	0.4%	0.5%	0.2%	-0.8%	0.2%	0.9%	-0.4%	0.2%
	-(1.7)	-(0.8)	(2.2)	(2.3)	(1.1)	-(2.1)	(0.9)	(2.3)	-(1.6)	(2.1)

Note: We calculate measures of trading behavior around limit days. HP is the net short position of hedgers divided by open interest. ΔHP is the change in hedging pressure before and after limit days. SP is the net long position of speculators divided by open interest. ΔSP is the change in speculator position before and after limit days. For each panel, "Avg" reports the average values of these measures, whereas "Limit Up" and "Limit Down" condition on limit events. "Limit Up – 90% Up" and "Limit Down – 90% Down" are the differences in the measures around limit days versus around 90% moves. White (1980) t-statistics are reported in parentheses. T-statistics are calculated for the difference between limit and 90% days for those indicated, and otherwise are calculated for the difference between limit days and the unconditional average.

To construct a baseline of the trading behavior of market participants around large price moves, we investigate the change in hedging pressure and speculative position around 90% limit move days. These are days that prices moved between 90% to 100% of the price limit without actually invoking the limit. After 90% up days, hedgers are more net short whereas speculators are more net long. After 90% down days, hedgers are more net long whereas speculators are more net short. In line with the limit day results, hedgers appear to be contrarian traders and speculators appear to be momentum traders.

If the trading behavior of hedgers and speculators only depend on past returns independent of limits, then hedgers should take larger contrarian positions after limit ups compared to 90% moves, because the price changes are larger for limit ups. By similar logic, speculators should take larger long positions after limit ups than 90% moves, hedgers should take larger long positions after limit downs than 90% moves, and speculators should take larger short positions after limit down than 90% moves. We test this hypothesis in Table 4.

Our results in Table 4 do not support the above hypothesis that hedgers and speculators base their trading behavior on past returns independent of limits. We find hedgers and speculators both trade less around limit days compared to 90% days. Across commodities, ΔHP is on average 0.5% lower after limit ups compared to 90% ups, and is on average 0.3% higher after limit downs compared to 90% downs – rather than increasing their contrarian positions as in the above hypothesis, hedgers reduce their contrarian positions around limit days compared to 90% days, rather than increasing them. We observe similar behavior for speculators: ΔSP is on average 0.3% lower after limit ups compared to 90% ups, and on average 0.3% higher after limit downs compared to 90% downs.

The trading behavior of hedgers and speculators we document is consistent with the finding in Kang et al. (2016) that hedgers act as contrarian traders and speculators act as momentum traders. However, unlike Kang et al. (2016), hedgers do not appear to be earning a liquidity premium as returns on the first day after limit days continue in the same direction, and do not reverse after the first day. Figure 1 shows that for the 90% move days, we do observe return reversals – after large moves up, returns are negative for the next week, after large moves down, returns are positive for the next week – consistent with the liquidity provision interpretation in Kang et al. (2016). This difference between the 90% and limit days is consistent with prices moving to their appropriate levels after limits. When price limits are not hit, hedgers provide liquidity to speculators and earn a liquidity premium. When price limits are hit, price discovery is restricted. Hedgers and speculators agree where the true price should be for the next day when trading resumes, such that there is no additional drift afterwards.

We also examine the gross positions of hedgers and speculators around limit days as well as 90% days. We calculate the gross position of hedgers, Gross HP, as the sum of all long and short positions of hedgers divided by open interest. Δ Gross HP is then the difference in Gross HP before and after limit days. We also calculate the change in gross speculator position, Δ Gross SP, similarly. These two variables inform us of the overall trading activity regardless of direction for speculators and hedgers. We present the results in Table 5.

For hedgers, the difference between their gross position before limit days and after limit up days is generally similar. Corn, cotton, and soybeans show some decreases in activity, whereas soybean oil, feeder cattle, and lean hogs show increases. Across all commodities, the difference around limit up days is economically and statically small. This behavior around limit up days is similar to around 90% up days. On average, gross position for hedgers are similar around 90% up days. For limit down days, gross hedger activity is somewhat higher by 0.5%, and this change is similar to what is also observed for 90% down days. Combining this result with those from Table 4 for Δ HP, it is evident hedgers hold increased long positions relative to short positions after limit down, possibly to take the other side of speculators and trade contrarian to past price moves.

Turning our attention to gross speculator positions, we see after both limit up and limit down days the gross speculator position is apparently reduced. After limit up days, the average speculator position is 0.8% lower; after limit down days, the average speculator position is 1.1%

lower. We also see position reductions around 90% move days, but the limit day reductions are 0.3% larger across commodities. At a glance, it appears limits are curbing speculation. However, closer inspection reveals that does not appear to be the case.

If price limit truly curbed speculation, we expect to see speculators decrease their long positions after limit up days and perhaps increase their short positions. In fact, Table 4 and 6 show that following limit up days, the net speculator position is increased whereas the gross speculator position is decreased. Our finding indicates speculators exit their short positions and increase their long positions, just the opposite of what curbing speculation would suggest.

Under the assumption that limits do curb speculation, we would also expect to see speculators decrease their short positions or increase their long positions after limit down days. Tables 4 and 5 show that the net speculator position on average decreased 1.2% after limit down days, whereas the gross speculator position on average decreased 1.1%. Speculators decreased their long positions and left their short positions unchanged, again contradicting the predictions of stopping speculation.

Table 5: Change in Gross Hedger and Speculator Positions around Limit Days.

	BO	C	CT	FC	LC	LH	S	SM	W	ALL
	Δ Gross HP									
Avg	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Limit Up	0.3%	-0.8%	-1.2%	0.9%	0.1%	0.5%	-1.7%	0.0%	0.0%	-0.1%
	(0.6)	-(2.6)	-(3.4)	(4.2)	(0.1)	(1.1)	-(6.7)	(0.0)	(0.1)	-(0.9)
Limit Up - 90% Up	0.7%	-0.1%	-1.4%	0.2%	-0.1%	0.0%	-0.2%	0.8%	-0.2%	-0.1%
	(1.3)	-(0.5)	-(4.1)	(1.2)	-(0.3)	(0.0)	-(0.8)	(0.9)	-(0.5)	-(0.8)
Limit Down	0.2%	0.3%	0.1%	0.0%	1.1%	0.6%	0.7%	2.4%	-0.1%	0.5%
	(0.3)	(0.7)	(0.4)	(0.1)	(5.0)	(1.2)	(1.4)	(4.2)	-(0.2)	(3.7)
Limit Down - 90% Down	-0.7%	-0.2%	0.5%	-0.3%	0.6%	0.6%	-0.2%	1.1%	0.1%	0.1%
	-(1.3)	-(0.6)	(1.5)	-(1.5)	(2.7)	(1.3)	-(0.5)	(2.0)	(0.2)	(1.0)
	Δ Gross SP									
Avg	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Limit Up	-1.1%	-0.5%	-0.8%	-0.8%	-0.9%	-1.2%	-0.2%	0.0%	-0.7%	-0.8%
	-(1.8)	-(2.0)	-(3.2)	-(2.9)	-(2.8)	-(3.2)	-(0.7)	(0.0)	-(3.1)	-(6.7)
Limit Up - 90% Up	-0.6%	0.0%	-0.4%	-0.3%	-0.2%	-0.1%	-0.2%	-0.1%	0.4%	-0.3%
	-(1.0)	(0.0)	-(1.6)	-(0.9)	-(0.7)	-(0.2)	-(0.8)	-(0.3)	(1.7)	-(2.4)
Limit Down	-1.5%	-0.7%	-0.8%	-1.5%	-1.1%	-0.7%	-1.0%	-1.4%	-1.2%	-1.1%
	-(7.2)	-(3.7)	-(3.6)	-(5.4)	-(4.9)	-(1.9)	-(4.9)	-(3.7)	-(3.2)	-(10.4)
Limit Down - 90% Down	-0.7%	-0.2%	-0.7%	-0.1%	0.1%	-0.9%	0.2%	0.3%	-1.0%	-0.3%
	-(3.2)	-(0.9)	-(3.2)	-(0.2)	(0.4)	-(2.3)	(1.2)	(0.7)	-(2.8)	-(2.8)

Note: We calculate measures of trading behavior around limit days. Gross HP is the gross position of hedgers obtained by summing up long and short positions. Δ Gross HP is the change in Gross HP before and after limit days. Gross SP is the sum of long and short positions of speculators. Δ Gross SP is the difference between Gross SP before and after limit days. For each panel, "Avg" reports the average values of these measures, whereas "Limit Up" and "Limit Down" condition on limit events. "Limit Up – 90% Up" and "Limit Down – 90% Down" are the differences in the measures around limit days versus around 90% moves. White (1980) t-statistics are reported in parentheses.

We verify our results are not driven by the scaling by open interest or netting out long and short positions. We repeat the above analysis using short hedger positions and long speculator positions, without scaling by open interest. Hedgers build up more short positions after limit up and fewer short positions after limit down, and speculators have more long positions after limit up and fewer long positions after limit down. In the interest of space, we do not include these results here.

The results in this section show that limits do not alter the trading behavior of hedgers or speculators. Both groups operate similarly around limits and large 90% price moves. Furthermore, considering both net and gross positions of hedgers and speculators reveals that price limits do not appear to be stopping speculators – they increase their long positions after limit ups and do not change their short positions after limit downs. These findings provide convincing evidence that limits do not appear to be effective in curbing speculation.

5.2 Migration of Trading to Options Markets

When price limits are hit in the futures market, options continue to trade. It is possible that the usual trading activity in the futures market moves over to the associated options market in which trading is unrestricted. Specifically, hedgers continue to hedge in the options market by synthesizing the appropriate positions, and speculators continue to take the opposite side of hedgers. If market participants switch from futures to options, we expect to see the open interest and the gross positions in the options market increase relative to the futures market after limit days. We test this hypothesis below.

The CFTC COT report distinguishes between futures and options market participants. The CFTC publishes two reports – one for futures only and another that includes delta-adjusted futures positions along with futures. Taking the difference between these reports gives us the futures equivalent options positions of market participants. We form ratios of trading positions in the combined futures and options markets to trading positions in only the futures market. Trading positions are measured with 1) Open interest, 2) Hedger gross positions formed by summing up the long and short positions of hedgers, and 3) Speculator gross positions formed by summing up the long and short positions of speculators. We calculate the change in the ratios computed using these three variables and report the results in Table 6.

In the top panel of Table 6, we present the difference in the ratio of futures and options open interest to the open interest in the futures market. Coefficients have been multiplied by 100 for the ease of exposition. On average from one week to the next, there is no change in this ratio for any of the commodities. However, after limit up or limit down days, this ratio apparently increases, consistent with the hypothesis that trading activity migrates to the options market when the futures market hits price limits. There is also a large difference between the changes in ratios for limit days and 90% move days, which suggests it is limit days rather than large price moves that prompt swift switches from trading futures to trading options. Our finding is consistent with Evan and Mahoney (1996, 1997).

The middle panel shows the change in the ratio of hedger gross positions. The results for hedgers are similar to those for overall open interest. Hedgers' gross position in the options market is larger relative to their positions in the futures market after limit days. There are also significance differences for changes in position ratios for limit days and 90% move days.

The lower panel presents the difference between position ratios of speculators. For speculators, trading does not appear to move from the futures market to the options market, as indicated by the negative difference between the positions ratio before and after limit days. It does not appear to be the case that speculators effectively translate their futures positions into options positions.

Table 6: Change in the Ratio of Positions around Limit Days.

Coeff*100	BO	C	CT	FC	LC	LH	S	SM	W	ALL
	Δ Ratio of Open Interest									
Avg	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Limit Up	0.83	2.22	5.31	-1.23	-0.09	0.94	4.56	1.08	2.30	1.67
	(2.0)	(4.3)	(5.7)	(-3.5)	(-0.2)	(3.9)	(5.4)	(2.7)	(3.9)	(6.3)
Limit Up - 90% Up	0.00	-0.09	3.89	-0.48	0.03	0.16	0.61	-0.07	0.75	0.81
	(0.0)	(-0.2)	(4.1)	(-1.4)	(0.1)	(0.7)	(0.7)	(-0.2)	(1.2)	(3.1)
Limit Down	1.75	0.34	0.69	3.53	1.73	0.60	0.69	0.86	1.60	1.67
	(5.8)	(0.5)	(0.7)	(6.7)	(3.1)	(1.9)	(1.0)	(3.4)	(3.8)	(7.0)
Limit Down - 90% Down	1.87	-0.62	0.36	0.73	-0.08	-0.75	0.28	0.58	1.20	0.52
	(6.2)	(-1.3)	(0.4)	(1.4)	(-0.1)	(-2.5)	(0.4)	(2.3)	(2.8)	(2.2)
	Δ Ratio of Hedger Gross Positions									
Avg	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Limit Up	0.47	0.94	2.75	-0.78	-0.29	0.19	2.03	0.81	0.83	0.72
	(1.4)	(3.8)	(4.8)	(-3.0)	(-0.9)	(1.0)	(3.6)	(2.7)	(2.3)	(4.3)
Limit Up - 90% Up	-0.46	-0.07	1.92	-0.03	0.04	0.13	0.18	0.45	0.58	0.44
	(-1.4)	(-0.3)	(3.3)	(-0.1)	(0.1)	(0.8)	(0.3)	(1.5)	(1.6)	(2.6)
Limit Down	1.65	0.04	0.09	1.48	1.38	0.57	0.66	0.49	0.85	0.90
	(6.3)	(-0.1)	(0.2)	(5.3)	(3.1)	(3.7)	(1.4)	(2.1)	(2.4)	(5.9)
Limit Down - 90% Down	1.64	-0.27	-0.14	0.21	0.41	-0.24	0.58	0.23	0.51	0.32
	(6.3)	(-1.4)	(-0.2)	(0.8)	(0.9)	(-1.6)	(1.2)	(1.0)	(1.4)	(2.1)
	Δ Ratio of Speculator Gross Positions									
Avg	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Limit Up	-2.10	-0.30	-1.35	-0.68	-0.15	-1.18	-1.03	-0.30	-0.03	-0.82
	(-2.1)	(-1.2)	(-2.8)	(-2.1)	(-0.4)	(-3.8)	(-1.5)	(-0.4)	(0.0)	(-5.0)
Limit Up - 90% Up	-1.97	-0.21	-1.25	-0.06	-0.12	-0.18	-1.30	-0.49	0.42	-0.58
	(-2.0)	(-0.7)	(-2.6)	(-0.2)	(-0.4)	(-0.6)	(-2.0)	(-0.7)	(0.8)	(-3.6)
Limit Down	-3.09	-0.35	-0.48	-0.21	-0.35	0.51	-1.80	-1.84	-3.27	-0.75
	(-3.9)	(-1.1)	(-1.0)	(-0.8)	(-1.2)	(1.6)	(-2.6)	(-2.2)	(-4.2)	(-4.7)
Limit Down - 90% Down	-1.79	-0.25	-0.83	0.25	0.55	0.17	0.41	-1.22	-3.23	-0.30
	(-2.2)	(-0.7)	(-1.8)	(1.0)	(1.9)	(0.5)	(0.6)	(-1.4)	(-4.1)	(-1.9)

Note: We calculate ratios of open interest, gross position of hedgers, and gross positions of speculators. Each ratio is formed by taking the CFTC combined futures and options value and dividing the futures only value. "Δ Ratio of Open Interest" is the difference between the ratio in open interest before and after limit days. "Δ Ratio of Hedger Gross Positions" is the change in the ratio of the sum of hedgers' long and short positions in futures and options to the same sum in futures. "Δ Ratio of Speculator Gross Positions" is the ratio of the sum of speculators' long and short positions. Coefficients are multiplied by 100. White (1980) t-statistics are reported in parentheses.

5.3 Basis

We turn our attention to a fundamental driver of commodity prices, inventories. Basis – the percent difference between one contract and the next nearest contract – reflects the prevailing inventory and provides fundamental information about commodities (Gorton et al, 2013). We compare the before and after basis around limit events to shed light on how fundamentals may be related to limits.

The basis increases after limit up and decreases after limit down. This pattern holds for every commodity. The difference in basis is reflective of the price continuation after the limit day for the limit contract but not for other contracts along the curve. For limit up events, the limit contract tends to be even higher the following day, such that basis is more positive on the day after limit up days. The reverse is true for limit down days: The limit contract tends to be lower the following day, so basis is more negative in the day after.

For most commodities, the day before both limit up and limit down days exhibit elevated basis. Corn (C), cotton (CT), live cattle (LC), soybean (S), soybean meal (SM), wheat (W), and to a lesser extent feeder cattle (FC) all exhibit higher basis before limit days. Basis is informative about the inventory levels of the underlying commodity, so Table 7 suggests that inventory is low before extreme price moves in either direction. Perhaps this is not surprising, as low current inventory is associated with increased stockout risk and higher futures price volatility (Deaton and Laroque, 1992). Since price limits are not a function of volatility, high conditional volatility leads to more limit occurrences.

Table 7: Basis Before and After Limit Days.

	BO	C	CT	FC	LC	LH	S	SM	W
Avg	-1.95%	-2.73%	-2.95%	1.53%	0.23%	0.78%	4.09%	5.83%	-2.71%
	Limit Up - All								
Day Before	-2.68%	11.30%	15.13%	1.04%	5.06%	-0.08%	7.65%	11.97%	14.76%
Day After	-1.83%	13.69%	18.02%	2.45%	9.77%	3.38%	8.87%	14.57%	17.07%
	Limit Down - All								
Day Before	-2.21%	9.60%	12.41%	2.98%	12.71%	0.34%	10.93%	18.13%	20.32%
Day After	-2.54%	7.10%	9.70%	2.06%	9.72%	-5.83%	8.66%	15.79%	17.25%
	Limit Up - No Consecutive								
Day Before	-2.68%	7.35%	8.28%	1.21%	3.41%	2.11%	7.25%	11.79%	14.58%
Day After	-1.86%	9.66%	10.62%	2.79%	7.21%	5.27%	8.33%	14.49%	16.36%
	Limit Down - No Consecutive								
Day Before	-1.97%	4.72%	8.30%	1.63%	11.66%	2.27%	11.63%	16.47%	21.34%
Day After	-2.28%	2.73%	5.89%	0.73%	9.25%	-3.63%	9.13%	14.45%	18.12%

Note: Basis is calculated as the percent difference between one contract and the next nearest contract, and scaled to be in annual terms.

We find direct evidence of low inventory just before limit days using data from USDA and New York Board of Trade. Table 8 compares the nearest inventory numbers before limit days to the average inventory in the past year. For example, for soybean oil (BO), prior to limit up days, its inventory is -6.41% lower compared to the average inventory level in the past year. We observe decreased inventory for six of the nine commodities before limit up, and five before limit down. It does appear the high basis prior to limits is associated with low inventory. The evidence in Tables 7 and 8 combined support the arguments in Deaton and Laroque (1992). After a period of low inventory and high basis, futures prices are especially volatile, which leads to limit events. Feeder cattle, live cattle, lean hogs, and soybean meal do not show drastically decreased inventories. These are also the commodities that show relatively smaller changes in basis immediately before limit events in Table 7.

Table 8: Inventory Relative to Historical Average before Limit Days.

	BO	C	CT	FC	LC	LH	S	SM	W
Limit Up - All	-6.41%	-16.69%	-38.44%	2.72%	-2.00%	2.51%	-23.34%	0.65%	-4.88%
Limit Up - No Cons	-6.49%	-15.09%	-29.24%	3.35%	-0.30%	1.80%	-26.16%	0.31%	-7.59%
Limit Down - All	-5.38%	-13.24%	-13.70%	4.54%	0.86%	5.96%	-17.72%	6.16%	-6.21%
Limit Down - No Cons	-6.06%	-13.07%	-9.11%	2.25%	1.81%	3.90%	-20.71%	6.16%	-5.93%

Note: We compute the percent difference between the most recent month's inventory before limit days and the average inventory in the past year. A positive number indicates the most recent month's inventory is higher compared to the average inventory. Inventory data for BO ends in July 2011, August 2014 for C, July 2014 for S, September 2011 for SM, and August 2014 for W.

5.4 Implied Volatility

If price limits serve the purpose of curbing speculation, not only would we expect to see subsequent prices reverse, but also subsequent volatility fall (Ma et al., 1989). The difficulty is that using prices subject to limits to calculate volatility leads to mismeasured volatility (Miller, 1989). If we keep all of the limit days, the realized volatility computed using these data is automatically biased downwards. If we throw out the consecutive limit days and only focus on the one-day limit events, we are by construction left with smaller price moves and bias downwards the volatility estimates even further. Miller (1989) asserted this problem to be a "possibly insuperable obstacle".

We circumvent this selection bias issue by using implied volatility instead of realized volatility. Because there are no price limits in the options markets, implied volatility does not suffer from the same selection bias as in the futures markets. We examine implied volatility before and after limit days using data from the Commodity Research Bureau. Table 9 contains the results.

Compared to unconditional averages, implied volatilities are generally higher immediately before limit days. This finding is consistent with increased basis and lower inventory. Deaton and Laroque (1992) demonstrate that low inventory is associated with stockout risk, which is associated with elevated volatility. Limits are more likely to be hit when volatility is high, as they do not adjust to changing volatility. Table 9 provides additional evidence to Tables 7 and 8 and paints a coherent picture that low inventory, high volatility, and high basis sets the stage for limit events.

Increased volatility before limit events may also be consistent with excessive speculation. However, this interpretation does not appear to be plausible when we compare implied volatility before and after limit events. Whereas the hypothesis that price limits are able to curb speculation implies lower volatility after limit events, implied volatilities before and after limit days are remarkably similar. There is no material difference before and after limit days. This finding holds even if we limit our sample to only non-consecutive limit days. The level of implied volatility is lower compared to including all limit days, because throwing out consecutive limit days mechanically reduces the next-day volatility, but implied volatility before and after limit days are still similar. We repeat the above analysis for average weekly volatility around limit days and find similar results. We compare the average of the implied volatilities in the five days prior and five day after limit days. The weekly average volatility before and after are both somewhat lower compared to the day before and after, because volatility is likely the highest right around limit days. Similar patterns emerge: The average implied volatilities in the week before limit events are similar to those in the week after.

Table 9: Implied Volatility around Limit Days.

	BO	C	CT	FC	LC	LH	S	SM	W
Avg	23.0	25.7	22.1	17.2	14.0	28.0	23.2	23.5	28.9
	Limit Up - All								
Day Before	29.4	31.8	34.1	17.5	20.4	48.0	28.1	26.4	39.6
Day After	30.6	33.3	35.8	16.9	20.8	45.9	32.1	31.0	42.9
	Limit Down - All								
Day Before	32.8	34.8	31.9	18.4	20.1	41.2	34.8	34.1	39.5
Day After	32.6	34.3	31.8	19.9	21.4	43.2	32.7	34.7	37.5
	Limit Up - No Consecutive								
Day Before	29.6	31.6	31.1	17.2	20.2	37.3	27.3	25.9	37.9
Day After	30.8	33.0	32.2	16.4	20.1	34.6	30.3	30.0	40.9
	Limit Down - No Consecutive								
Day Before	30.6	34.6	30.6	16.3	17.5	37.9	34.4	33.4	39.8
Day After	31.3	34.1	30.7	17.5	19.1	38.8	32.2	34.1	37.3
	BO	C	CT	FC	LC	LH	S	SM	W
	Limit Up - All								
Week Before	28.7	31.4	32.0	17.9	20.0	41.9	27.2	25.8	39.0
Week After	30.7	32.9	33.2	16.5	20.3	42.5	31.0	30.8	42.0
	Limit Down - All								
Week Before	31.4	34.5	31.2	17.1	18.6	36.0	33.8	33.4	38.9
Week After	31.7	33.9	31.4	18.6	19.8	39.6	32.2	34.1	37.7
	Limit Up - No Consecutive								
Week Before	28.9	31.3	30.8	17.8	19.9	34.7	26.5	25.4	38.4
Week After	30.8	32.8	31.8	16.2	20.0	35.8	29.9	30.3	41.3
	Limit Down - No Consecutive								
Week Before	30.4	34.4	30.6	16.5	17.9	36.0	33.7	33.2	39.2
Week After	31.4	33.7	30.7	17.9	19.1	39.1	31.6	34.0	37.9

Note: We compute implied volatility for available contracts. “Day Before” and “Day After” are averages across the day before or after a limit day. “Week Before” and “Week After” are simple averages in the five days before and after a limit day.

The finding that volatility behaves similarly before and after limits suggests price limits are not stopping speculative behavior. If limits are effective in restricting speculation, we would expect to see lower volatility following limit hit days. We can reject this hypothesis. On the other hand, our volatility finding is consistent with limits disrupting price discovery, which would imply similar volatility measured before and after limits. A more plausible interpretation of Table 9 is that volatility is high before limit days because inventory is low, which leads to limit events.

6. Conclusion

In this paper, we examine a large data set of commodities limits of more than 25 years. We revisit key competing hypotheses about price limits and propose a new fundamentals-based explanation.

We first examine the return behavior around limit days, and find returns tend to continue in the same direction as the limit events. The fraction of daily positive returns is 10% higher after limit up days, and 10% lower after limit down days, compared to unconditional return distributions. Average returns are high after limit up days, and are low after limit down days. Regression analysis shows that returns on the day after limit up days are statistically higher compared to the unconditional returns by 40 to 60 basis points, and returns on the day after limit down days are statistically lower by 40 to 60 basis points. These effects are economically large and do not reverse even after a week. Price continuation suggests limits hinder price discovery rather than curb speculation.

The economic effect on prices around limit days is not due to large price moves. We compare days around limits with those around price moves exceeding 90% of the limit but do not actually hit the limit. We find the returns in the day after 90% move up or move down exhibit similar behavior to other trading days. In particular, 53% of returns are positive or zero, and average returns before or after the 90% move days are economically and statistically small.

We provide direct evidence that limits do not restrict speculation through examining the traders' positions from the CFTC COT report. After limit up events, speculators increase their long positions relative to short positions, just the opposite of what restricting speculation would imply. More generally, speculators and hedgers trade in a similar fashion around limit events compared to large price moves, indicating limits are not effective in altering the trading behavior of market participants.

We propose a fundamentals explanation of price limits based on the theory of storage. Immediately before price limits occur, inventory is relatively low. This state is reflected in the elevated futures basis and implied volatility, as well as direct evidence of lower inventory. Increased futures volatility due to stockout risk makes it more likely that futures prices will hit limits. Rather than limiting excess volatility from speculation, price limits are hit when time-varying volatility is high due to low inventory.

On limit days, price discovery which usually happens on the futures markets stops and may move to the options markets. We test this hypothesis by comparing the options-implied futures prices with the next day opening or closing prices. We find that options-implied prices contain useful information about next day opening prices, and moderately useful information for next day closing prices. These results support our hypothesis that options prices are closer to the informationally efficient prices. Options-implied prices more closely approximate the next day opening prices rather than closing prices because closing prices also include any new information that comes out during the day.

Limits restrict the maximum daily price movement. Our paper shows the restriction is often hit when fundamentals drive up futures volatility instead of speculators' frenzies trading leading to excess volatility. In this sense, our findings suggest it's not clear the extent to which the current limits regulation helps commodity futures market function more efficiently. However, limits may have other benefits not discussed in our paper. Limited price movement protects investors from large price swings and margin calls, thereby alleviating market liquidity spirals. Perhaps a structural model which includes both the effects of restricted price discovery and the cost of liquidity spirals can get closer to answering this question.

References

Bessembinder, H., Coughenour, J. F., Seguin, P. J., and Smoller, M. M. "Mean Reversion in Equilibrium Asset Prices: Evidence from the Futures Term Structure." *Journal of Finance* 50 (1995): 361-375.

Black, F. and Scholes, M. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (1973): 637-654.

Black, F. "The Pricing of Commodity Contracts." *Journal of Financial Economics* 3 (1976): 167-179.

Brennan, M. J. "A Theory of Price Limits in Futures Markets." *Journal of Financial Economics* 16 (1986): 213-233.

Deaton, A. and Laroque, G. "On the Behaviour of Commodity Prices." *Review of Economic Studies* 59 (1992): 1-23.

Evans, J. and Mahoney, J. M. "The Effect of Daily Price Limits on Cotton Futures and Options Trading." Working Paper. New York: Federal Reserve Bank of New York (1996).

Evans, J. and Mahoney, J. M. "The Effects of Price Limits on Trading Volume: A Study of the Cotton Futures Market." *Current Issues in Economics and Finance* 3 (1997).

Gorton, G. B., Hayashi, F., and Rouwenhorst, K. G. "The Fundamentals of Commodity Futures Returns." *Review of Finance* 17 (2013): 35-105.

Howell, L. D. "Fluctuations in Prices of Cotton Futures Contracts." (1934).

Kang, W., Rouwenhorst, K. G., and Tang, K. "A Tale of Two Premiums: The Role of Hedgers and Speculators in Commodity Futures Markets." Working Paper. *Available at SSRN 2449315* (2016).

Kim, K. A. and Rhee, S. "Price Limit Performance: Evidence from the Tokyo Stock Exchange." *Journal of Finance* 52 (1997): 885-901.

Kim, Y. H and Yang, J. J. "What Makes Circuit Breakers Attractive to Financial Markets? A Survey." *Financial Markets, Institutions & Instruments* 13 (2004): 109-146.

Lehmann, B. N. "Commentary: Volatility, Price Resolution, and the Effectiveness of Price Limits." *Journal of Financial Services Research* 3 (1989): 205-209.

Ma, C. K., Rao, R. P., and Sears, R. S. "Volatility, Price Resolution, and the Effectiveness of Price Limits." *Journal of Financial Services Research 3* (1989): 165-199.

Miller, M. H. "Commentary: Volatility, Price Resolution, and the Effectiveness of Price Limits." *Journal of Financial Services Research* 3 (1989): 201-203.

Mincer, J. A. and Zarnowitz, V. "The Evaluation of Economic Forecasts." *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance* NBER (1969): 3-46.

Moser, J. T. “Circuit Breakers.” *Economic Perspectives* 14 (1990): 2-13.

Phylaktis, K., Kavusannos, M., and Manalis, G. “Price Limit and Stock Market Volatility in the Athens Stock Exchange.” *European Financial Management* 5 (1999): 69-84.

Reiffen, D., Buyuksahin, B., and Haigh, M. S “Do Price Limits Limit Price Discovery in the Presence of Options?” *Available at SSRN 928013* (2006).

Reiffen, D. and Buyuhsain, B. “The Puzzle of Privately-Imposed Price Limits: Are the Limits Imposed by Financial Exchanges Effective?” *Aestimatio: The IEB International Journal of Finance 1* (2010): 110-143.

Telser, L. G. “Margins and Futures Contracts.” *Journal of Futures Markets* 1 (1981): 225-253.

White, H. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica* 4 (1980): 817-838.

Appendix A: Fraction of Positive and Negative Returns After Limit Days

We compare the unconditional fractions of positive and negative returns with the fractions after limit days. If there is price continuation, we expect to see a greater fraction of positive returns after limit ups and a greater fraction of negative returns after limit downs. If there is price reversal, we expect to see a lower fraction of positive returns after limit ups and a lower fraction of negative returns after limit downs.

Table A.1 presents the results. Looking across all contracts and all dates, these nine commodities have a 53% chance of being positive or zero (47% positive and 6% zero). In the day immediately following a limit up day, the probability of a positive or zero return is 10% higher at 63.3%. In comparison, following a limit down day the probability of a positive or zero return is more than 10% lower at 40.2%. These numbers illustrate that prices tend to continue in the same direction as the limit events, and are consistent with the hypothesis that price limits restrict the extent prices can move in a certain direction, and hinder price discovery.

Price continuation does not depend on multiple consecutive limit hit days. In the panels titled "Limit Up – No Consecutive Hits" and "Limit Down – No Consecutive Hits", we count the number of positive and negative returns after limit days excluding any consecutive limit days. If limit days are persistent, dropping consecutive limit days provides a conservative estimate for the fraction of price continuation. We observe the same patterns as the case with all limit days: Positive or zero returns are more likely after limit ups and negative returns are more likely after limit downs. As expected, the difference with the unconditional distribution is smaller: After non-consecutive limit ups, the fraction of positive returns is 6.8% greater compared to the unconditional counterpart. After non-consecutive limit downs, the fraction of positive returns is 9.3% lower. These are large differences that suggest our results are independent of inclusion or exclusion of consecutive limit days.

Among others, Reiffen et al. (2006) also find prices tend to continue in the same direction as the limit day. However, they find the continuation fraction is 75% to 80% rather than 55-63% that we find. This difference could be due to the different commodities and sample periods. In particular, Reiffen et al. (2006) include 488 limit hits for pork belly futures, a commodity we do not consider since it is no longer traded.

Table A.1: Fraction of Positive and Negative Returns Unconditionally and after Limit Days, 01/07/1991 – 05/23/2016.

	BO	C	CT	FC	LC	LH	S	SM	W	TOTAL
	Unconditional Returns									
% $r_{t+1} \geq 0$	50.4%	52.5%	51.9%	54.6%	54.8%	54.3%	52.9%	52.1%	52.3%	52.8%
% $r_{t+1} < 0$	49.6%	47.5%	48.1%	45.4%	45.2%	45.7%	47.1%	47.9%	47.7%	47.2%
	Limit Up - All									
% $r_{t+1} \geq 0$	54.8%	69.7%	65.4%	64.4%	66.9%	59.1%	65.4%	48.2%	66.0%	63.3%
% $r_{t+1} < 0$	45.2%	30.3%	34.6%	35.6%	33.1%	40.9%	34.6%	51.8%	34.0%	36.7%
	Limit Down - All									
% $r_{t+1} \geq 0$	33.0%	44.7%	30.7%	42.4%	44.8%	35.9%	49.2%	46.1%	39.0%	40.2%
% $r_{t+1} < 0$	67.0%	55.3%	69.3%	57.6%	55.2%	64.1%	50.8%	53.9%	61.0%	59.8%
	Limit Up - No Consecutive Hits									
% $r_{t+1} \geq 0$	54.2%	69.9%	58.4%	62.3%	62.2%	55.8%	62.6%	45.3%	61.4%	59.6%
% $r_{t+1} < 0$	45.8%	30.1%	41.6%	37.7%	37.8%	44.2%	37.4%	54.7%	38.6%	40.4%
	Limit Down - No Consecutive Limits									
% $r_{t+1} \geq 0$	34.0%	46.2%	32.3%	48.1%	49.2%	40.8%	49.2%	46.7%	39.7%	43.5%
% $r_{t+1} < 0$	66.0%	53.8%	67.7%	51.9%	50.8%	59.2%	50.8%	53.3%	60.3%	56.5%
	90% Move Up									
% $r_{t+1} \geq 0$	55.2%	46.8%	49.0%	59.7%	61.0%	51.7%	40.3%	43.8%	51.5%	52.5%
% $r_{t+1} < 0$	44.8%	53.2%	51.0%	40.3%	39.0%	48.3%	59.7%	56.3%	48.5%	47.5%
	90% Move Down									
% $r_{t+1} \geq 0$	54.0%	62.3%	46.8%	50.4%	60.6%	53.5%	50.6%	50.9%	76.3%	54.0%
% $r_{t+1} < 0$	46.0%	37.7%	53.2%	49.6%	39.4%	46.5%	49.4%	49.1%	23.7%	46.0%

Note: This table shows the close-to-close one-day futures returns. "Unconditional Returns" displays the fraction of positive and negative one-day returns, regardless of limits, for all contracts and days. "Limit Up – All" and "Limit Down – All" show the fractions for positive and negative returns on the day following the limit hit. "Limit Up – No Consecutive Hits" and "Limit Down – No Consecutive Limits" repeat the calculation excluding all consecutive limit days. "90% Move Up" and "90% Move Down" restrict the sample to days in which the price moves were between 90% and 100% of the limit, without actually hitting the limit.

After 90% move up or down days, the fractions of the next day positive or negative returns are similar the unconditional proportions. For 90% limit ups, 52.5% of the next day returns are positive or zero and 47.5% are negative. The unconditional numbers are 52.8% and 47.2%. Similarly, for 90% move downs, these fractions are 54.0% and 46.0%. It does not appear that conditioning on large moves substantially changes the proportions of positive and negative returns. Rather it is the restriction that comes from hitting the limits that alter the return distribution.

Appendix B: Effect of Limits on Contracts along the Curve

We consider the effect of limit hits on contracts of the same commodity but different maturities. Contracts on the same commodity naturally tend to behave similarly. If a limit event reflects a common shock to all contracts of a commodity, we may see similar average returns for both contracts that hit limits and those that do not. For example, on 04/28/2016, the June 2016 live cattle contract hit limit down. Does this limit event affect next day returns on the April 2016 or August 2016 live cattle contracts, or is the effect concentrated on the June 2016 contract? We use a variation of Equation (1) to test this hypothesis:

$$r_{i,t+1} = a + b_{up} I^*_{Up} + b_{down} I^*_{Down} + \gamma FE_i + e_{i,t+1} \qquad (B.1)$$

Where the only difference compared to (1) is that we replaced I_{Up} and I_{Down} with I^*_{Up} and I^*_{Down}, which are indicator variables equal to one if on the previous day any contract along the curve aside from *i* hit the limit. In the above example, on 04/29/2016 I^*_{Down} would be equal to one for the April 2016 or August 2016 live cattle contracts, but equal to zero for the June 2016 contract. We present the results in Table B.1.

Table B.1: The Effect of Limits on Contracts along the Curve.

	BO	C	CT	FC	LC	LH	S	SM	W	TOTAL
	All Limit Days									
b_{up}	-0.44%	0.03%	0.10%	0.08%	-0.06%	-0.04%	0.45%	-0.04%	0.15%	0.02%
	(-1.32)	(0.14)	(0.77)	(0.95)	(-0.84)	(-0.43)	(1.30)	(-0.09)	(0.38)	(0.39)
b_{down}	-0.07%	0.30%	-0.03%	0.01%	0.01%	0.07%	-0.07%	-0.54%	-0.22%	-0.01%
	(-0.15)	(1.24)	(-0.21)	(0.10)	(0.07)	(0.61)	(-0.23)	(-1.04)	(-0.53)	(-0.10)
Fixed Effects	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract
	No Consecutive Limit Days									
b_{up}	-0.25%	-0.04%	0.10%	0.06%	-0.10%	-0.06%	0.35%	-0.23%	-0.05%	-0.01%
	(-0.70)	(-0.21)	(0.73)	(0.74)	(-1.31)	(-0.60)	(1.10)	(-0.68)	(-0.11)	(-0.20)
b_{down}	0.42%	0.27%	-0.30%	-0.05%	-0.05%	0.11%	0.13%	-0.37%	-0.62%	-0.06%
	(1.22)	(1.04)	(-1.68)	(-0.60)	(-0.43)	(1.05)	(0.47)	(-0.73)	(-1.45)	(-0.85)
Fixed Effects	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract	Contract

Note: We examine the effect of limit days on subsequent one-day returns. The regression specification is the following:

$$r_{i,t+1} = a + b_{up} I^*_{Up} + b_{down} I^*_{Down} + \gamma FE_i + e_{i,t+1}$$

Where $r_{i,t+1}$ is the close-to-close futures returns for contract *i*. I^*_{Up} (I^*_{Down}) is an indicator variable equal to one if the previous day is a limit up (limit down) day for a contract on the same commodity, but not for *i*. FE_i are contract fixed effects to account for the latent heterogeneities among different contracts. T-statistics in parentheses are clustered by time to allow for arbitrary cross-sectional correlation.

Across all nine commodities, the average next day futures returns for those contracts that did not hit limits are not statistically different from their unconditional means following limit up or limit down occurrences along the curve. In other words, there is no difference in the average returns for these contracts whether the previous day was a limit day for another contract along the curve or not. Evidently, price continuation is restricted to those futures contracts that experience limits and do not spill over to other contracts on the same commodities. Even if a limit event comes as a result of a common shock to all contracts, the effect of the shock is concentrated and much stronger for the contract that hit limit compared to those that do not.

For robustness, we also examine a sample without any consecutive limit days. The results are similar to those for all limit days: Contracts hitting limits do not have material impact on contracts of the same commodity that do not hit limits. Next day average returns are not different from their unconditional means for all nine commodities.