

Limits to Arbitrage and the Skewness Risk Premium in Options Markets*

Thomas Ruf
University of British Columbia

JOB MARKET PAPER

This version: November 22nd, 2011

Abstract

Option prices, particularly those of out-of-the-money equity index puts, are difficult to justify in a no-arbitrage framework. This paper shows how limits to arbitrage affect the relative pricing of out-of-the-money put vs. call options (option-implied skewness). More importantly, I am able to quantify the extent to which this price differential is justified by the physical return distribution of the underlying asset and to what extent it represents a skewness risk premium. Analyzing the market for commodity futures options, I find that implied skewness becomes more negative, but realized skewness remains unchanged, when i) arbitrageurs hold larger net long positions in options, ii) long positions in the underlying asset are concentrated among fewer traders, iii) arbitrageurs are more financially constrained, and iv) recent realized skewness has been more negative. Thus, the skewness risk premium increases as limits to arbitrage get more severe. Finally, trading strategies designed to exploit these limits yield up to 2 percent per month after risk-adjustment.

*Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver BC, Canada V6T 1Z2 (e-mail: thomas.ruf@sauder.ubc.ca). I am thankful to Oliver Boguth, Murray Carlson, Lorenzo Garlappi, Maurice Levi, Kai Li, Alberto Romero, Mikhail Simutin, Tan Wang for helpful comments. Special thanks go to Eric Meschke and David Nuelle at Buttonwood Group Trading, LLC in Chicago for offering insights into commodity trading. Financial support by the Research Foundation of the Canadian Securities Institute is gratefully acknowledged.

1 Introduction

Empirical studies have shown that writing equity and equity index options, most notably out-of-the-money (OTM) index puts, yields abnormally high returns (Bakshi and Kapadia, 2003)¹. The origins of these abnormal returns are hotly debated. No-arbitrage asset pricing models are generally unable to explain them assuming reasonable parameter choices for risk aversion and crash frequencies (see e.g. Bates, 2000; Bondarenko, 2003)². In his comprehensive survey, Bates (2003) concludes that no-arbitrage arguments alone are likely inconsistent with observed option prices. Instead he suggests to focus on frictions to financial intermediation in option markets, drawing parallels between the options market and the market for catastrophe insurance (Froot, 2001).

One friction faced by financial intermediaries, the exogeneous demand for options, is the subject of the literature on demand-based option pricing. Bollen and Whaley (2004) and Garleanu, Pedersen, and Poteshman (2009) find evidence that systematic demand pressure in the market for index and equity options affects the level as well as the slope of the implied volatility function. This price effect is due to market makers being capital-constrained and not fully hedged. Both papers suggest that returns of options will also be affected but thus far, this literature has not related abnormal option returns to demand pressure directly. In other words, there is evidence that demand pressure makes some options more expensive than others (e.g. OTM index puts relative to OTM calls), but no one has established that they yield abnormally low returns as a consequence of demand pressure. In addition, variation in demand across time and across strikes is only one side of a limits to arbitrage explanation (Gromb and Vayanos, 2010); variation in the ability and willingness of market makers or arbitrageurs to absorb that demand is the other. How the capacity of arbitrageurs affects both prices and returns in the options market therefore remains an open question.

My paper aims to fill those gaps. First, I consider not just demand pressure, but various aspects of limits to arbitrage as well as interactions between different market frictions. In particular, I emphasize the role of arbitrageurs' capacity to absorb hedging demand rather than demand pressure alone. Second, I simultaneously analyze option prices and option returns, rather than each in isolation. Third and most important, using

¹A number of studies also find abnormal returns in equity options, e.g. from writing options of stocks with high idiosyncratic volatility (e.g. Goyal and Saretto, 2009; Cao and Han, 2011).

²Exceptions are Liu, Pan, and Wang (2005) who suggest that historical index option prices can be matched if investors dislike model uncertainty surrounding the true frequency of large negative jumps, and Benzoni, Collin-Dufresne, and Goldstein (2011) who assume an extremely slow mean-reverting belief about the likelihood of a crash.

the methodology of Neuberger (2011), I can distinguish between effects that alter option prices due to changes in rational investors' beliefs about the physical return distribution of the underlying asset or due to a change in the risk premium, where the latter may be driven either by investors' risk aversion or by the limits to arbitrage faced by market makers.

The variance risk premium is the subject of a large number of previous studies (Carr and Wu, 2009; Bollerslev and Todorov, 2011, Driessen, Maenhout, and Vilkov, 2009). Only Cao and Han (2011) investigate the effects of limits to arbitrage in that context. The focus of this paper is the relative pricing of the tails of the return distribution, i.e. prices of OTM puts relative to OTM calls (option-implied skewness), as opposed to the absolute level of prices (option-implied volatility). Thus, I measure and describe the risk premium for skewness. Nevertheless, I also briefly investigate the variance risk premium as a robustness check to my findings on skewness.

My work empirically distinguishes between beliefs and risk premia as determinants of levels and changes in option prices. To intuitively illustrate the difference between prices responding to changes in beliefs or risk premia, it is instructive to compare the market for OTM options with the market for catastrophe insurance, as suggested by Bates (2003). If the price for earthquake insurance goes up, it could be due to one of three reasons. Either beliefs about the frequency of such an event have changed, households have become more risk-averse towards the possibility of an earthquake, or insurers are more constrained in their ability to provide such insurance. Similarly, observing a high price for OTM put options may be due to a) a high physical probability of a large decline, b) high risk aversion or c) large frictions in the intermediation of that risk. The existing literature has not been able to empirically distinguish the former from the latter two explanations because the extreme nature of OTM option returns and the low number of observations make statistical inference difficult.

I use non-parametric measures of option-implied skewness and realized skewness developed by Neuberger (2011) to distinguish among these possibilities. Implied skewness reflects the relative prices of OTM puts vs. OTM calls and is the expectation of realized skewness under the $\mathbb{Q}$ -measure. The difference between realized and implied skewness is the skewness risk premium and can be thought of as a return to a particular swap contract that is long OTM calls and short OTM puts.³ I conduct the empirical analysis on the market for commodity futures options rather than index or equity options for a number

³Kozhan, Neuberger, and Schneider (2011) use these measures to quantify more precisely than was previously possible the extent to which risk premia for both variance and skewness are present in options on the S&P 500 index.

of reasons.

All empirical work using S&P 500 index options is bedevilled by the paucity of data, making estimation of the physical return distribution difficult. Options on a broad market index are even less suitable for analyzing the effects of limits to arbitrage vis-à-vis other effects, because financial constraints to market makers are strongly pro-cyclical (Adrian and Shin, 2010), while risk aversion is strongly counter-cyclical (Campbell and Cochrane, 1999). In other words, more severe financial constraints, higher risk aversion and negative realizations in index returns tend to occur jointly.

Options on a cross-section of stocks resolve some of these problems in theory. In practice, option demand across stocks is highly correlated for similar stocks, e.g. large or small, and jointly driven by investor sentiment (Lemmon and Ni, 2010). Finally, the non-parametric measures employed, especially those for skewness, require a large number of options at each point in time for each underlying to adequately measure prices in both tails. Unfortunately, the available datasets (e.g. OptionMetrics) would limit this type of analysis to the very largest stocks only.

Commodity futures options have none of these problems. My dataset consists of a cross-section of 25 U.S.-listed optionable commodities over a time period of 20 years, exhibiting low cross-correlations in both prices and demand patterns. Option chains of some commodities have quotes for more than 100 options at a time. I merge this dataset with data on the aggregate futures and options positions of different groups of traders, notably speculators, e.g. market makers and arbitrageurs, and commercial traders that trade for hedging purposes.

While the origin of the exogenous demand may be different for commodity futures options, the market structure is very similar. In particular, market makers face the same funding constraints and risks from the unhedgeable portions of their inventory regardless of the underlying. In fact, market makers in different options exchanges are often subsidiaries of the same parent companies. Further, trading desks of investment banks actively seek out arbitrage opportunities across all financial markets so that many of the arbitrageurs are identical across markets as well.

I consider four categories of potential limits to arbitrage effects: demand variables, measures of trader concentration, financial constraints of arbitrageurs and lastly, variables that proxy for recent profits and losses by arbitrageurs. For each set of variables, I investigate how they affect the relative pricing of OTM puts vs. calls, i.e. the option-implied skewness, as well as the measure of realized skewness and the difference between the two, the skewness risk premium. If a particular variable, X , is found to decrease

implied and realized skewness equally, it will have no effect on the risk premium. In other words, the change in prices caused by X is reflective of a change in the physical distribution and investors' rational beliefs about it. If, on the other hand, an increase in variable X decreases implied skewness, but does not affect realized skewness, its effect carries through to the risk premium. In this case, we can conclude that prices change due to a deterioration in the arbitrage capacity of speculators (or a change in risk aversion, depending on the context); we may say that e.g. OTM puts are too expensive relative to OTM calls in the sense that a portfolio that is long OTM calls and short OTM puts has negative returns on average when X is high.

I begin by showing some evidence from the cross-section of commodities where cross-sectional dispersion in average hedging pressure is correlated with average implied skewness and the average skewness risk premium. Motivated by this finding, I conduct the main analysis in a large panel dataset, focusing on the time series dynamics of skewness measures and variation in limits to arbitrage. My main findings are that implied skewness, but not realized skewness, becomes more negative when i) speculators hold larger net long positions in the corresponding options, ii) long (short) positions in the underlying futures market are concentrated among fewer (more) traders, iii) speculators are more financially constrained, and iv) recent realized skewness has been more negative. In addition, increases in financial constraints make the effects in i) and iv) more pronounced. In other words, prices are affected but (beliefs about) the dynamics of the underlying asset are not, so that (almost) all of the price change translates into a skewness risk premium. As limits to arbitrage become more severe the absolute size of the skewness risk premium increases.

As a robustness check, I investigate whether similar effects can be detected in the absolute level of option prices and in the variance risk premium. I confirm that financial constraints and recent realized volatility increase both option prices and the variance risk premium, thereby providing additional evidence in favor of a limits to arbitrage explanation.

While the use of non-parametric measures has great advantages over using option returns directly, they make it difficult to interpret the economic magnitude of each effect. Does a strategy designed to exploit limits to arbitrage give rise to significant returns? Is it implementable? To answer these questions, in the final part of the paper I form zero net investment portfolios of particular assets related to skewness. I find that a long/short portfolio of delta-hedged risk reversals formed on the basis of the size of net option exposure of speculators, for instance, yields very significant 2 percent per month,

even after risk-adjustment. Portfolios formed on the basis of trader concentration and lagged realized skewness yield smaller, but still significant returns.

The remainder of this paper is structured as follows. Section 2 provides some background information on commodity options and describes the sample used for the analysis as well as the variables related to hedging pressure. Section 3 motivates the limits to arbitrage explanation by looking at skewness and skewness risk premia in the cross section of commodities. Section 4 describes the skewness measures. Section 5 contains the design and results of the main empirical analysis conducted on a large panel dataset. Section 6 examines portfolios constructed on the basis of limits to arbitrage criteria. Section 7 concludes.

2 Background on Commodity (Option) Markets

While the academic finance literature has intensively investigated options on stocks and equity indices for decades, there is surprisingly little research on commodity futures option markets. In part, this is due to lack of widely available data sources on commodity futures options, but it may also reflect the relative lack of research into commodities in general and the obscurity in which commodities existed in the public's mind in the past.

Physical commodities as an asset class, let alone their respective option markets, have been literally unknown to the average investor until a few years ago. Markets for financial derivatives on physical commodities have existed for several hundred years and at least as long as stock markets, yet compared to financial asset classes (such as stocks and bonds) the demand for which can be characterized as purely speculative, markets for physical commodities are mostly dominated by commercial traders who participate with the intention of mitigating risk inherent to their line of business, in other words hedging, not speculation. Further, each commodity is somewhat unique because of fundamental differences in the dynamics of supply and demand of the physical product (seasonalities in production or consumption, storability).

A particular impediment with regards to commodity options is the lack of electronic trading venues in stark contrast to equity or even currency options. To the present, the majority of trading in commodity options takes place via open outcry in pits rather than electronically.⁴

Prior to the advent of commodity-related exchange-traded funds (ETFs) and exchange-

⁴For recent information on relative volumes on the CME, see <http://www.cmegroup.com/resources-for/files/options-update.pdf>.

traded commodities (ETCs), the investing public had neither the knowledge nor the tools to participate in these markets, even though the (historic) diversification benefits of commodities have been known for quite some time, e.g. Bodie and Rosansky (1980). A widely circulated paper by Gorton and Rouwenhorst (2006) on the equity-like performance of being long commodity futures further advanced the popularity of commodity investing. Domanski and Heath (2007) call this development the 'financialization' of the commodity space and find evidence that commodities have started to behave more like financial assets.

[Figure 1 about here.]

Figure 1 provides a rough idea of the size of commodity derivatives markets, and commodity options in particular, over time. The two panels contrast the market for exchange-traded commodity derivatives with the market for over-the-counter (OTC) derivatives. The left panel depicts the aggregate open interest (in \$B) of the 25 exchange-traded commodities in my sample. Because the sample is limited to U.S. listed, optionable commodities with a long history it under-reports the true level of open interest across all commodity exchanges. On the other hand, open interest does not correct for spread positions held within or across markets by the same agent. The panel distinguishes between open interest of Futures and Options and shows that the dollar amount of options contracts has at times reached up to half that of Futures, peaking at \$376B at the end of the first half of 2008 (compared with \$633B for Futures). The right panel shows the aggregate notional amount outstanding in OTC commodity derivatives as reported by the Bank of International Settlements in their 'Semiannual OTC derivatives statistics'⁵. At the end of June 2008, the BIS reported notional amounts for OTC forwards and swaps (OTC options) to be peaking at \$7.5T (\$5.0T). The OTC market is bigger by an order of magnitude. Nevertheless, it appears that both markets have exponentially grown (and subsequently ebbed) in a similar fashion over recent years. Further, OTC market makers such as commodity trading desks of large investment banks will in turn hedge their net exposure from the OTC market in the liquid futures market, thus leading to some level of integration between the two markets.

2.1 Futures and Options Data

To my knowledge, this paper is the first to use the futures options database of the Commodity Research Bureau in addition to their more commonly used futures database. To-

⁵Available at <http://www.bis.org/statistics/derstats.htm>.

gether, they contains end-of-day closing prices for a large number of U.S. and international futures and futures options markets covering financial indices, interest rates, currencies and commodities. For the purpose of this paper, I focus on major commodity futures markets in the U.S. for which exchange-listed futures options exist and have been liquid over the majority of the sample period.

[Table 1 about here.]

Table 1 lists the sample containing 25 commodities which can be roughly divided into 5 major groups: Agricultural, Energy, Meat, Metal and Soft.⁶ The table shows the first and last maturity of option contracts by commodity and the number of individual option chains having the necessary data to create option-implied and realized skewness measures at a remaining time to maturity of at least 30 (90) days.⁷ My data stops at the end of August 2010, so that the last option expiration is either the August or September contract of that year with two exceptions. At the end of 2006, unleaded gasoline drops out of the sample due to lack of options and COT data and is replaced by a related commodity, RBOB gasoline, reflecting a change in economic relevance. For the same reason, data on pork bellies ends in early 2008. The appendix describes the structure of the data, the cleaning process and construction of implied and realized measures from option prices in more detail.

2.2 COT Variables

In addition to a long history of options quotes, I require data on the position of traders for these commodities as published by the U.S. Commodity Futures Trading Commission (CFTC) in their Commitment of Traders (COT) reports. Because the CFTC only publishes this data if the number of active participants in the market and the size of open interest are large enough, this further ensures that the commodities in the sample are liquid and economically important.⁸

⁶My set of commodities is similar to the U.S. based commodities used in Gorton, Hayashi, and Rouwenhorst (2007). Compared to Szymanowska, De Roon, Nijman, and Van den Goorbergh (2010), I exclude platinum and palladium, because there are no liquid options available, but I include natural gas.

⁷To compute implied measures, I require a sufficient number of option quotes of out-of-the-money options in both tails for a particular maturity. The computation of realized skewness requires changes in an implied variance contract as inputs, thus the same requirements have to be fulfilled for all days leading up to expiration.

⁸Currently, the report is published every Friday detailing data as of Tuesday of the same week. Prior to October 1992, COT reports were published twice a month with a reporting lag of 6 business days. Markets are excluded if there are less than 20 traders present with positions above a commodity-specific reporting

In the COT reports, traders are divided into 3 groups: commercials, non-commercials and non-reporting. The first group is thought to mostly consist of producers and consumers of the underlying commodity that use the Futures market to hedge future production or consumption, thus they are typically called hedgers. The second group consists of large institutions that trade in the commodity for financial gain, i.e. trading desks, hedge funds, commodity trading advisors (CTA) and commodity pools, hence commonly called speculators. The last group consists of small traders with positions below the reporting limits, also called the public or retail investors. Apart from total open interest, number of reporting traders, the report details long and short positions held by members of each group as well as concentration ratios for the long and short side of the market.

Hedging pressure in futures: The most commonly used measure derived from COT data in the literature on commodity futures is hedging pressure (HP), defined as the normalized net *short* exposure via futures contracts in commodity i by commercial traders as a group:

$$HP_{i,t} = \frac{\sum \text{comm. short fut.} - \sum \text{comm. long fut.}}{\sum \text{comm. short fut.} + \sum \text{comm. long fut.}} \quad (1)$$

By construction, ‘HP’ lies between -1 and $+1$. Because futures positions must be secured by margin accounts holding cash or Treasuries, the denominator can be interpreted as being in proportion to the total amount of capital at risk for hedgers as a group. ‘HP’ has been found to positively predict futures returns (see e.g. Bessembinder (1992), De Roon, Nijman, and Veld (2000)). A corresponding measure of net *long* exposure can also be constructed for the group of speculators, denoted by ‘Sp’:

$$Sp_{i,t} = \frac{\sum \text{spec. long fut.} - \sum \text{spec. short fut.}}{\sum \text{spec. long fut.} + \sum \text{spec. short fut.}} \quad (2)$$

Because hedgers and speculators generally constitute the overwhelming majority of total futures open interest, the numerators are quite similar for most commodities and, due to the slight change in definition, of the same sign. However, even in the absence of retail traders as a third group, HP differs from ‘Sp’ due to the denominators being the sum of total positions held by each group.

Options hedging pressure: Apart from a Futures-only report, Combined Futures-and-options reports are available since April 1995, where option positions are transformed

limit. Futures-only COT data generally begins in 1986, while the so-called combined COT reports only start in April of 1995.

into Futures-equivalents using each option’s delta (based on a pricing model). To get a measure of demand pressure in the commodity options space, I back out the long and short positions of commercials in the options market alone and compute an analogue hedging pressure via options only (OHP):

$$OHP_{i,t} = \left(\frac{\sum \text{comm. short options} - \sum \text{comm. long options}}{\sum \text{comm. short options} + \sum \text{comm. long options}} \right) \quad (3)$$

As with futures, the flipside to ‘OHP’ is the normalized net *long* exposure of speculators ‘SpO’.

$$SpO_{i,t} = \left(\frac{\sum \text{spec. long options} - \sum \text{spec. short options}}{\sum \text{spec. long options} + \sum \text{spec. short options}} \right) \quad (4)$$

As before, ‘OHP’ and ‘SpO’ will differ due to different denominators even in the absence of small investors holding option positions. In summary, both HP and OHP give an indication of how “one-sided” hedgers’ aggregate positions are relative to their total exposure in each market, while Sp and SpO do the same for speculators.

Long and short trader concentration: The COT report contains another item that, to my knowledge, has not been investigated before. In addition to all the previously mentioned items, the CFTC also reports concentration ratios separately for the long and the short side independent of trader classification, i.e. the proportion of long (short) open interest that is held by the largest N long (short) traders for that commodity. I use the concentration ratios reported in the Futures-only data for $N = 8$, e.g. for the long side

$$CRL_{i,t} = \frac{1}{OI_{Fut,i,t}} \sum_{j \in \text{Top } 8} (\text{Long Fut.})_j \quad (5)$$

Ceteris paribus, if a larger fraction of futures is held by a small number of traders, those traders may be less inclined to add to their already concentrated positions, while the absence of other traders means that a given demand shock has to be absorbed by fewer individual traders. The last 4 columns of Table 1 show time series averages for the two concentration ratios and the two options-only demand pressure variables.

3 Cross-sectional Evidence

To see how skewness (as well as variance) are related to limits to arbitrage, I begin by computing time series averages of implied and realized variance and skewness and their

differences, the risk premia for variance and skewness, for the cross section of commodities in my sample. While the focus of this paper is on skewness and the skewness risk premium, it may be helpful to readers unfamiliar with the notion of the latter to go through the corresponding terms of variance first. In addition, I will discuss the connection between limits of arbitrage and the variance risk premium in Section 5.6 as a robustness check of my findings for the skewness risk premium. In the interest of brevity, I reserve the exact definitions and computational details of all the measures involved for a later section as well as the technical appendix.

3.1 Variance Risk Premium

The existing literature (e.g. Britten-Jones and Neuberger, 2000; Jiang and Tian, 2005) defines the variance risk premium (VRP) as the difference between ex-ante model-free implied variance (MFIV) and ex-post realized variance (RV). Carr and Wu (2009) propose to consider those three terms in the context of a financial contract, the variance swap. While not traded on exchanges, variance swaps can be entered into over the counter with investment banks. The buyer of a variance swap pays the implied variance at the time of transaction as the fixed leg of the swap and receives as floating leg the variance realized over the term of the contract. If on average implied exceeds realized, the buyer of the swap tends to overpay in the form of a risk premium.

In the case of equity markets, Carr and Wu (2009) find a large negative VRP and suggest that investors dislike states of high volatility and are willing to pay the variance risk premium in order to insure against those states. The evidence on VRP in individual stocks is more mixed. While Carr and Wu (2009) find some negative VRPs in their sample of 30 large cap stocks, Driessen et al. (2009) report marginally positive VRPs on average for the constituents of the S&P 100 index.

[Table 2 about here.]

Table 2 presents time series averages of implied and realized variance and the variance risk premium for my sample of commodities as well as the S&P 500 index. For readability, the variance measures are translated into annual volatilities. The VRP is also annualized and multiplied by 100. Measures are based on options with maturities of between 90 and 99 days, with the exception of the second row for the S&P which is based on 30 to 36 days to maturity. This row is included because prior to 2008, only four S&P 500 option series per year were listed more than 90 days prior to expiration. The introduction of

longer maturities with the beginning of the Great Financial Crisis is biasing the sample towards high realized volatility observations, in turn causing the VRP to appear smaller.

For commodities, average implied volatility is always larger than the realized variety. Consequently, the variance risk premium, the difference between realized and implied variance, is always negative and (Newey-West adjusted) standard errors indicate that almost all estimates are very significantly different from zero. In unreported tests, I find that short-term VRPs generated by options with 30 to 36 days to maturity are even more consistently negative. This means that just like for equity indices, but unlike individual stocks, volatility in commodities and thus commodity options are on average overpriced. It might be interesting to investigate whether part of this premium is connected to exposure to some form of systematic variance factor as Carr and Wu (2009) do for the cross-section of stocks. However, as the focus of this paper is limits to arbitrage, I will leave this topic to future research.

3.2 Skewness Risk Premium

While measures of option-implied skewness have been proposed in the literature (Bakshi, Kapadia, and Madan, 2003), until recently no realized counterpart was known. Neuberger (2011) and Kozhan et al. (2011) (henceforth KNS) overcome this problem and propose a pair of measures, implied and realized skewness, whose difference can be interpreted as a risk premium for skewness. The realized skewness of KNS is mostly determined by the interaction between returns and forward-looking implied variance. If variance increases concurrent with negative returns, the realized measure becomes consequently negative. The implied measure is the risk-neutral expectation of realized skewness.

Again, one may think of this in the context of a swap. The buyer of a skewness swap pays as fixed today's price, i.e. the implied skewness, and receives as floating the realized covariation between returns and variance over the term of the contract. If on average implied lies below (above) realized, i.e. the buyer expects more negative (positive) skewness, the skewness risk premium (SRP) is positive (negative).

[Table 3 about here.]

Table 3 shows measures of average implied and realized skewness as well as the SRP for the cross-section of commodities in the sample. Numbers are based on maturities between 90 and 99 days (except the second line for the S&P).

The column 'BKM' contains an implied, normalized measure of skewness of log returns (based on Bakshi et al., 2003) and shows that most commodities are far less negatively

skewed than the equity market. In fact, the agricultural commodities in particular tend to have positive skewness. The model-free implied skewness (MFIS) (based on Neuberger, 2011) is a non-normalized measure representing risk-neutral expectations of the third moment of the percentage return, not log return, over the entire term. The difference between log and percentage returns is the reason why BKM skewness is negative for more commodities. Further, because MFIS is not divided by some function of variance, higher average volatility will make MFIS larger in absolute terms. The next two columns show the average sum of realized cubed returns ‘ r^3 ’ and the KNS measure of realized skewness ‘RSkew’, which is the sum of two terms: cubed returns as well as the covariation between returns and changes in implied variance.

Note that for most commodities and the index, the majority of realized skewness does not come from cubed daily returns, in fact the two terms occasionally have opposite signs. Heating oil, for instance, seems to have more large negative returns than positive ones on average. However, the difference between the two measures suggests that more than 100 percent of total skewness comes from the positive covariation between returns and implied variance. This indicates that, unlike for the equity market, large positive returns tend to occur together with increasing volatility more than large negative returns for many commodities.

Lastly, the skewness risk premium, just like implied skewness itself has varying signs across commodities in the sample. The S&P 500 index has a strongly significant and positive SRP when considering 30-day variance swaps.⁹ Implied skewness is on average much more negative (at $-.60$) than realized skewness (at $-.36$). This is what makes strategies that sell OTM index puts so profitable: index puts are much more expensive relative to index calls beyond what can be justified by the ex-post physical distribution of index returns. Commodities, on the other hand, exhibit both positive and negative SRPs. The meats as well as natural gas are similar to the S&P 500 in that implied skewness is more negative on average than realized skewness, while agricultural commodities and the precious metals in particular exhibit the opposite behavior.

3.3 Limits to arbitrage

What can explain the notable dispersion of average implied skewness as well as SRPs across commodities in Table 3? Cross-sectional variation in implied skewness could potentially be due to corresponding differences in beliefs about the return distribution of

⁹As with variance swaps, the lack of power for 90-day skewness swaps is due to the smaller sample size and relatively larger number of observations that cover the recent financial crisis.

the underlying asset. The price for natural gas certainly behaves quite differently from that for corn in many ways. However, the existence of and variation in the skewness risk premium show that differences in the price dynamics of the asset alone are not sufficient as an explanation.

In what follows, I argue that the cross-sectional pattern apparent in Table 3 is the consequence of limits to arbitrage (Shleifer and Vishny, 1997) in the options market. After showing some supporting evidence for the cross-section of commodities in this section, I conduct a sequence of tests in a panel dataset to show in more detail what types of frictions are present in the options market and how they affect the process of financial intermediation.

In any financial market, arbitrageurs face constraints with regards to the amount of capital they can commit to any one position, the amount of losses they can tolerate before they are forced to liquidate or the ease with which they can access additional funding. In the face of demand pressure - be it in futures or options - from institutional investors (Bollen and Whaley, 2004), from ‘end-users’ (Garleanu et al., 2009) or from commercial hedgers (e.g. Hirshleifer, 1988), asset prices will depend on arbitrageurs’ ability to raise funds and their willingness to add to their existing positions to accomodate that demand. Both ‘end users’ and hedgers can be considered as price takers and consumers of liquidity. Market makers and large speculators, on the other hand, tend to be liquidity providers.

Already Keynes (1930) posits in the Theory of Normal Backwardation that futures markets help commodity producers to hedge part of their future production by selling it forward at prices determined today. Since futures are in zero net supply, someone else has to be long the corresponding amount of contracts. In Keynes’ model, speculators fill this gap, but they demand compensation in the sense that the futures price must lie below the expected future spot price, thereby introducing a futures risk premium into the market. The extant literature has generally found results supportive of the theory (Bessembinder, 1992; Bodie and Rosansky, 1980; De Roon et al., 2000).

[Figure 2 about here.]

The implicit assumption in Keynes’ theory, however, is that speculators must be risk averse to some extent. Otherwise no risk premium would be necessary. Only recently has the literature begun to explicitly model this.¹⁰ For instance, Etula (2010) and Acharya, Lochstoer, and Ramadorai (2011) both assume VaR constraints for speculators/financial

¹⁰Hirshleifer (1988) uses quadratic utility for all agents in his model, but does not discuss the reasons for this choice.

intermediaries, which makes them effectively risk averse, giving rise to limits to arbitrage in both models.

The same principle may apply in the corresponding options market. Assume that instead of via futures, hedgers would like to hedge their natural long position with short positions in options, e.g. by buying puts or selling calls or any other strategy with a negative delta. Someone has to write the put or buy the call, i.e. take a net long position via options. Just like in the futures market, market makers or, more generally, speculators fill that void, but option prices are affected because of the inability of speculators to fully hedge themselves (Garleanu et al., 2009). *Ceteris paribus*, it thus seems sensible to assume that for commodities in which hedgers demand more short exposure via options, (OTM) puts would be more expensive while calls would be relatively less expensive.¹¹

Consider the top left panel of Figure 2, which plots implied skewness from Table 3 against the OHP measure from Table 1. It appears that the level of implied skewness is significantly negatively related to the extent to which hedgers on average hold short positions via options in the commodity. Thus, as expected, commodities in which hedgers generally hold larger short options positions, i.e. have a more positive OHP, tend to have more negative skewness and vice versa. This supports the notion that prices are affected in a way that is consistent with limits to arbitrage and demand pressure. Does net short option demand also affect the risk premium for skewness? The top right panel shows the relation between hedging pressure and the SRP, more specifically the Newey-West adjusted t-statistic testing the significance of the SRP estimate. This choice is made because the skewness measures and thus the premium are not normalized by some function of variance, and thus the absolute size of premia varies by an order of magnitude across commodities. We see that more net short hedging pressure via options is related to a more positive SRP, indicating that implied skewness is on average too negative compared to realized skewness for those commodities, i.e. OTM puts are overpriced relative to OTM calls.¹² OHP explains slightly more than one third of total cross-sectional variation of both implied skewness and the SRP, respectively.

A key point I want to emphasize is that it is not the amount of hedging pressure itself that determines prices. The same amount of hedging demand may lead to different prices depending on how binding the aforementioned frictions, i.e. limits to arbitrage, are

¹¹It is difficult to make statements about the absolute expensiveness of options solely based on the delta exposure of speculators or hedgers. Rather, one would require information on the vega exposure, i.e. whether hedgers are long or short volatility, which is not available in the COT data.

¹²Relative overpricing means that a portfolio that is long puts and short calls has negative returns in excess of the risk-free rate on average.

for arbitrageurs at that time. Thus it is likely more relevant to judge the impact from net hedging demand not relative to the total positions of hedgers, but relative to those of speculators. In addition, it is the net exposure actually absorbed by speculators that is relevant in the context of limits to arbitrage. I replicate the previous two panels at the bottom of Figure 2, this time using the net exposure of speculators ‘SpO’, scaled by the total capital at risk for speculators as a group. While the general inference does not change, the increase in fit is notable for both plots. An R-squared of 70 percent implies a correlation of over 83 percent in the case of implied skewness. The direct comparison of the two measures of hedging pressure lends strong support to a limits to arbitrage interpretation because it shows that speculators’ capacity (to absorb hedging demand) can explain the observed patterns better than hedging pressure.

The cross-sectional finding presented in Figure 2 merely shows a correlation and does not prove causality. But it serves as a starting point to investigate the different dimensions of limits to arbitrage as they exist in the market for commodity futures options. For instance, it is possible that some unknown factor causes both the observed cross-sectional patterns in skewness as well as different propensities by commercial traders to use options to hedge different commodities. Differences in individual industry or market structure are potential candidates. Therefore, in order build additional support for a limits to arbitrage explanation, I investigate in Section 5 how different frictions facing arbitrageurs affect the price of skewness (i.e. option-implied skewness) as well as the return on skewness (i.e. the skewness risk premium) *over time*, with the amount of hedging pressure being only one of those frictions.

4 Measures of Skewness

The literature on option-implied moments has brought forth a number of measures for skewness. It is still common to use the difference in implied volatility between OTM put options and ATM options, i.e. the volatility slope or volatility smile, as a proxy for risk-neutral skewness implied in option prices (Duan and Wei, 2008; Lemmon and Ni, 2010). This is valid because a one-to-one mapping exists between the volatility smile and the risk-neutral density of the underlying asset return (Rubinstein, 1994). Bakshi et al. (2003) (henceforth BKM) propose a more direct measure of option-implied skewness. Via the spanning approach by Bakshi and Madan (2000), BKM use options data to replicate the risk-neutral expectations of the first 4 uncentered moments of the log return. They then systematically build expressions for standardized skewness (and kurtosis) from

those building blocks. The construction uses Taylor series expansions and is thus only approximate. This measure has been employed e.g. by Dennis and Mayhew (2002), Conrad, Dittmar, and Ghysels (2009), Rehman and Vilkov (2010) for index and equity options. In practice, BKM skewness is more prone to produce negative outliers due to the log transformation. Note that the skewness implied by this measure is the skewness of the return over the life of the underlying option chain, not the average daily skewness as is commonly measured in asset pricing (Boyer, Mitton, and Vorkink, 2010) using daily stock returns.

Kozhan et al. (2011) and Neuberger (2011) (henceforth KNS) propose a different measure for longer-horizon skewness that measures the risk-neutral expectations of the third power of percentage returns, the so-called model-free implied skewness (MFIS). Under the assumption of continuous rebalancing, it is equal to the integral of the product of asset return and changes in implied variance (for details see the appendix):

$$\text{MFIS}_{0,T} = \mathbb{E}^{\mathbb{Q}} \left[3 \int_0^T dv_{t,T}^E \left(\frac{dS_t}{S_t} \right) \right] \approx \mathbb{E}^{\mathbb{Q}} [r_{0,T}^3] \quad (6)$$

It, too, is based on the spanning approach. Its main drawback is that it is not standardized by some power of volatility, making the interpretation of values less intuitive. It has the advantage, however, that only one integral needs to be numerically approximated, while the BKM measure (see Equation 18 in the appendix) is a ratio of several such integrals. Most important and key to the focus of this paper, the KNS measure has a realized counterpart. The question how much of the risk-neutral expectation reflects the true return distribution and how much reflects risk aversion by investors insuring against adverse outcomes was previously hard to answer. No natural counterpart to implied skewness was available that was able to measure realized skewness over some time period. Equation 6 reveals that the realized counterpart can be recovered as long as the integral can be computed. This only requires the existence (or at least, the replication) of a particular variance contract. A risk premium for skewness can be defined as the difference between the expectations of realized skewness under the risk-neutral and the physical measure, i.e.

$$\text{SRP}_{0,T} = \mathbb{E}^{\mathbb{Q}} \left[3 \int_0^T dv_{t,T}^E \left(\frac{dS_t}{S_t} \right) \right] - \mathbb{E}^{\mathbb{P}} \left[3 \int_0^T dv_{t,T}^E \left(\frac{dS_t}{S_t} \right) \right] \quad (7)$$

In practice, the $\mathbb{P}$ -measure expectation is replaced by its ex-post realization, abbreviated RSkew, to compute a realization of the skewness risk premium.

The empirical analysis in Section 5 is based on the KNS measures of implied and realized skewness, abbreviated MFIS and RSkew. Because MFIS (just like MFIV) indicates the expected distribution of returns over the remainder of the life time of the option chain, it gets naturally smaller as time progresses and ultimately converge to zero at expiration. As a first step, I annualize MFIS, then I take an average over a short range of days to avoid outliers on individual days from impacting results. For one-month (3-month) skewness, I average MFIS over dates with remaining maturities in days of $I_{1\text{month}} = (30, \dots, 36)$ ($I_{3\text{months}} = (90, \dots, 99)$). More formally, for commodity i and expiration date T , given a maturity m , e.g. 1-month skewness, I average over observations at times t where $T - t \in I_m$.

$$\text{MFIS}_{i,m,T} = \frac{1}{|I_m|} \sum_{T-t \in I_m} \text{MFIS}_{i,t,T} \quad (8)$$

Finally, because MFIS has an excess kurtosis of 64 in my sample, I take the signed third root of this average (excess kurtosis -0.4).

$$\text{MFIS}_{i,m,T}^{1/3} := \text{sign}(\text{MFIS}_{i,m,T}) (|\text{MFIS}_{i,m,T}|)^{1/3} \quad (9)$$

Just like raw MFIS has an extremely heavy-tailed distribution, so do the raw realized skewness as well as the skewness risk premium (excess kurtosis is 44 for the latter). I define the signed third root of realized skewness in an identical fashion to $\text{MFIS}^{1/3}$, i.e.

$$\text{RSkew}_{i,m,T}^{1/3} := \text{sign}(\text{RSkew}_{i,m,T}) (|\text{RSkew}_{i,m,T}|)^{1/3} \quad (10)$$

where $\text{RSkew}_{i,m,T}$ represent the average annualized realized skewness over some small time window. To get a measure for the skewness premium that is better suited for analysis than raw SRP, we compute $\text{SRP}^{1/3}$ as follows: Within the estimation window (e.g. 30-36 days before maturity for 1-month measures), I compute the signed third root of realized and implied skew separately, take the difference and only then compute an average. The sequence is important to ensure that values for realized skewness are only included on days where the implied measure is non-missing.

$$\text{SRP}_{i,m,T}^{1/3} := \frac{1}{|I_m|} \sum_{T-t \in I_m} [\text{RSkew}_{i,t,T}^{1/3} - \text{MFIS}_{i,t,T}^{1/3}] \quad (11)$$

Note that in the case of SRP, the superscript $^{1/3}$ is merely for notational purposes to distinguish it from the raw premium measure. The transformation of signed roots takes away the extreme nature of the distribution of the skewness measures even better than scaling by some power of variance does. The goal is to get a sense of how certain factors affect skewness and the SRP on average without letting outliers dominate the analysis. The downside of this method is that the economic importance of each effect is difficult to ascertain. Section 6 aims to provide some answers in that regard.

5 Time Series Results

In Section 3, I provide evidence of substantial dispersion in the average skewness across commodities. The cross-sectional results suggest that the net option exposure of speculators may play a role in the average level of skewness. This is visible in scatter plots for both option-implied skewness as well as the risk premium for skewness. In this section, I focus on the time series dynamics of skewness rather than cross-sectional averages using a large panel data set. All variables involved, in particular the skewness measures, are de-measured separately for each commodity in order to filter out any level effects that may obfuscate the results.

5.1 Design of the Time Series Analysis

Each set of results in this section is presented in the form of two panels, one for ex-ante implied skewness and one for the ex-post realized skewness risk premium, or to be precise, the transformations defined in Equations (9) and (11). By taking the difference between the coefficients within each column, the corresponding coefficient for the realized skewness can be recovered. If a factor that affects implied skewness significantly does not show significantly (with the opposite sign) in the results for the skewness premium, it means that realized skewness was affected in a similar fashion as implied skewness, leading to the effect being insignificant for the difference, i.e. the premium. Thus, the factor seems to be properly priced into option prices and its effects on skewness expected. If, on the other hand, a factor affects both implied skewness and the premium significantly (with opposite signs), we conclude that the factor does not influence realized skewness and thus option prices appear to include a premium in response to that effect.

All dependent variables that are used to explain skewness are known prior to the time window where measurement of implied skewness, i.e. the $\mathbb{Q}$ -measure expectation of future

realized skewness, takes place. Consequently, they are also known prior to the realization of actual skewness over the corresponding time period. Thus, the regressions have predictive character and just like forecasting regressions for equity returns generally low R^2 . The choice of 3-month horizons are a compromise between having timely predictors and less noise in the realized measure.

A number of control variables are included in each regression, but have been omitted in all of the following tables to conserve space and bring into focus the key results. These are the percentage change in futures open interest over the last 6 month, the lagged 3-month return of the underlying futures contract (i.e. momentum) and the current convenience yield that is implied by the front futures contract and the futures contract that is closest to a 6-month maturity. Growth in open interest may be driven by new speculators or hedgers coming to the market, potentially leading to a price impact on options. The inclusion of last two variables is motivated by the existing literature. Brunnermeier, Nagel, and Pedersen (2008) use momentum in their analysis of skewness in carry trade returns. The convenience yield reflects the potential scarcity of the physical commodity and is affected by the probability of low inventories (Gorton et al., 2007). It seems natural to assume that skewness would react to the probability of low-inventory, high-price states. These variables add some explanatory power to the regression, but are mostly insignificant. The exception is the growth in open interest, which increases both implied as well as realized skewness by a similar extent, leaving the SRP unchanged.

Further, as the findings below reveal, there are 3 particularly strong effects on the relative pricing of options: the measure of speculators' net exposure in options (SpO), the long concentration ratio (CRL) and the lagged realized skewness. Each of these is included in the regressions that follow unless its effect is the focus of the investigation for that particular set of results. For instance, the table in Section 5.2 displays the effects of different types of demand pressure, among them the SpO measure. Thus, SpO is only included when specifically indicated, while the other two measures are included in every column.

5.2 The Effect of Demand Pressure

The literature on hedging pressure in futures (e.g. Bessembinder and Chan, 1992) finds the futures risk premia to be predictable by the measure HP. Ceteris paribus, higher pressure leads to greater price concessions by hedgers in the futures market and increases the return that speculators can expect on average. Here, I test if demand pressure variables affect any of the skewness measures. My analysis includes the hedging pressure and speculator

net exposure present in the underlying futures market as well as the equivalent measures in the futures options market.

[Table 4 about here.]

Table 4 shows the results from regressing implied skewness as well as the skewness risk premium on a sequence of demand pressure variables. The bottom three lines indicate the dimensions of the sample and the number of non-missing observations by column. Within each column, the sample is identical for the two panels.

The first two columns indicate that neither hedging pressure (HP) nor the scaled net exposure of speculators (Sp) has any affect on skewness. More interesting are the following 2 columns containing the corresponding measures based on the option market, OHP and SpO. Both have a significantly negative impact on implied skewness and a positive impact on the premium of nearly identical magnitude, which means that realized skewness is not affected. The sign of the effects agree with the results shown in the scatter plots. Not only do commodities with on average high levels of SpO have more negative skewness on average, there is also evidence that commodity skewness reacts to the level of SpO over time. Further, the effect is absent for realized skewness. Puts are thus overpriced and calls relatively undervalued in times where SpO is high, i.e. where speculators as a group have written a lot of puts and bought a lot of calls as a proportion of their total overall positions.

A key result is visible in column 5, containing both demand pressure variables for the options market. Notably, SpO dominates OHP both for implied skewness as well as the risk premium. This impressively supports the idea that it is not alone the amount of demand pressure coming from hedgers, but how this demand is absorbed by speculators. If total open positions by speculators are small, the same amount of net hedging pressure will have a larger effect. This result can also help delineate between the effects of changes in risk aversion and limits to arbitrage. The hedging pressure in options scaled by hedgers total open interest can be regarded as a proxy for the effective risk aversion of that group, while SpO, i.e. absorbed hedging pressure scaled by total speculators' capital at risk, proxies for limits to arbitrage. We see that the latter matters more for both prices and premia.

Column 6 conducts a robustness check splitting the SpO variable conditional on its sign. At least for implied skewness, a difference in slope is discernable. To test potential non-linearity of the effect of SpO, column 7 adds a square term. The fact that the linear term remains highly significant in the presence of the square term indicates an asymmetric

effect. But more importantly, the significance of the square term provides evidence that the marginal effect of one more unit of hedging pressure absorbed by speculators has an increasing price effect.

5.3 The Effect of Trader Concentration

A limits to arbitrage explanation of the dynamics of relative option prices where prices depend on the ability and willingness of speculators to enter positions opposite to those of hedgers would suggest that prices may be impacted by both the number of (large) speculators present in the market and the share of their portfolio that is allocated to any one market. On the one hand, it may affect their ability to take on additional exposure if they are already heavily exposed to a particular market. On the other hand, those knowing that they face less competition in a market may strategically offer to trade at less favorable prices (Oehmke, 2009). Both channels affect prices in the same way.

One way to catch this notion of individual exposure and strategic effects is the proportion of long (short) open interest that is concentrated in the hands of the largest N traders in that market. This measure would appear to be superior to the pure number of traders, which does not say anything about their size or capability. Looking at a concentration ratio offers a better glimpse into the supply and demand picture. For instance, if almost all open interest is concentrated among the largest traders it suggests that the market is not deep and that possibly those large traders are near or at their limits with regards to the exposure they are allowed or willing to have. On the other hand, a market where large traders make up a small share of total open interest indicates that many traders have positions and none is dominating; competition is high and the individual exposure of traders is relative low.

[Table 5 about here.]

Table 5 contains a sequence of regression results relating both skewness measures to a set of variables of trader concentration. The first 3 columns show that the level of concentration appears only to play a role on the long side of the futures market. The effect there is significantly negative for implied skewness and even more significant for the premium. In other words, in times when concentration is particularly high on the long side of the underlying futures market, puts are particularly expensive relative to calls and in fact overpriced. The explanation in the context of limits of arbitrage would be that the small number of long speculators currently present is less willing or able to add long

positions to their portfolios, i.e. write puts or buy calls. This pushes the prices of puts up and down for calls.

Column 4 conditions the long concentration on the sign of speculators' exposure in the futures market. While implied skewness exhibits little difference between the two subsets, the skewness risk premium shows that the pricing effect is stronger in times when speculators are long. Consistent with the notion of limits to arbitrage, the concentration ratio plays a stronger role when it is the speculators that are highly concentrated.

Using levels of concentration may be able to capture persistent effects, but may fail to uncover transitory effects. In a similar fashion as some demand pressure has only a temporary effect on prices until the demand is absorbed, it may be the case that it takes a small number of traders some time to adjust their portfolios to demand shocks. The last 2 columns test for temporary effects. Notably, after controlling for level effects, I find that skewness does react to changes in trader concentration on the short side, but that the effect is only temporary. It is not immediately clear, why the effects are permanent for one side and temporary for the other.

5.4 The Effect of Financial Constraints

The recent literature has shown that arbitrageurs, and as a result asset prices, are sensitive to changes in their ease to access funding, i.e. credit (Adrian and Shin, 2010; Adrian, Etula, and Muir, 2011). Acharya et al. (2011) and Etula (2010) show that the balance sheets of broker-dealer firms offer a good glimpse at the their ability to engage in arbitrage in the context of commodity futures. I use the measures they propose, i.e. the effective risk aversion $\hat{\Phi}$ of broker-dealers and its year-over-year change $d\hat{\Phi}$ as in Etula (2010), and the related measure of growth of household assets relative to the growth of broker-dealer assets employed by Acharya et al. (2011), also on a 12-month rolling basis. The data is extracted from the quarterly flow of funds database of the Federal Reserve Board. For all three measures, higher values mean a deterioration in financial conditions.

[Table 6 about here.]

The first 3 columns of Table 6 show those results. Consistent with the empirical artifact that speculators are generally more active on the long side, I find that deteriorations in the access to funding pushes implied skewness to be more negative, i.e. puts rise in price relative to calls. It seems plausible to assume that in times of less credit availability, realized returns of financial assets may be more negatively skewed as well. Notably,

comparing implied skewness and the risk premium, I find that this is not the case for commodities. Almost all of the price effect translates into a premium.

In addition to the pure effects from changes in funding constraints of arbitrageurs, I am also interested whether other effects related to limits to arbitrage are amplified under increasing constraints. For instance, does the effect from SpO in Table 4 get larger in times where arbitrageurs suffer from reductions in the size of their balance sheets? This is what the next 3 columns in Table 6 try to answer by including the interaction terms between proxies of financial constraints and SpO. The coefficients for the constituent terms stay remarkably stable. More importantly, recent increases in financial constraints for broker-dealers amplify the effect of hedging pressure in the options market. This makes intuitively sense, as I would expect that speculators react with larger price adjustments in times of stress for a given amount of hedging demand.

The final column investigates the effect of recent volatility on skewness. Past volatility by itself is not useful with regards to skewness, because assets that tend to have positive (negative) skewness will be more positive (negative) skewed in times of higher volatility. Thus, a linear specification is confounded by opposing effects. However, the interaction term between SpO and lagged volatility is significantly negative. This suggests that speculators react more sensitively to hedging pressure for commodities that recently experienced high volatility. Again this makes sense, as trading capital is allocated on the basis of perceived risk. The fact that the cross-term is not significant and the coefficient only have as large for the skewness premium indicates that speculators are at least partially right in their pricing policy.

5.5 The Effect of Lagged Realized Skewness

The effect of lagged volatility suggests the possibility that recent experience for speculators in the individual asset may affect their abilities to take on future risk in the same asset.

[Table 7 about here.]

I test this by looking at recent realized skewness of the individual commodity. The top row in Table 6 shows that recent realized skewness has a very substantial positive effect on current implied skewness. Negative recent skewness leads market participants to price in more negative skewness in the future (under the $\mathbb{Q}$ measure) and vice versa. The corresponding coefficient on the premium suggests that more than half of the effect in implied skewness is due to an actual effect on realized skewness going forward. Yet the

premium is still significant, so that option prices react by more than what is justified on average.

Columns 2 to 4 contain interaction terms between the financial constraint proxies and lagged skewness. As with SpO, the effect of recent realized skewness on current prices is amplified in times of financial stress for broker-dealers, but the effect carries through to the premium to a smaller extent.

The last two columns, investigate if the pricing response to lagged skewness depends on speculators aggregate positions during the time the skewness was realized. The cross term between lagged skewness and the lagged average level of ‘Sp’ yields a weakly significant effect on the premium. I get stronger results when splitting lagged skewness into times when speculators were net long vs. net short in futures. Notably, the premium effect due to lagged realized skewness is limited exclusively to those times where speculators were long. For instance, after experiencing negative realized skewness while being long, speculators as a group demand especially large price concessions, a large part of which constitutes a risk premium.

In summary, all of the findings in this section are consistent with a limits to arbitrage explanation for the case of commodity options markets, whereby ability and willingness of speculators to absorb demand shocks from hedgers, and thus option prices, vary with their existing positions in options, the concentration of traders in futures markets, with recent realized skewness and with the financial constraints they face. In addition increases in financial constraints amplify those effects.

5.6 Robustness: The Variance Risk Premium

Thus far, I have exclusively focused on the relative pricing of OTM puts relative to OTM calls and the extent to which it is justified by rational beliefs. While I present the cross-sectional averages of volatilities and the variance risk premium in Table 2, I have left out any discussion about what drives the absolute level of option prices and the variance risk premium. This paper has to remain largely silent on this issue because the COT data on aggregate positions is based on directional exposure (i.e. delta) only and does allow to extract volatility exposure for traders. Therefore demand-related variables cannot be used to explain time-variation in the level of volatility or its premium. However, one can make the assumption that hedgers and retail customers that are not sophisticated financial firms may generally prefer to buy options rather than writing them because short positions in options require strong risk management techniques (margin calls, short vega and short gamma exposure) while long positions do not. This intuition is certainly

confirmed by the fact that all commodities on average do exhibit a large negative variance risk premium.

[Table 8 about here.]

Albeit limited, I can still conduct some basic analysis using other variables that are not demand-based. Table 8 shows that lagged realized volatility and financial constraints have an effect of both the level of option prices (i.e. implied volatility) and the variance premium embedded in options. The dependent variables in this table are the model-free implied variance (MFIV) and the variance risk premium (as used by Jiang and Tian, 2005; Carr and Wu, 2009). Unlike the skewness measures, I do not transform the variance measures because the inference below essentially does not change between volatilities and variances.

In Table 6, broker-dealers' risk aversion (both level and changes) and the inverse asset growth of broker-dealers negatively affected implied skewness. Almost the entire magnitude of these effects carried through to the skewness risk premium. In the case of variance, I find that the proxies for financial constraints increase implied variance, i.e. make options more expensive in general, and at least partially make the variance risk premium more negative¹³. Consistent with limits to arbitrage, more severe financial constraints increase both prices and the premium on options.

Table 7 shows that lagged realized skewness has a greater positive effect on implied skewness going forward than on future realized skewness, thus leading to a negative effect on the skewness risk premium. For instance, a negative realization of skewness made puts more expensive relative to calls than is justified by next period's realized skewness. In a similar fashion, the last column of Table 8 shows that lagged realized volatility affects implied volatility more severely than realized volatility, because the VRP is more negative. A high realization of volatility in one period leads to options being more expensive in the next period and in fact too expensive when compared to future realized volatility. Again, this is consistent with arbitrageurs limiting additional volatility exposure after experiencing an adverse shock. High current volatility may lead to a mechanical trimming of exposure to that particular market, i.e. it reduces arbitrageurs' willingness; alternatively, recent losses to existing short volatility exposure may have decreased the capability of arbitrageurs to absorb additional demand.

¹³The effect for the level of effective risk aversion, $\hat{\Phi}$, is more significant in the case of volatilities than variances.

6 Returns to Marginal Liquidity Provision

The results in the previous section provide evidence that time-variation in the severity of limits to arbitrage facing speculators has predictive power for the skewness risk premium. Large current exposure in options by speculators and high concentration of long futures positions forecast a positive skewness risk premium in the future. In other words, in those instances OTM put options are overpriced and OTM calls are relatively underpriced on average. Recent realized skewness forecasts a negative skewness risk premium, i.e. after a positive realization of skewness, OTM calls tend to be overpriced relative to OTM puts.

Unfortunately, the non-parametric nature of the measures employed, not to mention the non-linear transformation to limit excess kurtosis, make it hard to judge how economically important these factors really are. Can a trader create abnormal returns by taking on exposure to skewness in commodities based on some dimension of limits to arbitrage? And second, what kind of risks does he subject himself to?

As for the second question, consider a trader that aims to exploit the predictive power of the ‘SpO’ variable. To profit, he would go long (short) skewness in commodities that, at the current time, have above-normal (below-normal) levels of speculative long exposure. Since his positions would align with those of the existing speculators, he essentially represents the marginal investor on the speculators’ side. In other words he takes on the position of the marginal liquidity provider in the commodity options market. As such, he should be particularly vulnerable to shocks to the funding constraints of other speculators. If those investors on the same side of the trade as himself are forced to liquidate, his position will suffer accordingly. Thus, the return from such a strategy can be considered compensation for liquidity provision. In a future version of this paper, I will investigate more thoroughly the extent to which any returns are compensation for risk factors related to liquidity risk. For now, I consider the standard risk factors only.

6.1 Practical Implementation

The most direct way to implement these strategy would be with long and short positions in skewness swaps. Unfortunately, the use of skewness swaps is, at least at present, largely academic as they are not offered by investment banks. Even variance swaps, which are available over the counter, will likely be available for a small number of equity indices only, not for the cross section of stocks or commodities.

It is possible to form long/short portfolios of skewness swaps synthesized from options, but it is unclear of what magnitude the bid/ask spreads would be and further, given

the extreme number of short option positions that need to be taken, what the margin requirements would be. In addition, the replication method of KNS requires continuous rebalancing in the underlying futures contract. If one chooses instead to rebalance daily, the floating leg, i.e. the realized skewness, becomes very noisy for short periods of time. After all, in the preceeding section I choose 90 day maturities precisely because shorter maturities of e.g. one month are unreliable. All these issues make it difficult to compute returns on skewness swaps in a set up of overlapping, re-balancing portfolios. Nevertheless, in unreported tests, I do form portfolios of skewness swaps, albeit without intermittent rebalancing. The signs and significances do line up with the results reported below.

Instead, I form relatively simple portfolios made of options in an attempt to catch the basic notion of skewness. Risk reversals, i.e. a long position in a OTM call option and a short position in a put option with the same absolute delta, are well-known in the realm of currency options and widely available over-the-counter. A risk reversal based on options with an absolute delta of .25 each, carries a total delta of $+.50$. To neutralize the valuation effect from a directional move in the underlying, I enter an offsetting position in the underlying futures contract. Bali and Murray (2010) use a similar construct, which they call ‘skewness asset’, but adjust the weights of the put in order to make the asset both delta and vega-neutral.

Using delta-neutral risk reversals (DNRR) as the asset of choice, I form zero net investment portfolios that are long and short an equal dollar amount of DNRRs in each commodity. The portfolios are held for either one month or half a month after which they are replaced by new DNRRs of the same commodities, but rebalanced to allocate equal dollar amounts between positions once more. This process is repeated until the underlying option series are close to expiration. To reduce standard errors of this strategy, several partially overlapping portfolios are held at each point in time, much in the way momentum portfolios of stocks are constructed (Jegadeesh and Titman, 1993).

The decision whether a given commodity receives a positive or negative weight in the portfolio is based on the relative rank after sorting the cross-section of commodities available at the time of formation according to some criterion such as ‘SpO’. The bottom third receives a negative, the top a positive and the middle a zero weight.

6.2 Portfolio Results

[Table 9 about here.]

Using the method deccribed above, I construct a time series of portfolio returns based

on the key variables that measure some aspect of limits to arbitrage. I use the maximum time period for which the sorting variable is available. Table 9 shows those results for the three sorting criteria that exhibited strong predictability of the skewness risk premium: the net exposure of speculators in options (SpO); the futures trader concentration on the long side (CRL); and lagged realized skewness. The analysis is done for 2 levels of delta, .25 and .10, as well as for monthly and bi-monthly rebalancing. This means that for monthly rebalancing there are 2 overlapping portfolios at any point in time, while there are 4 when rebalancing is done bi-monthly. For each commodity, the risk reversals are based on an expiration that lies between 2 and 4 months after first formation.

To compute returns in a more realistic fashion, assumptions about margin requirements have to be made. The rules governing margins of futures and options are complex and change constantly with market states, regulatory environments and over time more generally.¹⁴ For the case of a delta-hedged risk reversal, margins are required for the short side of the option trade and for the futures contract. I conservatively assume that a margin buffer of 10 percent of the nominal exposure in the futures contract has to be maintained. In most commodities and in most periods, this will exceed actual margin requirements¹⁵.

On the part of the option positions, I assume that the required margin is 150 percent of the sum of put and call price combined for the .25-delta risk reversal. In practice, no margin would be required on the long side of the risk reversal, but about 3 times the option price on the short side based on example calculations. For the more extreme .10-delta RRs, the option price is small relative to potential changes in value. Example calculations here suggest that about 6 times the short option value is required, which I proxy for by using 3 times the sum of both option prices as margin.¹⁶

Three patterns are discernable from Table 9. First, increasing the rebalancing frequency and, at the same time, the number of overlapping portfolios from 2 to 4 reduces standard errors. Second, using options that are further out-of-the-money (lower delta) may lead to larger returns, but increases noise in measurement due to the more extreme nature of far-OTM option returns. Third, equity-based risk factors (Fama-French factors and stock momentum) do not significantly contribute to an explanation of the returns of these portfolios.

¹⁴The CBOE provides numerous examples for margin requirements on option strategies e.g. in <http://www.cboe.com/LearnCenter/pdf/margin2-00.pdf>

¹⁵<http://www.cmegroup.com/clearing/margins/> provides current information on futures margins on the CME.

¹⁶I used <http://www.cboe.com/tradtool/mcalc/default.aspx> for the example calculations.

Specifically, I find that sorting on ‘SpO’ over time period from 1995 to 2010 yields up to 145 basis points per month in returns with a t-statistic in excess of 4 based on risk reversals with a delta of .25. The return increases to about 2 percent for the more extreme .10-delta risk reversals, while significance does not change materially. The strategy goes short (long) risk reversals of commodities with low (high) levels of SpO. A long position in a delta-hedged risk reversal proxies for a long position in skewness. Thus, the positive sign of the coefficient for SpO in forecasting the SRP translates into a significantly positive return when going long (short) skewness in commodities with high (low) SpO.

I repeat the same procedure for the long futures trader concentration. Going long (short) DNRRs in commodities with a high (low) level of trader concentration yields a positive return of between 68 and 85 basis points per month, except for the low-turnover, far-OTM case, which appears to be too noisy. The returns here are lower than expected given the large significance of the log of CRL variable in Section 5. The equal-weighting based on (linear) rank of CRL may not be able to fully account for the non-linear relationship with the skewness risk premium.

Lastly, given the negative sign for lagged realized skewness in the panel analysis in Section 5, I expect a negative return for a strategy that is long (short) DNRRs in commodities with a recently high (low) realization of skewness. Table 9 confirms this intuition. Returns reach up 97 basis points for .25-delta DNRRs and up to 131 basis points for the more extreme .10-delta DNRRs. Notably, it appears that this strategies loads positively on equity market risk. Thus, a trader that wants to profitable trade based on lagged skewness would actually have negative beta risk.

Unfortunately, the dataset underlying this analysis does not provide bid and ask quotes, only mid-quotes. Thus, it is hard to judge how much of the observed abnormal returns can be realistically attained. Note, however, that in a delta-hedged risk reversal, the proportion of capital used for the options is relatively small given the large position in the futures contract. Transactions costs in futures are generally on the order of a few basis points.

Lastly, the strategy presented here is rudimentary in the sense that the weights are equal across commodities and across time. Given the results that the effects of limits to arbitrage increase in times of financial constraints one might consider varying over time the dollar amount put at risk accordingly. One might also consider putting more weight on risk reversals of commodities for which the sorting variable ranked very high or very low to improve returns.

7 Conclusion

Asset pricing models have some difficulty to rationalize the abnormal returns that have been found for a number of strategies involving equities and equity index options. Among these, the out-of-the-money index put option puzzle has received the most attention (see, among others, Bondarenko, 2003; Liu et al., 2005; Benzoni et al., 2011). Rather than assuming unrealistically large risk aversion parameters or biases beliefs, one strand of the literature looks for alternative explanations in the microstructure of the options market, more specifically, the task of financial intermediation fulfilled by market makers and arbitrageurs in general (Bollen and Whaley, 2004; Garleanu et al., 2009).

I contribute to this literature by investigating several aspects of limits to arbitrage that affect this intermediation activity: hedging demand pressure in options, trader concentration in the underlying futures market, funding constraints of arbitrageurs and recent losses to speculators. In a large panel dataset containing 25 commodities over 20 years, I find that OTM puts become more overpriced and OTM calls more underpriced when i) speculators hold larger net long positions in the corresponding options, ii) long (short) positions in the underlying futures market are concentrated among fewer (more) traders, iii) speculators are more financially constrained, and iv) recent realized skewness was more negative while speculators were net long futures. In addition, increases in financial constraints make the effects in i) and iv) more pronounced. Lastly, to judge the economic importance of these limits to arbitrage, I form portfolios that aim to exploit the different returns to OTM calls and puts. I find these strategies to yield significant returns on the order of 1 to 2 percent per month that are unrelated to common risk factors.

As a side product, I document that commodity option prices consistently contain a large negative variance risk premium. For lack of suitable data on volatility exposure by traders, I am unable to investigate this premium more thoroughly in the context of limits to arbitrage. Lastly, my findings suggest that limits to arbitrage in the futures market spill over into the options market and impact prices and returns there. Conversely, the option market may contain information that affect futures returns. I am leaving both topics for others to explore.

References

- V. V. Acharya, L. A. Lochstoer, and T. Ramadorai. Limits to arbitrage and hedging: Evidence from commodity markets. Working Paper 16875, National Bureau of Economic Research, March 2011. URL <http://www.nber.org/papers/w16875>.
- T. Adrian and H. S. Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418 – 437, 2010.
- T. Adrian, E. Etula, and T. Muir. Financial intermediaries and the cross-section of asset returns. Staff Reports 464, Federal Reserve Bank of New York, 2011.
- G. Bakshi and N. Kapadia. Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, 16(2):pp. 527–566, 2003.
- G. Bakshi and D. Madan. Spanning and derivative-security valuation. *Journal of Financial Economics*, 55(2):205 – 238, 2000.
- G. Bakshi, N. Kapadia, and D. Madan. Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies*, 16(1):101–143, 2003.
- T. G. Bali and S. Murray. Implied risk-neutral skewness and the cross-section of option returns. Working paper, Baruch College, CUNY, 2010.
- D. S. Bates. Post-’87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1-2):181 – 238, 2000. ISSN 0304-4076.
- D. S. Bates. Empirical option pricing: a retrospection. *Journal of Econometrics*, 116(1-2):387 – 404, 2003.
- L. Benzoni, P. Collin-Dufresne, and R. S. Goldstein. Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics*, 101(3):552 – 573, 2011.
- H. Bessembinder. Systematic risk, hedging pressure, and risk premiums in futures markets. *Review of Financial Studies*, 5(4):637–667, 1992.
- H. Bessembinder and K. Chan. Time-varying risk premia and forecastable returns in futures markets. *Journal of Financial Economics*, 32(2):169 – 193, 1992.

- Z. Bodie and V. I. Rosansky. Risk and return in commodity futures. *Financial Analysts Journal*, 36(3):pp. 27–39, 1980.
- N. P. B. Bollen and R. E. Whaley. Does net buying pressure affect the shape of implied volatility functions? *The Journal of Finance*, 59(2):pp. 711–753, 2004.
- T. Bollerslev and V. Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66(6): forthcoming, December 2011.
- O. Bondarenko. Why are put options so expensive? Working paper, University of Illinois at Chicago, 2003.
- B. Boyer, T. Mitton, and K. Vorkink. Expected Idiosyncratic Skewness. *Review of Financial Studies*, 23(1):169–202, 2010.
- M. Britten-Jones and A. Neuberger. Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, 55(2):839–866, 2000.
- M. Broadie and J. Detemple. American option valuation: New bounds, approximations, and a comparison of existing methods. *The Review of Financial Studies*, 9(4):1211–1250, 1996.
- M. K. Brunnermeier, S. Nagel, and L. H. Pedersen. Carry trades and currency crashes. Working Paper 14473, National Bureau of Economic Research, November 2008. URL <http://www.nber.org/papers/w14473>.
- J. Y. Campbell and J. H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):pp. 205–251, 1999.
- J. Cao and B. Han. Cross-Section of Option Returns and Idiosyncratic Stock Volatility. McCombs Research Paper Series No. FIN-15-09, 2011.
- P. Carr and D. Madan. Optimal positioning in derivative securities. *Quantitative Finance*, 1(1):19 – 37, 2001.
- P. Carr and L. Wu. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.
- J. S. Conrad, R. F. Dittmar, and E. Ghysels. Ex Ante Skewness and Expected Stock Returns. *SSRN eLibrary*, 2009.

- F. A. De Roon, T. E. Nijman, and C. Veld. Hedging pressure effects in futures markets. *The Journal of Finance*, 55(3):1437–1456, 2000.
- P. Dennis and S. Mayhew. Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis*, 37(03):471–493, 2002.
- D. Domanski and A. Heath. Financial Investors and Commodity Markets. *BIS Quarterly Review*, March 2007, pages 53–67, 2007.
- J. Driessen, P. J. Maenhout, and G. Vilkov. The price of correlation risk: Evidence from equity options. *Journal of Finance*, 64(3):1377 – 1406, 2009.
- J.-C. Duan and J. Wei. Systematic risk and the price structure of individual equity options. *Review of Financial Studies*, 22(5):1981–2006, 2008.
- E. Etula. Broker-Dealer Risk Appetite and Commodity Returns. *SSRN eLibrary*, 2010.
- K. A. Froot. The market for catastrophe risk: a clinical examination. *Journal of Financial Economics*, 60(2-3):529 – 571, 2001.
- N. Garleanu, L. H. Pedersen, and A. M. Poteshman. Demand-based option pricing. *Review of Financial Studies*, 22(10):4259–4299, 2009.
- G. Gorton and K. G. Rouwenhorst. Facts and fantasies about commodity futures. *Financial Analysts Journal*, 62(2):pp. 47–68, 2006.
- G. B. Gorton, F. Hayashi, and K. G. Rouwenhorst. The fundamentals of commodity futures returns. Working Paper 13249, National Bureau of Economic Research, July 2007. URL <http://www.nber.org/papers/w13249>.
- A. Goyal and A. Saretto. Cross-section of option returns and volatility. *Journal of Financial Economics*, 94(2):310 – 326, 2009.
- D. Gromb and D. Vayanos. Limits of arbitrage: The state of the theory. Working Paper 15821, National Bureau of Economic Research, March 2010. URL <http://www.nber.org/papers/w15821>.
- D. Hirshleifer. Risk, futures pricing, and the organization of production in commodity markets. *Journal of Political Economy*, 96(6):pp. 1206–1220, 1988.
- N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):pp. 65–91, 1993.

- G. J. Jiang and Y. S. Tian. The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4):1305–1342, 2005.
- J. M. Keynes. *A Treatise on Money*. MacMillan, London, 1930.
- R. Kozhan, A. Neuberger, and P. G. Schneider. The Skew Risk Premium in Index Option Prices. *SSRN eLibrary*, 2011.
- M. L. Lemmon and S. X. Ni. The effects of investor sentiment on speculative trading and prices of stock and index options. Working paper, 2010.
- J. Liu, J. Pan, and T. Wang. An equilibrium model of rare-event premia and its implication for option smirks. *Review of Financial Studies*, 18(1):131–164, 2005.
- A. Neuberger. Realized Skewness. *SSRN eLibrary*, 2011.
- M. Oehmke. Gradual arbitrage. Working paper, Columbia University, 2009.
- Z. Rehman and G. Vilkov. Risk-Neutral Skewness: Return Predictability and Its Sources. *SSRN eLibrary*, 2010.
- M. Rubinstein. Implied binomial trees. *The Journal of Finance*, 49(3):pp. 771–818, 1994.
- A. Shleifer and R. W. Vishny. The limits of arbitrage. *The Journal of Finance*, 52(1):pp. 35–55, 1997.
- M. Szymanowska, F. A. De Roon, T. E. Nijman, and R. Van den Goorbergh. An Anatomy of Commodity Futures Risk Premia. *SSRN eLibrary*, 2010.
- S. B. Thompson. Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics*, 99(1):1 – 10, 2011.

A Datasets

A.1 Commodity Futures and options

The main source of data are the 'InfoTech CDs' provided by the Commodity Research Bureau (CRB) covering i) futures and cash markets and ii) futures options for a large cross-section of U.S. and international futures and commodities markets. The futures data contains OHLC prices, while the options data only provides a daily closing price. Volume and open interest data is not available for options and for futures only from the year 2001 onwards. I substitute volume and open interest data from Thomson Reuters DATASTREAM for all futures contracts in the sample, going back to 1980. The option closing price is either the price of the last trade or if no trade occurred it is a 'nominal settlement' as determined by an options pricing model.

Details on cleaning process

For a few commodities in the sample, option prices are rounded and are missing a crucial last digit. For instance, Feeder Cattle has a minimum tick size of 0.025, i.e. option prices must be multiples of this tick size, but prices are given with only 2 digits after the decimal point. In those cases, if a price ends in x.x20 or x.x30 (x.x70, x.x80) it is corrected to end in x.x25 (x.x75). Second, prices below the minimum tick size are deleted from the sample.

Further, the data is cleaned from stale option quotes. If an option price remains stale for more than 2 days, every further instance at the same price is deleted. This filter is waived for prices at or below 5 times the minimum tick size, because for far out-of-the-money options the tick size prevents frequent adjustment to movements in the underlying. Lastly, going from at-the-money to out-of-the-money for calls and puts separately, option quotes with nominal value are deleted if they follow another option with nominal value.

An additional quirk in the data is that the strike is always given as a 4-digit number regardless of actual strike price. For instance, the strike 7500 could be 0.75 cents in the case of gasoline, 7.50 cents in the case of sugar, ...or even 7500 cents in the case of silver. I identify the true strike using the following algorithm: First the closing price of the underlying futures contract is added to the raw options data. I visually inspected the data for each commodity to identify the smallest and largest strike that ever occurred over the course of the sample as well as the greatest common factor (GCF) of strikes.

For puts, I start with the highest possible, i.e. assuming a 4-digit strike price. By simple no-arbitrage for American-type options, it must hold that $K - P < F$, allowing for

some small amount of tolerance. If this inequality does not hold, the strike is divided by 10 until it holds or until the strike falls below the minimum strike price or is not a multiple of the GCF (maximum 5 iterations). For calls, I start with the smallest strike possible within the sample data, which is a number between 1 and 0.1. By no-arbitrage, it must hold that $K + C > F$. The strike is multiplied by 10 until the condition holds for the first time or the strike exceeds the maximum strike price. This algorithm identifies the only valid strike price given the price of the underlying and the option price. Additional no-arbitrage conditions as well as a maximum on implied volatility ensure that the strike-quote tuple is sensible.

Construction of discount factors

I construct discount factors as suggested by the manual to the Ivy DB US options Database, which outlines an algorithm for discount factors based on BBA LIBOR rates and CME Eurodollar futures. The CRB data on Eurodollar futures starts in 1982 which coincides with the exchange listing of that contract. BBA LIBOR data is available from Thomson Reuters DATASTREAM starting in 1986.

Eurodollar futures represent the present value of a 3-month time deposit of \$1m USD at a bank outside the U.S. starting at the expiration of the futures contract. In other words, they are forward rate agreements. They expire during the last month of each quarter and are available up to 10 years in the future. It is thus possible to construct discount rates (zero rates) up to 10 years into the future based on a strategy that sequentially rolls over 3-month bank deposits. The anchoring of these forward agreements is provided via linear interpolation of LIBOR spot rates.

Step 1: Transform BBA LIBOR rates (for $T \in 1w, 1m, 2m, \dots, 12m$) into discount factors (DF) using an actual/360 day count convention:

$$DF_T = (1 + r_T \frac{d}{360})^{-1}$$

Step 2: Transform the DF back into continuously-compounded rates using an actual/365 convention:

$$r_{c,T} = -\frac{365}{d} \log(DF_T)$$

Step 3: Linearly interpolate the two LIBOR rates surrounding the front Eurodollar futures (expiration > 7 days). Transform the interpolated rate of the front futures back into a discount factor DF_0

$$DF_0 = \exp\left(\frac{d_0}{365}r_{c,T_0}\right)$$

Step 4: The Eurodollar implied forward rates are 100 minus the settlement prices divided by 100. Compute subsequent discount factors by discounting the previous DF with the implied forward rate:

$$F_{i,i+1} = \frac{100 - ED_i}{100} \quad (12)$$

$$DF_{i,i+1} = \left(1 + F_{i,i+1} \frac{d_{i+1} - d_i}{360}\right)^{-1} \quad (13)$$

$$DF_{i+1} = DF_i \cdot DF_{i,i+1} \quad (14)$$

Step 5: Transform all Eurodollar discount factors back to continuously-compounded rates (as in step 2).

Construction of implied measures

Given a clean set of option prices and discount rates, implied volatilities are computed following the BBSR algorithm proposed by Broadie and Detemple (1996). It combines the Binomial Black Scholes (BBS) method, whereby the option prices in the penultimate nodes of the tree are replaced by the Black Scholes value, with the Richardson interpolation. In the latter, a binomial tree is constructed twice, once with N_1 nodes yielding a price C_1 and then with $N_2 = 2N_1$ nodes yielding a price C_2 . The price $C = 2C_2 - C_1$ gives a much more accurate estimate of the true price than C_2 alone, because of the oscillation property of the binomial tree estimation. For further details, see Broadie and Detemple (1996).

Given IVs of American-type options, I proceed in accordance with the literature on the computation of option-implied measures of variance and skewness. IVs are interpolated linearly on a fine grid for moneyness levels of up to 8 standard deviations around the money. The IVs are translated into European-type option prices, which are then used according the summation formulas put forth in Bakshi, Kapadia, and Madan (2003) (BKM) and Kozhan, Neuberger, and Schneider (2011) (KNS) to compute model-free implied variance, implied BKM skewness and model-free implied skewness (KNS) as described in those papers and in the technical appendix of this paper.

Realized measures

Following KNS, the realized counter-parts to MFIV and MFIS can be computed as sums of functions of daily (futures/stock) returns and option price data. KNS provide formal proof that under the risk-neutral measure the expectations of these sums converge to the implied measures in the limit. The formula to compute the realized measures are also in the technical appendix.

A.2 SPX cash settled index options

A second dataset was acquired from 'Market Data Express' covering all options on the S&P 500 cash index. I use cash options rather than futures options, because the latter was only available at quarterly expirations until 2006. The data set covers the time period 1990 to 2009. The data requires some filtering for errors which can be inferred from the documentation provided by the vendor. The computation of implied volatilities is done in the same way as for the commodity options, the only difference being that data on dividends is required. I infer those from the difference in returns between the total return and the price index for the S&P 500 as provided by Thomson Reuters DATASTREAM. Lastly, implied and realized measures are computed just as above.

B Theory of option-implied measures

B.1 Spanning Approach

Carr and Madan (2001) derive a neat way of replicating any (twice differentiable) time-T payoff function of an underlying price process by taking an initial time-0 position in the risk-free asset, the underlying asset (stock, forward, Futures) and in a continuum of European options with maturity T. Call the stochastic time-T value of underlying $S = S_T$ and today's value S_0 . The payoff is some function $H(S)$ which can be replicated as follows:

$$H(S) = [H(S_0) - H_S(S_0)S_0] + H_S(S_0)S + \int_0^{S_0} H_{SS}(K) (K - S)^+ dK + \int_{S_0}^{\infty} H_{SS}(K) (S - K)^+ dK \quad (15)$$

The derivation can be found in the appendix of Carr and Madan and is based on the fundamental theorem of calculus. The time-0 price of the payoff must then equal to the value of the replicating portfolio, i.e.

$$V_0[H(S)] = \mathbb{E}_0^{\mathbb{Q}} [e^{-r\tau} H(S_T)] = [H(S_0) - H_S(S_0)S_0] e^{-r\tau} + H_S(S_0)S_0 + \int_0^{S_0} H_{SS}(K)P_0(K)dK + \int_{S_0}^{\infty} H_{SS}(K)C_0(K)dK \quad (16)$$

Here, we used that $\mathbb{E}_0^{\mathbb{Q}} [e^{-r\tau} S_T] = S_0$. This approach has been used by, among others, Bakshi et al. (2003), Britten-Jones and Neuberger (2000), and Jiang and Tian (2005) to derive option-implied expectations for payoffs of higher moments of returns.

B.2 Computing implied moments

In an oft-cited paper, Bakshi et al. (2003) (BKM) derive non-parametric formula for option-implied, risk-neutral skewness (and kurtosis) of log returns in the following way. Using the method of Bakshi and Madan (2000), they replicate in turn the square, cubic and quartic contract of log return $r_{r,T} = \log(S_T/S_t)$. Then the risk-neutral skewness of

the log return over the period $\tau = [t, T]$ is given by

$$\text{BKMSKEW}_{t,T} = \frac{\mathbb{E}_t^{\mathbb{Q}} \left[\left(r_{t,T} - \mathbb{E}_t^{\mathbb{Q}}[r_{t,T}] \right)^3 \right]}{\mathbb{E}_t^{\mathbb{Q}} \left[\left(r_{t,T} - \mathbb{E}_t^{\mathbb{Q}}[r_{t,T}] \right)^2 \right]^{3/2}} \quad (17)$$

$$= \frac{e^{r\tau} W_{t,T} - 3\mu_{t,T} e^{r\tau} V_{t,T} + 2\mu_{t,T}^3}{[e^{r\tau} V_{t,T} - \mu_{t,T}^2]^{3/2}} \quad (18)$$

where

$$\mu_{t,T} = \mathbb{E}_t^{\mathbb{Q}} \left[\log\left(\frac{S_T}{S_t}\right) \right] \quad (19)$$

$$= e^{r\tau} - 1 - \frac{e^{r\tau}}{2} V_{t,T} - \frac{e^{r\tau}}{6} W_{t,T} - \frac{e^{r\tau}}{24} X_{t,T} \quad (20)$$

and V, W, X are the time- t prices of the replicating portfolios of the square, cubic and quartic contract respectively. Note that the relationship for μ is based on a Taylor expansion and thus the resulting expressions are not exact.

B.3 Aggregation Property

The existence of risk premia for higher moments of stock market returns has received considerable attention in recent years. One tool that has proven very useful in evaluating the presence of one such premium, the variance risk premium, is the variance swap. For recent applications, see e.g. Carr and Wu (2009). The basic idea of the variance swap is that the buyer of a variance swap pays some fixed amount that represents today's expectation of future variance of the period return and then receives the actually realized variance as measured by returns of a higher frequency

Neuberger (2011) points out that the definitions of the variance swap used in practice are not fully consistent as the expectation of the period return variance do not equal the variance of daily returns. He derives the 'Aggregation Property' (AP) which ensures that the risk-neutral expectation of a function g of the period return is equal to the sum of functions g of returns at a higher frequency. Denote S_t as the underlying price process, $s_t = \log S_t$, $\delta S_t = S_t - S_{t-1}$ as the price change and $\delta s_t = \log(S_t/S_{t-1})$ as the log return. Following Neuberger (2011), if (g, X) has the Aggregation Property (where X could be a

price process or a log price process) then

$$\mathbb{E}_t^{\mathbb{Q}} [g(X_T - X_0)] = \mathbb{E}_t^{\mathbb{Q}} \left[\sum^T g(\delta X) \right] \quad (21)$$

for any partition of the period $[0, T]$. The left side of the equation is called the implied characteristic and is written in terms of a function of the price process over the whole period, while the right hand side is called the realized characteristic and can be computed using price changes or returns of a higher frequency. This definition can be extended to not only cover one-dimensional price processes, but also a tuple (X, v) where v is a so-called generalized variance of the process X .

B.4 Variance Swap

Variance swaps can be defined in a number of ways. The version commonly used in practice is based on log returns, i.e. $g(s_T - s_0) = (s_T - s_0)^2$. The implied variance can be replicated using the square contract of the log return as in BKM. While this definition satisfies the definition of a generalized variance, it does not have the aggregation property, which means that

$$\mathbb{E}_0^{\mathbb{Q}} \left[\log \left(\frac{S_T}{S_0} \right)^2 \right] = \mathbb{E}_0^{\mathbb{Q}} \left[2 \sum^T \log \frac{S_t}{S_{t-1}} \right] \quad (22)$$

is not always exact. As pointed out in Jiang and Tian (2005), this relationship only holds for fully continuous processes without jumps. Neuberger (2011) proposes an alternative functional form, which has both the property of generalized variance and the AP. Using $g^L(s) = 2(e^s - 1 - s)$, it holds (even in the presence of jumps) that

$$v_{0,T}^L = \mathbb{E}_0^{\mathbb{Q}} \left[L \left(\log \left(\frac{S_T}{S_0} \right) \right) \right] \quad (23)$$

$$= 2\mathbb{E}_0^{\mathbb{Q}} \left[\frac{S_T}{S_0} - 1 - \log \left(\frac{S_T}{S_0} \right) \right] \quad (24)$$

$$= 2\mathbb{E}_0^{\mathbb{Q}} \left[\sum^T \left[\frac{S_t}{S_{t-1}} - 1 - \log \left(\frac{S_t}{S_{t-1}} \right) \right] \right] \quad (25)$$

Using the spanning approach, the implied characteristic can be replicated using weights $H_{SS}^L(K) = \frac{2}{K^2}$ for the option contracts. This variance measure is identical to the so-called model-free implied variance (MFIV) used in Britten-Jones and Neuberger (2000). The

floating leg of the variance swap is different from the common definition, but in practice the two measures of realized variance are very close. A second alternative proposed by Neuberger (2011) can be obtained using $g^E(x) = 2(s \cdot e^s - e^s + 1)$. The replicating weights are $H_{SS}^E(K) = \frac{2}{S_0 K}$. Neuberger (2011) calls the first the implied variance of the log contract and the second the implied variance of the entropy contract.

B.5 Skewness Swap

Neuberger (2011) also proposes an analogue measure for the skewness of the log return over the period $[0, T]$ using the 2 previously defined measures of variance. Note that

$$v_{0,T}^L = -2\mathbb{E}_0^{\mathbb{Q}} \left[\log \left(\frac{S_T}{S_0} \right) \right] \quad (26)$$

and similarly, one can show that

$$v_{0,T}^E = 2\mathbb{E}_0^{\mathbb{Q}} \left[\left(\frac{S_T}{S_0} \right) \log \left(\frac{S_T}{S_0} \right) \right] \quad (27)$$

Defining $g^Q(s, v^E) = 3v^E(e^s - 1) + 6(se^s - 2e^s + s + 2)$, g^Q has the aggregation property. The implied skewness is

$$\text{MFIS}_{0,T} = \mathbb{E}_t^{\mathbb{Q}} [g^Q(s_T - s_0, v^E(s_T - s_0))] \quad (28)$$

$$= 6\mathbb{E}_t^{\mathbb{Q}} \left[\left(\frac{S_T}{S_0} + 1 \right) \log \frac{S_T}{S_0} \right] \quad (29)$$

$$= 3(v_{0,T}^E - v_{0,T}^L) \quad (30)$$

The option replicating weights are $H_{SS}^Q(K) = \frac{2(K-S_0)}{S_0 K^2}$ using the spanning approach. Realized skewness over the period $[0, T]$ can be computed exactly as

$$\sum^T g^Q(\delta s, \delta v^E) = \sum^T (3\delta v^E e^{\delta s} + 6(\delta s(e^{\delta s} + 1) - 2(e^{\delta s} - 1))) \quad (31)$$

The second term can be shown to approximate cubed returns, i.e. re-writing returns as $e^{\delta s} - 1 = r_{\delta t}$,

$$6(\log(1 + r_{\delta t})(r_{\delta t} + 2) - 2r_{\delta t}) = r_{\delta t}^3 + O(r_{\delta t}^4) \quad (32)$$

This is closely related to the commonly used definition of skewness as the average skewness of daily returns. As the frequency of realized returns is increased the second term tends to become smaller and ultimately vanishes as long as the underlying process is reasonably close to a continuous diffusion process. It turns out that the first term is far more important. Using slightly different notation, Kozhan et al. (2011) write implied skewness under the assumption of continuous rebalancing as

$$\text{MFIS}_{0,T} = 3\mathbb{E}^{\mathbb{Q}} \left[\int_0^T dv_{t,T}^E \left(\frac{dS_t}{S_t} \right) \right], \quad (33)$$

which emphasizes the fact that the skewness of the return over a longer horizon equals the covariation of returns and changes in expected future variance over the period. Using a Taylor series expansion, it follows that

$$\text{MFIS}_{0,T} = 3(v_{0,T}^E - v_{0,T}^L) = \mathbb{E}^{\mathbb{Q}} [r_{0,T}^3 + O(r_{0,T}^4)] \quad (34)$$

In practice, computing realized skewness requires that the entropy contract is traded or, equivalently, that its price can be constructed from a range of options on the underlying.

Table 1: SAMPLE OVERVIEW

Commodity	Futures Options Data				COT data			
	First Exp.	Last Exp.	#Exp. ≥ 30	w/TTM ≥ 90	Concentration Long	Short	Net Exposure OHP	SpO
Agricultural								
Soybean Oil	1989/05	2010/09	152	139	22%	37%	-24%	-25%
Corn	1989/05	2010/09	180	108	21%	17%	-16%	-14%
Oats	1990/12	2010/09	84	78	34%	59%	-21%	-23%
Rough Rice	1992/09	2010/09	105	108	42%	49%	3%	-5%
Soybeans	1989/05	2010/09	201	150	19%	23%	-2%	-5%
Soybean Meal	1989/05	2010/09	186	163	22%	35%	-5%	-12%
Wheat (CBOT)	1989/05	2010/09	165	107	29%	28%	9%	0%
Energy								
Crude Oil (WTI)	1989/04	2010/09	258	248	21%	19%	6%	9%
Heating Oil No. 2	1989/04	2010/09	258	245	24%	25%	1%	-10%
Unl. Gasoline	1989/07	2006/12	209	192	32%	27%	7%	3%
Natural Gas	1992/12	2010/09	214	212	18%	26%	6%	1%
RBOB Gasoline	2007/02	2010/09	44	42	22%	23%	3%	1%
Meat								
Feeder Cattle	1987/03	2010/08	186	186	33%	20%	21%	35%
Live Cattle	1991/04	2010/08	122	117	31%	24%	32%	38%
Lean Hogs	1997/02	2010/08	104	104	38%	32%	37%	27%
Pork Bellies	1987/02	2008/02	103	79	31%	30%	5%	13%
Metal								
Gold (NYMEX)	1989/04	2010/09	244	128	29%	38%	12%	7%
Copper (HG)	1990/03	2010/09	212	194	33%	43%	14%	-10%
Silver (NYMEX)	2001/02	2010/09	116	49	26%	51%	18%	1%
Soft								
Cocoa	1990/05	2010/09	208	150	39%	42%	-16%	-15%
Cotton No. 2	1990/05	2010/09	116	112	27%	35%	14%	8%
Orange Juice	1990/05	2010/09	160	143	34%	52%	3%	-2%
Coffee C	1990/05	2010/09	209	162	22%	32%	-2%	-4%
Lumber	1987/09	2010/09	137	117	32%	36%	37%	12%
Sugar No. 11	1990/05	2010/09	221	160	25%	37%	12%	6%

This table lists all U.S.-exchange listed commodities in the sample. The table contains information on the first and last expiration month available and the number of option series which have sufficient data allowing the computation of implied and realized measures up to a time to maturity of at least 30 (90) days. The last 4 columns list the time series averages of four COT variables: the long and short futures concentration ratios, i.e. the share of total interest held by the largest 8 traders on either side of the futures market; the short net exposure of hedgers in the options market (delta-weighted and scaled by hedgers' total open interest) and the long net exposure of speculators (delta-weighted and scaled by speculators' total open interest).

Table 2: VOLATILITY AND VARIANCE RISK PREMIUM

Commodity	# Obs.	Implied Vol.	Realized Vol.	Variance avg.	Premium t-stat
Equity Market					
S&P 500	105	22.0%	18.3%	-0.96	[-1.07]
S&P 500 (30 days)	238	19.8%	16.2%	-1.11	***[-2.99]
Agricultural					
Soybean Oil	130	25.3%	22.6%	-1.23	***[-3.44]
Corn	107	26.1%	21.7%	-2.07	***[-6.50]
Oats	72	32.9%	29.1%	-2.01	***[-2.80]
Rough Rice	97	28.0%	22.6%	-2.92	***[-7.33]
Soybeans	149	25.3%	21.8%	-1.63	***[-5.21]
Soybean Meal	160	25.4%	23.8%	-0.87	*[-1.87]
Wheat (CBOT)	106	26.5%	24.7%	-0.86	**[-2.57]
Energy					
Crude Oil (WTI)	244	34.3%	31.0%	-2.00	**[-2.34]
Heating Oil No. 2	236	33.3%	30.6%	-1.88	***[-3.73]
Unl. Gasoline	184	31.4%	29.6%	-1.08	**[-2.58]
Natural Gas	205	47.3%	44.8%	-2.16	*[-1.92]
RBOB Gasoline	41	42.3%	37.9%	-2.93	[-0.86]
Meat					
Feeder Cattle	185	14.4%	11.4%	-0.82	***[-8.20]
Live Cattle	117	15.5%	13.2%	-0.73	***[-5.18]
Lean Hogs	103	25.9%	22.9%	-1.35	***[-2.94]
Pork Bellies	74	39.7%	33.4%	-4.47	***[-6.85]
Metal					
Gold (NYMEX)	126	17.9%	14.4%	-1.17	***[-4.49]
Copper (HG)	187	30.1%	26.3%	-2.26	***[-2.73]
Silver (NYMEX)	47	32.5%	29.4%	-1.25	[-1.00]
Soft					
Cocoa	147	34.2%	29.7%	-2.78	***[-6.54]
Cotton No. 2	110	25.1%	23.9%	-0.53	[-1.52]
Orange Juice	138	34.4%	29.1%	-3.74	***[-4.58]
Coffee C	161	40.5%	33.1%	-5.14	***[-4.99]
Lumber	102	30.4%	27.6%	-1.76	***[-3.77]
Sugar No. 11	159	33.5%	30.2%	-2.18	***[-4.00]

This table depicts time series averages of annualized option-implied and realized volatility for the sample of commodities as well as the S&P 500 Equity Index. Implied volatility is the square root of the MFIV measures (as in Jiang and Tian, 2005; Britten-Jones and Neuberger, 2000) and realized volatility is the square root of realized variance (RV), i.e. the sum of daily square returns of the underlying futures contract. The last two columns show the sample estimates and the Newey-West adjusted t-statistics of the variance risk premium (VRP), defined as the difference between RV and MFIV (annualized and multiplied by a factor of 100). All measures are derived from options with a remaining maturity of 90 to 99 days.

Table 3: IMPLIED VS. REALIZED SKEWNESS

Commodity	# Obs.	Implied BKM	Implied MFIS	Realized r^3	Realized RSkew	Skewness Premium avg.	Premium t-stat
Equity Market							
S&P 500	105	-1.91	-1.24	-0.01	-1.11	0.13	[0.37]
S&P 500 (30 days)	238	-1.93	-0.60	0.00	-0.36	0.25	***[3.31]
Agricultural							
Soybean Oil	130	0.29	0.50	0.01	0.30	-0.20	***[-3.12]
Corn	107	0.31	0.76	-0.01	0.47	-0.29	***[-4.59]
Oats	72	0.41	1.30	-0.02	0.30	-1.00	***[-5.07]
Rough Rice	97	0.18	0.69	0.00	0.42	-0.28	**[-2.21]
Soybeans	149	0.48	0.78	-0.01	0.55	-0.23	***[-3.27]
Soybean Meal	160	0.28	0.01	-0.01	0.40	0.38	[0.80]
Wheat (CBOT)	106	0.13	0.57	0.00	0.58	0.01	[0.08]
Energy							
Crude Oil (WTI)	244	-0.49	-0.32	-0.16	0.02	0.33	[0.89]
Heating Oil No. 2	236	-0.09	0.47	-0.02	0.09	-0.38	***[-3.55]
Unl. Gasoline	184	-0.17	0.25	-0.05	0.18	-0.07	[-0.49]
Natural Gas	205	-0.05	2.05	0.12	2.99	0.94	*[1.94]
RBOB Gasoline	41	-0.11	0.16	-0.11	-0.76	-0.92	**[-2.14]
Meat							
Feeder Cattle	185	-1.73	-0.21	0.00	-0.04	0.18	***[3.58]
Live Cattle	117	-1.46	-0.23	0.00	-0.05	0.18	***[3.42]
Lean Hogs	103	-0.95	-0.65	-0.02	-0.28	0.37	***[3.70]
Pork Bellies	74	-0.47	-1.03	0.00	-0.25	0.78	**[2.38]
Metal							
Gold (NYMEX)	126	-0.18	0.30	0.00	0.13	-0.18	***[-3.29]
Copper (HG)	187	-0.30	-0.11	-0.01	-0.05	0.06	[0.29]
Silver (NYMEX)	47	0.59	1.54	-0.19	0.06	-1.48	***[-5.64]
Soft							
Cocoa	147	0.01	0.52	-0.01	0.28	-0.25	*[-1.82]
Cotton No. 2	110	-0.25	0.14	-0.01	0.02	-0.12	*[-1.85]
Orange Juice	138	0.22	3.17	0.04	0.40	-2.78	*[-1.95]
Coffee C	161	0.61	4.30	0.08	1.83	-2.47	***[-5.28]
Lumber	102	-0.26	0.05	0.01	0.11	0.06	[1.25]
Sugar No. 11	159	-0.18	0.36	-0.02	-0.01	-0.37	*[-1.94]

This table depicts time series averages of a number of measures of implied and realized skewness: ‘BKM’ skewness is the unit-free, normalized skewness of log returns computed as in Bakshi, Kapadia, and Madan (2003). ‘MFIS’ is the model-free implied skewness and ‘RSkew’ is the realized skewness, both computed following Kozhan, Neuberger, and Schneider (2011), ‘ r^3 ’ is the part of RSkew that consists of the sum of cubed returns only. The table further shows the sample estimates and the Newey-West adjusted t-statistics of the skewness risk premium (SRP), defined as the difference between RSkew and MFIS. All estimates except BKM skewness are annualized and multiplied by 100. Measures are derived from options with a remaining maturity of 90 to 99 days.

Table 4: DEMAND PRESSURE AND SKEWNESS

Panel A: Implied skewness							
HP	0.11						
	[0.69]						
Sp		0.06					
		[0.63]					
OHP			-0.23		-0.05		
			**[-2.44]		[-0.41]		
SpO				-0.38	-0.35		-0.47
				***[-3.74]	**[-2.54]		***[-4.99]
SpO+						-0.57	
						***[-3.22]	
SpO−						-0.24	
						*[-1.77]	
SpO ²							-0.43
							**[-2.55]
<i>R</i> ²	12.8%	12.8%	11.6%	12.5%	12.5%	12.7%	13.3%
Panel B: Skewness risk premium							
HP	-0.11						
	[-1.00]						
Sp		-0.04					
		[-0.51]					
OHP			0.19		-0.04		
			***[2.85]		[-0.65]		
SpO				0.40	0.42		0.45
				***[3.59]	***[3.38]		***[3.88]
SpO+						0.43	
						***[2.98]	
SpO−						0.37	
						***[3.57]	
SpO ²							0.24
							***[4.24]
<i>R</i> ²	2.9%	2.8%	4.5%	5.6%	5.6%	5.6%	5.8%
# comm.	25	25	24	24	24	24	24
# months	257	257	183	183	183	183	183
Total # Obs.	3,563	3,562	2,768	2,771	2,767	2,771	2,771

This table shows the results from a panel regression of implied skewness and the skewness risk premium on a number of variables related to demand pressure. The dependent variable in Panel A is $MFIS^{1/3}$, the signed third root of average implied skewness (as in Neuberger, 2011), based on option quotes with a remaining maturity of between 90 and 99 days. In panel B, the dependent variable is the difference between $RSkew^{1/3}$ and $MFIS^{1/3}$ averaged over the same range of maturities, where $RSkew^{1/3}$ is the signed third root of realized skewness (as in Neuberger, 2011). HP (Sp) is the scaled net short (long) exposure of hedgers (speculators) in futures. OHP (SpO) is the scaled net short (long) exposure of hedgers (speculators) in options, delta-weighted. SpO+ (SpO−) is SpO conditional on it being positive (negative). Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Table 5: TRADER CONCENTRATION AND SKEWNESS

Panel A: Implied Skewness						
logCRL	-0.21 **[-2.52]		-0.22 **[-2.73]		-0.21 **[-2.37]	
logCRS		-0.02 [-0.18]	0.02 [0.23]		-0.05 [-0.45]	
logCRL+				-0.17 *[-1.96]		
logCRL-				-0.28 **[-2.31]		
dCRL(3m)					-0.36 [-1.69]	0.02 [0.05]
dCRS(3m)					0.64 **[2.40]	0.70 **[2.43]
R²	11.1%	10.2%	11.2%	11.2%	10.8%	11.7%
Panel B: Skewness Risk Premium						
logCRL	0.18 ***[3.33]		0.18 ***[3.41]		0.16 **[2.75]	
logCRS		0.08 [0.99]	0.04 [0.57]		0.11 [1.07]	
logCRL+				0.20 ***[2.83]		
logCRL-				0.14 [1.28]		
dCRL(3m)					0.39 *[1.73]	0.10 [0.41]
dCRS(3m)					-0.47 *[-2.02]	-0.61 *[-1.92]
R²	4.6%	4.1%	4.6%	4.6%	4.4%	4.9%
# comm.	24	24	24	24	24	24
# months	183	183	183	183	183	183
Total # Obs.	2,637	2,631	2,631	2,637	2,626	2,626

This table shows the results from a panel regression of implied skewness and the skewness risk premium on a number of variables related to the concentration of traders on the long and short side in the underlying futures market. The dependent variable in Panel A is $\text{MFIS}^{1/3}$, the signed third root of average implied skewness (as in Neuberger, 2011), based on option quotes with a remaining maturity of between 90 and 99 days. In panel B, the dependent variable is the difference between $\text{RSkew}^{1/3}$ and $\text{MFIS}^{1/3}$ averaged over the same range of maturities, where $\text{RSkew}^{1/3}$ is the signed third root of realized skewness (as in Neuberger, 2011). $\log\text{CRL}$ ($\log\text{CRS}$) is the log of the concentration ratio of traders being long (short) in the futures market. $\log\text{CRL+}$ ($\log\text{CRL-}$) is $\log\text{CRL}$ conditional on speculators being net long (short) in futures. dCRL (dCRS) is the net change in the long (short) concentration ratio relative to 3 months earlier. Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Table 6: FINANCIAL CONSTRAINTS AND SKEWNESS

Panel A: Implied Skewness							
SpO	-0.39 ***[-3.91]	-0.37 ***[-3.51]	-0.37 ***[-3.42]	-0.39 ***[-3.90]	-0.35 ***[-3.35]	-0.34 ***[-3.17]	-0.36 ***[-3.48]
$\hat{\Phi}$	-0.13 *[-1.90]			-0.13 *[-1.91]			
$d\hat{\Phi}$		-0.11 *[-1.96]			-0.11 *[-1.89]		
Rel. AG			-0.48 *[-1.86]			-0.46 *[-1.85]	
RVol(3m)							-0.01 [-0.02]
X-term w/ SpO				-0.01 [-0.10]	-0.34 **[-2.74]	-1.24 **[-2.47]	-2.04 **[-2.37]
R^2	12.3%	12.3%	12.2%	12.3%	12.8%	12.5%	11.7%
Panel B: Skewness Risk Premium							
SpO	0.40 ***[3.67]	0.39 ***[3.45]	0.38 ***[3.36]	0.40 ***[3.68]	0.37 ***[3.63]	0.35 ***[3.35]	0.40 ***[3.68]
$\hat{\Phi}$	0.10 **[2.28]			0.10 **[2.33]			
$d\hat{\Phi}$		0.09 **[2.42]			0.08 **[2.29]		
Rel. AG			0.45 **[2.69]			0.43 **[2.67]	
RVol(3m)							-0.40 **[-2.58]
X-term w/ SpO				0.04 [0.33]	0.28 **[2.24]	1.28 **[2.22]	0.99 [0.96]
R^2	5.1%	5.1%	5.3%	5.1%	5.4%	5.6%	5.0%
# comm.	24	24	24	24	24	24	24
# months	183	183	183	183	183	183	183
Total # Obs.	2,637	2,637	2,637	2,637	2,637	2,637	2,637

This table shows the results from a panel regression of implied skewness and the skewness risk premium on a number of variables and interaction terms related to the degree of financial constraints of arbitrageurs. The dependent variable in Panel A is $MFIS^{1/3}$, the signed third root of average implied skewness (as in Neuberger, 2011), based on option quotes with a remaining maturity of between 90 and 99 days. In panel B, the dependent variable is the difference between $RSkew^{1/3}$ and $MFIS^{1/3}$ averaged over the same range of maturities, where $RSkew^{1/3}$ is the signed third root of realized skewness (as in Neuberger, 2011). SpO is the scaled, delta-weighted net long exposure of speculators in options. $\hat{\Phi}$ is the most recent quarterly value for the effective risk-aversion of broker-dealers as in Etula (2010), $d\hat{\Phi}$ is its year-over-year change. Rel. AG is the year-over-year asset growth in balance sheets of households relative to that of broker-dealers as in Acharya et al. (2011). RVol is the realized futures volatility over the last 3 months. X-term is the interaction term between SpO and the other variable included in the same column. Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Table 7: LAGGED REALIZED SKEWNESS

Panel A: Implied Skewness						
$\text{Rskew}_{t-m}^{1/3}$	0.26 ***[4.90]	0.26 ***[4.80]	0.24 ***[4.67]	0.24 ***[4.46]	0.27 ***[5.13]	
$\hat{\Phi}$		-0.12 *[-1.83]				
$d\hat{\Phi}$			-0.11 *[-1.97]			
Rel. AG				-0.46 *[-1.83]		
Sp_{t-m}					-0.03 [-0.34]	
X-term w/ $\text{Rskew}_{t-m}^{1/3}$		0.13 *[1.93]	0.12 **[2.47]	0.47 **[2.51]	0.09 [0.58]	
$\text{Rskew}_{t-m}^{1/3} (\text{Sp}_{t-m} > 0)$						0.33 ***[4.33]
$\text{Rskew}_{t-m}^{1/3} (\text{Sp}_{t-m} < 0)$						0.20 ***[3.53]
R^2	11.1%	12.8%	12.8%	12.6%	11.2%	11.5%
Panel B: Skewness Risk Premium						
$\text{Rskew}_{t-m}^{1/3}$	-0.10 ***[-2.99]	-0.10 ***[-3.10]	-0.09 **[-2.42]	-0.08 **[-2.14]	-0.12 ***[-3.52]	
$\hat{\Phi}$		0.09 **[2.26]				
$d\hat{\Phi}$			0.09 **[2.24]			
Rel. AG				0.43 **[2.54]		
Sp_{t-m}					0.11 [1.25]	
X-term w/ $\text{Rskew}_{t-m}^{1/3}$		-0.05 [-0.68]	-0.10 *[-1.91]	-0.40 [-1.60]	-0.25 *[-1.83]	
$\text{Rskew}_{t-m}^{1/3} (\text{Sp}_{t-m} > 0)$						-0.20 ***[-3.71]
$\text{Rskew}_{t-m}^{1/3} (\text{Sp}_{t-m} < 0)$						-0.02 [-0.38]
R^2	4.6%	5.2%	5.4%	5.5%	5.2%	5.3%
# comm.	24	24	24	24	24	24
# months	183	183	183	183	183	183
Total # Obs.	2,637	2,637	2,637	2,637	2,637	2,637

This table shows the results from a panel regression of implied skewness and the skewness risk premium on lagged realized skewness and interaction terms related to the degree of financial constraints of arbitrageurs. The dependent variable in Panel A is $\text{MFIS}^{1/3}$, the signed third root of average implied skewness (as in Neuberger, 2011), based on option quotes with a remaining maturity of between 90 and 99 days. In panel B, the dependent variable is the difference between $\text{RSkew}^{1/3}$ and $\text{MFIS}^{1/3}$ averaged over the same range of maturities, where $\text{RSkew}^{1/3}$ is the signed third root of realized skewness (as in Neuberger, 2011). $\text{Rskew}_{t-m}^{1/3}$ is the most recent lagged value of $\text{RSkew}^{1/3}$. The next three variables are as in the previous table. Sp_{t-m} is the mean speculator net exposure in futures during the time that $\text{Rskew}_{t-m}^{1/3}$ was measured. X-term is the interaction term between $\text{Rskew}_{t-m}^{1/3}$ and the other variable included in the same column. The last two variables are $\text{Rskew}_{t-m}^{1/3}$ conditioned on the sign of Sp_{t-m} . Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Table 8: FINANCIAL CONSTRAINTS AND VARIANCE


Panel A: Implied Variance					
Conv. Yield(6m)	-2.56	-2.49	-1.51	-1.53	-3.43
	[-1.67]	[-1.56]	[-1.04]	[-1.07]	*[-2.04]
$\Delta\text{OI}(6\text{m})$	-1.10	-1.09	-1.00	-0.91	1.10
	[-1.20]	[-1.18]	[-1.15]	[-1.08]	[1.58]
Ret(3m)	0.33	0.36	-0.61	-0.80	1.66
	[0.14]	[0.16]	[-0.26]	[-0.35]	[1.11]
$\hat{\Phi}$		-0.38			
		[-0.53]			
$d\hat{\Phi}$			2.68		
			***[4.35]		
Rel. AG				11.37	
				***[4.55]	
RVol(3m)					41.88
					***[9.12]
R^2	0.6%	0.7%	4.5%	5.1%	39.0%
Panel B: Variance Risk Premium					
Conv. Yield (6m)	3.06	3.23	2.57	2.61	3.38
	***[3.77]	***[4.09]	***[3.08]	***[3.20]	***[3.45]
$\Delta\text{OI}(6\text{m})$	-0.67	-0.65	-0.81	-0.80	-1.49
	[-1.51]	[-1.50]	*[-1.75]	*[-1.74]	***[-3.15]
Ret(3m)	-2.05	-1.91	-1.55	-1.54	-2.54
	*[-1.83]	*[-1.75]	[-1.41]	[-1.39]	**[-2.21]
$\hat{\Phi}$		-1.08			
		*[-1.96]			
$d\hat{\Phi}$			-1.48		
			***[-3.11]		
Rel. AG				-4.98	
				**[-2.79]	
RVol(3m)					-15.60
					***[-7.24]
R^2	0.9%	1.7%	2.7%	2.2%	8.7%
# comm.	25	25	25	25	25
# months	257	254	242	254	257
Total # Obs.	3,436	3,425	3,360	3,425	3,436

This table shows the results from a panel regression of implied variance and the variance risk premium on variables related to the degree of financial constraints of arbitrageurs. The dependent variable in Panel A is MFIV, the average implied variance (as in e.g. Jiang and Tian, 2005), based on option quotes with a remaining maturity of between 90 and 99 days. In panel B, the dependent variable is the difference between realized variance RV and MFIV averaged over the same range of maturities. $\hat{\Phi}$ is the most recent quarterly value for the effective risk-aversion of broker-dealers as in Etula (2010), $d\hat{\Phi}$ is its year-over-year change. Rel. AG is the year-over-year asset growth in balance sheets of households relative to that of broker-dealers as in Acharya et al. (2011). RVol is the realized futures volatility over the last 3 months. Control variables are the convenience yield (implied by 6-month futures rel. to front month), the lagged growth in open interest and the lagged 3-month return. Standard errors are clustered by month and commodity following Thompson (2011).


Table 9: PORTFOLIO RETURNS

Sorting Criterion	α	β	SMB	HML	UMD	R^2	α	β	SMB	HML	UMD	R^2	N
	2 Monthly Turnovers, Delta = .25						2 Monthly Turnovers, Delta = .10						
SpO	1.42% ***[4.06]					0.00%	2.02% ***[3.76]					0.00%	183
	1.29% ***[3.58]	0.14 *[1.68]	0.03 [0.29]	0.12 [1.17]	0.03 [0.43]	2.36% [0.43]	1.84% ***[3.32]	0.18 [1.38]	0.01 [0.05]	0.17 [1.03]	0.06 [0.60]	1.48% [0.60]	183
CRL	0.85% ***[2.61]					0.00%	0.03% [0.05]					0.00%	247
	0.78% **[2.31]	0.02 [0.24]	0.07 [0.62]	-0.06 [-0.55]	0.10 [1.45]	1.53% [1.45]	-0.22% [-0.31]	0.26 [1.50]	-0.01 [-0.03]	0.05 [0.24]	0.17 [1.17]	1.09% [1.17]	247
lag. Rskew	-0.39% [-1.06]					0.00%	-0.51% [-0.74]					0.00%	247
	-0.64% *[-1.68]	0.17 *[1.81]	-0.07 [-0.55]	0.15 [1.22]	0.18 **[2.29]	2.72% [1.22]	-0.60% [-0.85]	-0.06 [-0.35]	0.15 [0.67]	0.24 [1.08]	0.03 [0.21]	0.67% [0.21]	247
	4 Bi-Monthly Turnovers, Delta = .25						4 Bi-Monthly Turnovers, Delta = .10						
SpO	1.41% ***[4.86]					0.00%	2.15% ***[4.80]					0.00%	367
	1.45% ***[4.91]	0.03 [0.75]	-0.08 [-1.12]	-0.05 [-0.76]	-0.03 [-0.81]	0.89% [-0.81]	2.21% ***[4.85]	0.00 [-0.02]	-0.13 [-1.28]	-0.07 [-0.69]	-0.02 [-0.37]	0.55% [-0.37]	367
CRL	0.68% ***[2.69]					0.00%	0.71% *[1.87]					0.00%	495
	0.69% ***[2.67]	0.03 [0.63]	0.01 [0.14]	-0.12 *[-1.88]	0.03 [0.78]	1.17% [0.78]	0.68% *[1.78]	0.05 [0.87]	0.03 [0.27]	-0.13 [-1.39]	0.06 [0.97]	0.96% [0.97]	495
lag. Rskew	-0.85% ***[-2.97]					0.00%	-1.08% **[-2.33]					0.00%	495
	-0.97% ***[-3.33]	0.09 *[1.92]	0.01 [0.19]	0.14 **[2.00]	0.04 [0.87]	1.32% [0.87]	-1.31% ***[-2.80]	0.24 ***[3.13]	-0.04 [-0.34]	0.24 **[2.11]	0.07 [0.94]	2.48% [0.94]	495

This table shows raw and risk-adjusted returns to strategies of going long/short delta-neutral risk reversals (DNRR) in commodity options based on one of three criteria: SpO, CRL or lagged realized skewness. Overlapping portfolios are held for 2 months and rebalanced either monthly or bi-monthly. The risk reversals are based on options with a delta of .25 or .10, respectively, and delta-hedged at each rebalancing date. All percentage returns reported are scaled to a monthly horizon.



(a) Exchange-traded commodity derivatives



(b) OTC commodity derivatives

Figure 1: Comparison of exchange-traded open interest and OTC notional amounts for commodity derivatives. The left figure shows the aggregate open interest (in \$B) of the 25 exchange-traded commodities in my sample. The data is derived from the CFTC commitment of trader (COT) reports. For each commodity in the sample, the open interest in number of contracts is multiplied by the price of the Futures front contract and the contract multiplier. The difference between open interest reported in Futures-only data and the combined report constitutes the open interest in the options market in delta-weighted Futures equivalents. The right figure shows the aggregate notional amounts outstanding (in \$T) of OTC commodity derivatives excluding Gold, separated into OTC Forwards & Swaps and OTC Options. The data is taken directly from the 'Semiannual OTC derivatives statistics' of the Bank of International Settlements.




Figure 2: Scatter plots of skewness measures vs. traders' net options exposure. The X-axis for the top two figures shows the sample period average of the relative net short exposure of hedgers (as a proportion of their total open interest) via options (OHP) for each of 25 commodities, while the X-axis for the bottom figures shows the sample period average of the relative net long exposure of speculators (as a proportion of their total open interest) via options (SpO). Against these measures of exposure, the figures on the left plot the time series average of implied skewness (as defined by Bakshi et al. (2003)), while the figures on the right plot the (t-statistics of the) skewness risk premium as defined by Kozhan et al. (2011). Both measures are based on options data with a remaining maturity of 90 to 99 days.