Two Essays in Empirical Asset Pricing

by

Thomas Ruf

Diplom, Universität Ulm, 2005
M.Sc., University of Southern California, 2005

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate Studies

(Business Administration)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

August 2012

Abstract

In the first essay, I empirically investigate the effect of financial frictions and exogenous demand pressure on both prices and returns of options. Historically, observed option returns have been a challenge for no-arbitrage asset pricing models, most notably in the case of out-of-the-money equity index puts. I propose that liquidation risk, defined as the possibility of forced selling of speculative positions following a liquidity shock, is a major driver of the relative price of out-of-the-money put vs. call options (option-implied skewness) in commodity futures options markets and gives rise to a skewness risk premium in option returns. Establishing speculative net long positions in options (OSP) as a key proxy for liquidation risk, I find that the skewness risk premium rises (falls), but realized skewness remains unchanged, when OSP is more positive (negative). I also provide direct evidence of the price effects during such liquidation events. Trading strategies designed to theoretically exploit the skewness premium yield up to 2.5 percent per month and load significantly on risk factors related to the ease of funding for financial intermediaries.

In the second essay, I investigate the pricing dynamics of a class of option-like structured products, bank-issued warrants, using a large, high-frequency data set. I provide evidence that issuers extract rents from investors due to 2 key features of these markets: Each issuer is the sole liquidity provider in the secondary market for her products, and short-selling is not possible. As a consequence, I find that warrants are more overpriced the harder they are to value, and the fewer substitutes are available. Second, issuers are able to anticipate demand in the short term and preemptively adjust prices for warrants upwards (downwards) on days when investors are net buyers (sellers). Third, issuers decrease the amount of overpricing over the lifetime of most warrants, lowering returns for investors further. Lastly, while I find a negative relationship between issuer credit risk and overpricing, the effect is generally too small, is absent prior to the Lehman Brothers bankruptcy and does not conform to models of vulnerable options.

Preface

Both essays contained in this thesis, ‘Limits to Market Making, Liquidation Risk and the Skewness Risk Premium in Options Markets’ as well as ‘The Dynamics of Overpricing in Structured Products’, are solo-authored. Modified versions will be submitted for publication in finance journals.

Table of Contents

Appendix

List of Tables

List of Figures

Acknowledgements

Completing my degree at the Sauder School of Business would not have been possible without the continued support from a great number of faculty, colleagues, friends, and family during these seven years in Vancouver.

I owe an infinite amount of gratitude to my primary Ph.D. supervisor Maurice Levi, who provided support on so many levels – academically, financially, but foremost as a great mentor. I want to thank Tan Wang and Murray Carlson for their open doors and open ears as well as substantial encouragement and guidance on my research endeavors. Jan Bena, Lorenzo Garlappi and Ron Giammarino provided much welcome advice throughout my studies and especially while I was on the job market.

I was fortunate enough to be on the receiving end of a number of scholarships and grants. Special thanks go to the Canadian Securities Institute Research Foundation.

My colleagues and long-term friends Oliver Boguth and Mike Simutin have been helpful every step of the way and never minded discussing the latest out-in-left-field idea via Skype, not even on a Friday night or a Sunday afternoon. Alberto Romero and Vincent Grégoire deserve credit for enduring endless whiteboard sessions and philosophical discussions on the merits of representative agent asset pricing and the inner workings of Exchange-Traded Funds, respectively.

Finally, I cannot thank my family enough who never stopped believing, saw me through this long journey and will be just as overjoyed as I am that this chapter in our lives finds a successful ending.

Dedication

To My Parents.

Chapter 1

Introduction

Option pricing theory is inextricably linked to arguments of no-arbitrage and replication, i.e. the notion that two securities that provide the same payoff in all states of the world must have identical prices. Under this view, options are redundant securities whose payoffs, in the standard model, can be exactly replicated using the underlying security and a money market instrument. Option pricing models do not allow room for demand effects or financial frictions on prices. Yet the assumptions that have to be made to support the theory are plenty and many do not hold in practice (Bates, 2003).

The two following essays have surprisingly much in common, despite the fact that the data employed cover very different products that are subject to very different option market structures and competitive environments. At the heart of the investigation in both essays lies the question how demand for options by 'end users' (Garleanu et al., 2009) affects the pricing of options given certain financial frictions. The focus of the first is a risk-based explanation of certain limits to arbitrage effects in option markets, while the second investigates the consequences of a particular non-competitive market structure on prices.

The essays make use of the circumstance that the exogenous net demand for call options is negatively correlated with the net demand for put options for certain types of market participants. Whenever those groups of investors buy calls they tend to be sellers of puts in aggregate at the same time. Assuming less than elastic supply from intermediaries, this will increase the price of calls and lower the price of puts. This allows me to employ measures of option-implied skewness that compare prices of (out-of-the-money) puts with (out-of-the-money) calls, which perfectly captures the price effect one is looking for in this context. Finally, the modern non-parametric and model-free versions of option-implied measures being used in both cases avoid the pitfalls of using returns or prices of individual options (Bakshi et al., 2003; Kozhan et al., 2011).

1.1 Limits to Market Making, Liquidation Risk and the Skewness Risk Premium in Options Markets

The asset pricing literature has not yet reached a consensus on how to explain option returns of particular strategies such as writing index put options (Bakshi and Kapadia, 2003). No-arbitrage based asset pricing models (Benzoni et al., 2011) are forced to employ non-standard assumption in order to match the historically observed data. Deviating from the no-arbitrage framework, a small number of studies (Bates, 2003; Bollen and Whaley, 2004; Garleanu et al., 2009) advance the idea of frictions in the intermediation process of the options market, but are only able to link them to *option prices*.

The first essay concerns itself with establishing a firm link between observed *option returns* and potential limits to arbitrage experienced by option market makers, while deepening our understanding of the exact nature of these frictions. A newly developed measure of realized skewness (Neuberger, 2011; Kozhan et al., 2011) allows me to distinguish between option prices being affected by limits to arbitrage rather than reacting to informed demand and to capture the effect of these limits on option returns in a risk premium for skewness.

In the spirit of the recent limits to arbitrage literature that documents the effects of margin requirements (Brunnermeier and Pedersen, 2009), arbitrageurs capital scarcity (Hu et al., 2010) or financial intermediaries' balance sheets (Adrian and Shin, 2010) on prices in different asset classes, I focus on the effects of liquidity-related shocks (changes in margins, ease of wholesale funding, etc.) on option prices.

I propose a channel of liquidation risk, whereby financial traders that rely on leverage and short-term financing are forced to sell their option holdings following such an exogenous shock. Because financial intermediaries are capital constrained in the short term, upward sloping supply curves result in temporary price pressure in options potentially causing losses to other traders as well (Brunnermeier and Pedersen, 2009). Over time the price pressure dissolves, but rational market makers pre-emptively incorporate the possibility of such liquidation events into prices giving rise to a permanent price effect, similar in spirit to the liquidity run model by Bernardo and Welch (2004).

As a key proxy for the probability of such an event occurring I suggest the size of the net long positions of speculators in the options market (OSP) and show that it explains option prices as well as returns both in the cross-section and over time. Further support for the liquidation risk hypothesis comes from the asymmetric response to changes in OSP where an *on average* large price effect due to speculators reducing previously held positions points to occasionally even larger price effect during episodes of forced selling.

In the last part of the essay, net zero investment strategies of delta-hedged option portfolios allow me to quantify returns in options that are due to liquidation risk. I find very significant returns of up to 2.5 percent per month. Further, the employed strategies are found to load significantly on risk factors related to market-wide funding conditions.

1.2 The Overpricing of Structured Products

Retail structured products have become extremely popular among retail investors in many European and Asian countries. Being traded on easily accessible venues with low transaction costs and offering innovative payoffs not to be had otherwise (Stoimenov and Wilkens, 2005), this seems hardly surprising. The downsides of these products are less well-known (Henderson and Pearson, 2011; Bernard et al., 2011), especially among retail investors themselves[1]. The second essay focuses on one particular class of structured products, so called bank-issued warrants, whose payoffs are identical to those of cash-settled put or call options on any one of a variety of underlying assets (with one important caveat).

Especially intriguing about these instruments is the market structure they trade in. Each security is issued by one particular investment bank and can only be redeemed through that same issuer acting as a monopolist in the liquidity provision of its products. In addition, short-selling is not allowed. I investigate the pricing implications that arise due to these deviations from a typical competitive, no-arbitrage market environment. The existing literature generally limits itself to determining the average amount of overpricing in these instruments (Bergstresser, 2009). Instead, the main contribution of this essay is to investigate a) cross-sectional differences among warrants and b) price changes over time in response to retail demand. Using high-frequency tick data allows me to distinguish whether prices react to the arrival of demand due to limits to arbitrage faced by the issuer, which is similar to the focus of the first essay, or whether in fact issuers exploit the predictable components of retail order flow by adjusting prices in advance. My results point to the latter, contributing a rather novel finding to the literature that may also have implications for regulators in these markets.

The final part of the essay is concerned with yet another specialty of structured products, namely that they carry the credit risk of the issuer. This point has previously not received much attention, neither by academics (Bartram and Fehle, 2007) nor by retail investors[2]. Rather than starting from a model that automatically incorporates credit risk to compute theoretically fair values (Baule et al., 2008), I investigate if changes in issuer credit risk are properly reflected through changes in actual prices. Interestingly, I find that only after the Lehman Brothers bankruptcy is there a significant effect from credit risk. However, additional investigations reveal that issuer credit risk is still not fully priced.

[1] Wall Street Journal, May 28th, 2009, 'Twice shy on Structured Products?'
[2] New York Times, October 14th, 2008, 'Lehman's Certificates Proved Risky in Germany'

Chapter 2

Limits to Market Making, Liquidation Risk and the Skewness Risk Premium in Options Markets[3]

2.1 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)[4]. 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)[5]. 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 exogenous 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. They suggest that this price effect may be due to market makers being capital-constrained and unable to fully hedge themselves, giving rise to a particular kind of limits to arbitrage, i.e. limits to market making.[6]

These studies leave open a number of questions that this paper attempts to answer. First, assume demand was informed about the future distribution of the underlying instead of being purely exogenous. Then we would still expect to see the same effect in prices due to market makers

[3] A version of this chapter will be submitted for publication. Ruf, Thomas (2012) Limits to Market Making, Liquidation Risk and the Skewness Risk Premium in Options Markets.

[4] A 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).

[5] Exceptions 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.

[6] The literature uses the term 'limits to arbitrage' in many contexts, most often when describing mis-pricing of assets across two markets (Shleifer and Vishny, 1997). This paper suggest 'limits to market making' to describe a situation when market makers and liquidity providers more generally are limited in their ability to provide securities at prices that would prevail in the absence of certain frictions.

avoiding adverse selection, independent of the presence of financial frictions or limits to arbitrage. Thus, it is not clear ex ante that demand is a valid proxy for limits to arbitrage or more specifically, to market making.

Second, the literature has not related abnormal *option returns* to demand pressure directly.[7] An ex-ante decrease in the option price will not automatically translate into an ex-post increase in the option return if it is correlated with changes in the physical return distribution. Put differently, while there is evidence that demand pressure makes some options more expensive than others (e.g. OTM index puts relative to OTM calls), no one has established that they yield abnormally low returns as a consequence of demand pressure.

Third, when it comes to how the relative pricing of the tails, i.e option-implied skewness, is affected by demand, Garleanu et al. (2009) focus solely on jumps in the underlying asset as a possible channel. By contrast, this paper investigates the probability and the severity of liquidation events in the options market, of which jumps in the underlying are just one possible cause.

I consider an options market with commercial hedgers on the one side and financial traders on the other, whereby the latter are composed of market makers and outright speculators. Typically, the former group is liquidity consuming while the latter are liquidity providers. Importantly, financial traders, especially speculators, are much more prone than hedgers to external liquidity shocks such as changes in the ease of funding, losses in other markets or changes in margin requirements, largely as a result of their use of leverage and short term financing.

The arrival of such a shock may force some speculators that were initially, say, net long in options to quickly close out positions (sell calls and buy puts) while temporarily affecting option prices in an adverse manner, e.g. depressing call prices and raising put prices, as market makers struggle to accommodate this order imbalance. This in turn leads to losses for other traders with similar positions and possibly further selling (e.g. liquidity spirals as in Brunnermeier and Pedersen, 2009). In equilibrium, the potential for these liquidation events is anticipated and priced in by rational traders (Bernardo and Welch, 2004), in particular by well-informed market makers, so that the price of calls (puts) will fall (rise) as speculators increase their net long positions and vice versa. In this context, the aggregate net position of speculators emerges as an important state variable that determines the price of the option smile and the relative return of calls and puts via the mechanism of liquidation risk.

The notion of liquidation risk can explain permanent price effects without having to assume barriers to entry or slow-moving capital in these markets. Besides the jump risk faced by market makers, this constitutes another risk-based limit to market making in that it pushes prices away from those that would prevail in the absence of potential liquidation shocks affecting the options market.[8]

[7] Cao and Han (2011) find option open interest affects option returns and regard it as a proxy for demand pressure, but do not provide any evidence for this assumption.

[8] When market makers and non-market making financial traders hold similar net positions, the inventory risk effect of the first group (Garleanu et al., 2009) and the liquidation risk effect of the second group align and may reinforce each other.

My methodology allows me to identify what factors truly represent frictions to financial intermediaries in the options market as opposed to rational price setting by unconstrained market makers e.g. in the face of informed demand. To this end, I decompose the price of skewness into ex-post realized skewness and a skewness risk premium using the newly developed skewness measures of Neuberger (2011). Economic intuition would suggest that limits to arbitrage within the options market should influence only the skewness risk premium, but not affect ex-post realized skewness, which is determined by the physical price dynamics in the underlying asset market.[9] By contrast, beliefs about the distribution of physical returns, e.g. informed demand or the convenience yield, should be reflected in realized skewness.

The variance risk premium is the subject of a large number of previous studies (Carr and Wu, 2009; Bollerslev and Todorov, 2011; Driessen et al., 2009; Bollerslev et al., 2011). 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.

My work empirically distinguishes between beliefs and risk premia as determinants of 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.[10]

I conduct the empirical analysis on the market for commodity futures options rather than index or equity options because they do not have many of the disadvantages of equity and equity index options when it comes to measuring the price of tail risks and identifying the effect of demand on prices. My dataset consists of a cross-section of 25 U.S.-listed optionable commodities over a time period of 20 years. I merge this dataset with data on the aggregate futures and options positions of different groups of traders, notably commercial traders that trade for hedging purposes and

[9] The same can be said about factors reflecting the economy's risk aversion. Thus, allowing for a delineation of risk aversion effects from financial market frictions may depend on careful economic motivation.

[10] Kozhan et al. (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.

financially motivated, non-commercial traders, also called speculators. It can be assumed that the latter group also contains market makers and arbitrageurs that behave similarly to market makers by providing liquidity and taking no directional exposure in the underlying. However, the data does not allow me to disentangle market makers from outright speculators.

The central new variable is the (delta-weighted) net position of speculators as a share of their total positions in options (henceforth OSP). I consider extreme readings of this variable as indicative of at least one of two situations. First, if speculators absorb large outright positions from hedgers they grow more susceptible to external shocks. The probability of forced liquidations thus increases and to insure against adverse price movements in such an event, new entrants preemptively demand price concessions. Second, if speculators are largely acting as liquidity providers, extreme levels of OSP show that market makers carry large unbalanced inventories and may be close to capital and/or risk constraints similar to Garleanu et al. (2009). Besides exposing them to jump risk in the underlying asset, this also makes them less able to accommodate future demand imbalances and may make them susceptible to forced liquidations themselves. Note that the risk of forced liquidation increases regardless of which group the majority of traders in the non-commercial category belong to.

I begin by showing supporting evidence from the cross-section of commodities where cross-sectional dispersion in average OSP is strongly related to average implied skewness as well as 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. There I find that the skewness risk premium becomes more positive (negative), but realized skewness is unaffected in times of high positive (negative) values of OSP. In further support of the liquidation risk hypothesis, I find that, controlling for the *level* of existing demand, price pressure has a decidedly asymmetric effect on option prices, and as a consequence on the skewness risk premium. When speculators increase their net long exposure having been long already, there is no additional price impact. By contrast, when speculators reverse previously held long positions, a strong adverse price effect can be observed *on average*. Because financial traders optimally avoid price impact in normal times, this provides evidence for an even larger price impact during episodes of rushed selling by speculators.

Further, I find that reductions in market-wide liquidity and loss of funding for financial intermediaries (Adrian et al., 2011a,b) magnify the effect of OSP on prices. Lastly, a measure of trader concentration, which is another dimension of how likely forced liquidations are, also affects the price of skewness and the skewness premium, but not realized skewness. The fact that realized skewness is unaffected in (almost) all cases above gives great confidence that the suggested factors are indeed related to limits of market making and do not proxy for investor beliefs or informed demand.

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? To answer this question, 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 2.5 percent per month. Portfolios formed on the basis of net long trader concentration yields up to 1.6 percent. Interestingly, both strategies load significantly on risk factors related to market-wide funding conditions.

The remainder of this paper is structured as follows. Section 2.2 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 2.3 describes the skewness measures and presents the historical skewness and skewness risk premia found in the data. Section 2.4 motivates the concepts of limits to market making, outlines the idea of forced liquidations and presents visual evidence in support. Section 2.5 contains the design and results of the main empirical analysis conducted on a large panel dataset. Section 2.6 examines portfolio strategies constructed on the basis of two limits to market making criteria. Section 2.7 concludes.

2.2 Commodity Futures (Options) 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).

Figure 2.1 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'[11]. 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.

2.2.1 Empirical Advantages over Equity Options Data

I conduct the empirical analysis on the market for commodity futures options rather than index or equity options for a number of reasons. All empirical work using S&P 500 index options is bedeviled by the paucity of data, making estimation of the physical return distribution difficult. Options on a broad market index are not well suited for analyzing the effects of frictions and limits to arbitrage vis-à-vis other effects, because funding constraints for financial intermediaries 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. by size, and jointly driven by investor sentiment (Lemmon and Ni, 2011). 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 do not exhibit these problems. My dataset was acquired from the Commodity Research Bureau (CRB) and 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.[12] Option chains of some commodities have quotes for more than 100 options at a time. In addition, replication and hedging arguments rely on the ease of short-selling the underlying, which is easily done in the case of futures.

The data that was available for the present study is, however, subject to criticism on other fronts. The lack of bid/ask quotes makes it difficult to judge a) how noisy individual quotes are and how much of this noise trickles down into the non-parametric skewness measures and b) how much of the out-performance α found in Section 2.6 is achievable in practice[13]. Second, as the CRB data lacks information on option volume and open interest, liquidity-based filters cannot be applied. Instead, motivated by the fact that option liquidity generally decreases away from the

[11] Available at http://www.bis.org/statistics/derstats.htm.

[12] This is not to say that some commodities, in particular in agriculture, do not exhibit at times large correlations. However, compared to the market-induced correlation between e.g. large-cap stocks they are mild and less persistent.

[13] A number of hedge fund analysts were quick to point this out in previous versions of the paper.

money, I exclude options with prices that were small relative to the minimum tick size. However, most available option datasets are affected by some of these issues as well. For further details, see Appendix A.1.1.1.

2.2.2 Futures and Futures 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. Together, they contain 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 2.1 lists the sample containing 25 commodities which can be roughly divided into 5 major groups: Agricultural, Energy, Meat, Metal and Soft.[14] 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 90 days.[15] The data ends at the end of August 2010, so that the last option expiration is the August or September contract of that year, in most cases. Note that 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. Appendix A.1.1 describes the structure of the data, the cleaning process and construction of implied and realized measures from option prices in more detail.

Table 2.1 also shows the interquartile range of the key variable in the empirical part, the net long exposure of speculators in options (short OSP) and its time series correlation with the net long exposure of speculators in futures. Both the level and the correlation exhibit substantial variation across the sample, an advantage of the data alluded to earlier.

2.2.3 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.[16]

[14] My set of commodities is similar to the U.S. based commodities used in Gorton et al. (2007). Compared to Szymanowska et al. (2011), I exclude platinum and palladium, because there are no liquid options available, but I include natural gas.

[15] To 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.

[16] Currently, 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 limit. Futures-only COT

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. market makers, trading desks, hedge funds, commodity trading advisors (CTA) and commodity pools, collectively called speculators in the COT reports. However, it appears likely that at least some of the traders in this group do not take on outright speculative positions, but are instead in the business of market making by accommodating the liquidity needs of other traders. 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. These and other measures derived from the COT data are define below.

Hedging pressure in futures and options: 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:

$$HP_{i,t} = \frac{\sum \text{comm. short fut.} - \sum \text{comm. long fut.}}{\sum \text{comm. short fut.} + \sum \text{comm. long fut.}} \tag{2.1}$$

By construction, HP lies between -1 and $+1$. HP has been found to positively predict futures returns (see e.g. Bessembinder (1992), De Roon et al. (2000)). Apart from a futures-only report, combined futures-and-options reports are available since April 1995, where options positions are transformed into futures-equivalents using each option's delta (based on an options 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):

$$OHP_{i,t} = \Big(\frac{\sum \text{comm. short options} - \sum \text{comm. long options}}{\sum \text{comm. short options} + \sum \text{comm. long options}}\Big) \tag{2.2}$$

Speculative net exposure in futures and options: For each of the two derivatives, futures and futures options, corresponding measures of net *long* exposure can also be constructed for the group of speculators, denoted by SP and OSP respectively.

$$SP_{i,t} = \frac{\sum \text{non-comm. long fut.} - \sum \text{non-comm. short fut.}}{\sum \text{non-comm. long fut.} + \sum \text{non-comm. short fut.}} \tag{2.3}$$

Because hedgers and large financial traders generally constitute the overwhelming majority of total futures open interest, the numerators are quite similar for most commodities and, due to the

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

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. As with futures, the flip side to OHP is the normalized net *long* exposure of arbitrageurs, OSP:

$$OSP_{i,t} = \Big(\frac{\sum \text{non-comm. long options} - \sum \text{non-comm. short options}}{\sum \text{non-comm. long options} + \sum \text{non-comm. short options}}\Big) \tag{2.4}$$

As before, OHP and OSP 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 demand is relative to their total exposure in each market, while SP and OSP do the same for speculators. For hedgers, a more extreme position does not necessarily equal more risk, as the general purpose is to lock in forthcoming production or consumption of a commodity at a certain price and eliminate risk. For speculators as a group, however, leaning to one side can be a risk in and by itself. Should financial traders as a whole be forced to sell out of their positions due to external shocks, it may cause price pressure and result in losses to the group.

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. These are available on a net and a gross basis, whereby the former means that only the residual positions for each trader after offsetting equal long and short positions are counted. To judge the true concentration of traders with unbalanced positions, the net version is preferable. I use the concentration ratios reported in the combined futures-and-options data for $N = 8$, e.g. for the long side

$$\text{CRL}_{i,t} = \frac{1}{\text{OI}_{Com,i,t}} \sum^{j \in \text{Top } 8} (\text{Long Com.})_j \tag{2.5}$$

Ceteris paribus, if a larger fraction of the net exposure on one side of the market is held by a small number of traders, the price impact will be relatively larger, should one of them be forced to sell quickly. I will motivate these measures further in Section 2.4.1.

2.3 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, 2011). 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). Practitioners, especially in currencies use so-called risk reversals (see also Brunnermeier et al., 2008) to express the asymmetry in price between an equally out-of-the-money put and call option, quoted as the difference between their implied volatilities for a given absolute level of their Δ.

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 un-centered 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 et al. (2009), Rehman and Vilkov (2012) for index and equity options. Note that the skewness implied by this measure is the normalized 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 et al., 2010) using daily stock returns.

2.3.1 Raw Skewness Measures

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 re-balancing, it is equal to the integral of the product of asset return and innovations to its forward-looking variance and is synthetically constructed as the difference between two measures of variance, v^E and v^L (for the details of derivation see the technical appendix as well as KNS):

$$\mathrm{MFIS}_{0,T} := 3\left(v^E_{0,T} - v^L_{0,T}\right) \tag{2.6}$$

$$= \mathbb{E}^{\mathbb{Q}}\left[3\int_0^T dv^E_{t,T}\left(\frac{dS_t}{S_t}\right)\right] \approx \mathbb{E}^{\mathbb{Q}}\left[r^3_{0,T}\right] \tag{2.7}$$

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 A.7 in the appendix) is a function of several such integrals. Most important and key to the focus of this paper, the KNS measure has a realized counterpart. The question how risk-neutral skewness compared to realized skewness was previously impossible to answer. No natural counterpart to implied skewness was available that was able to measures realized skewness over some time period. Equation 2.7 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.

$$SRP_{0,T} = \mathbb{E}^{\mathbb{P}}\left[3\int_0^T dv^E_{t,T}\left(\frac{dS_t}{S_t}\right)\right] - \mathbb{E}^{\mathbb{Q}}\left[3\int_0^T dv^E_{t,T}\left(\frac{dS_t}{S_t}\right)\right] \tag{2.8}$$

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

The empirical analysis in Section 2.5 is based on the KNS measures of implied and realized skewness, abbreviated MFIS and RSkew. Because MFIS (just like the model-free implied variance or MFIV of Britten-Jones and Neuberger, 2000) represents a moment of returns over the remainder of the life time of the option chain, it gets naturally smaller as time progresses and ultimately converges 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 $I_{3\text{months}} = (90, \ldots, 99)$. More formally, for commodity i and expiration date T, given a maturity m, e.g. 3-month skewness, I average over observations at times t where $T - t \in I_m$.

$$\text{MFIS}_{i,m,T} = \frac{1}{|I_m|} \sum_{T-t \in I_m} \text{MFIS}_{i,t,T} \tag{2.9}$$

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).

$$\text{MFIS}^{1/3}_{i,m,T} := \text{sign}(\text{MFIS}_{i,m,T}) \left(|\text{MFIS}_{i,m,T}|\right)^{1/3} \tag{2.10}$$

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.

$$\text{RSkew}^{1/3}_{i,m,T} := \text{sign}(\text{RSkew}_{i,m,T}) \left(|\text{RSkew}_{i,m,T}|\right)^{1/3} \tag{2.11}$$

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, 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.

$$\text{SRP}^{1/3}_{i,m,T} := \frac{1}{|I_m|} \sum_{T-t \in I_m} \left[\text{RSkew}^{1/3}_{i,t,T} - \text{MFIS}^{1/3}_{i,t,T}\right] \tag{2.12}$$

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 2.6 aims to provide some answers in that regard.

2.3.2 Normalized Skewness Measures

The extreme nature of the third moment requires some transformation to make it suitable for statistical analysis. Taking the third root is a natural solution, but it has drawbacks, in particular for the skewness premium. Just as it is not correct to measure a volatility risk premium as the difference between implied volatility and realized volatility due to Jensen's inequality, the same inconsistency arises for the third root of skewness.

The obvious other solution is to divide the skewness measures by some function of volatility to normalize it. However, the question of which volatility to use is not as straightforward as one might think. Potential candidates are historical realized volatility, currently implied volatility or even future expected realized volatility. Each choice constitutes an implicit weighting scheme regarding skewness risk, while the raw measure implicitly assigns equal weights to all time periods when e.g. computing average risk premia.

In the latest version of Kozhan et al. (2011), skewness is normalized by the implied log variance contract v^L (to the power of 3/2):

$$\mathrm{SMFIS}_{0,T} := \frac{3\left(v^E_{0,T} - v^L_{0,T}\right)}{\left(v^L_{0,T}\right)^{3/2}} \tag{2.13}$$

$$= \frac{\mathbb{E}^{\mathbb{Q}}\left[3\int_0^T dv^E_{t,T}\left(\frac{dS_t}{S_t}\right)\right]}{\mathbb{E}^{\mathbb{Q}}\left[-2\log\left(\frac{S_T}{S_0}\right)\right]^{3/2}} \approx \frac{\mathbb{E}^{\mathbb{Q}}\left[r^3_{0,T}\right]}{\mathbb{E}^{\mathbb{Q}}\left[r^2_{0,T}\right]^{3/2}} \tag{2.14}$$

Normalized realized skewness and the skewness risk premium are then normalized by the same denominator. In addition to having time-varying weights this choice allows the variance risk premium to directly affect the normalized measure via the denominator. It is difficult to judge how this should affect results. Using several standardizing techniques, I find that all of them exhibit massive problems with outliers much worse than when using the third roots of the raw measures. As a consequence, the empirical results using these standardized measures differ starkly depending on the denominator employed and are not reported. Instead, favoring simplicity and an equal-weighting scheme, I opt in favor of using the measures based on signed third roots outlined in Section 2.3.1.

However, even with the raw measures one might have to worry about an effect from the variance risk premium on the skewness risk premium via demand. As will be discussed later in more detail, the key focus of the paper is the effect of directional (e.g. long delta positions) demand by hedgers on relative option prices, i.e. skewness. Garleanu et al. (2009) show that 'absolute' demand for options (i.e. buying vs. writing them) drives the variance risk premium. Thus, if the directional demand was correlated with the absolute demand, then the skewness risk premium could partially pick up the variance risk premium. This can only occur if hedgers had a preference to implement their views mostly via puts or mostly via calls only. I am not aware of option markets where the open interest is asymmetrically concentrated among calls or puts on average. Further, for this to have a systematic effect on the results presented, hedgers would have to have the same preference for a majority of the commodities in the sample. I would argue that this, too, is unlikely to occur and thus there should be no confounding effect from the variance risk premium on the skewness risk premium using the raw measures.

2.3.3 Sample Variance and Skewness

To see how skewness is related to demand from different groups of investors and limits to market making, 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 the interest of brevity, I reserve the exact definitions and computational details of all the measures involved for the Technical Appendix A.2 to the extent that they were not already discussed in Section 2.3.1.

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 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 2.2 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. These findings extend the results of Trolle and Schwartz (2010), who investigated the variance risk premium of oil and natural gas only. 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 market making, I will leave this topic to future research.

Skewness Risk Premium

While a measure of option-implied skewness has been proposed in the literature (Bakshi et al., 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 skewness 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 2.3 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 a discrete approximation of 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 covariation between returns and changes to implied variance is positive. 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.[17] 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.

2.4 The Economics of Option Markets

What can explain the notable dispersion of average implied skewness as well as risk premiums for skewness across commodities in Table 2.3? Cross-sectional variation in implied skewness could potentially be due to corresponding differences in beliefs about the return distribution of the underlying asset. The price for natural gas certainly behaves quite differently from that for corn in many ways, which should be reflected in physical moments like volatility and skewness. 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 particular, different signs for the SRP suggest that for some commodities returns to risk reversals, i.e. long OTM calls and short OTM puts, will on average be positive while they are negative for others over time. What drives these differences in returns across assets?

In what follows, I argue that the cross-sectional pattern apparent in Table 2.3 is the consequence of what I will call liquidation risk, i.e. the risk that arises from the possibility of sudden, forced selling of positions held by financial traders.

2.4.1 Exogenous Hedging Demand and Liquidity Provision

I consider a framework of a commodity futures option market with three notable groups of traders: commercial hedgers, speculators and market makers. Hedgers have an exogenous reason to be in the market due to their line of business, i.e. they are exposed to the price of a commodity they either produce or use as input and require insurance against adverse price movement. Hedgers tend to be price takers as well as liquidity takers in that they are less driven by the current price but

[17] As 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.

by immediate operational concerns. Speculators and market makers, on the other hand, tend to be large, sophisticated financial institutions, hedge funds or trading firms that opportunistically engage in financial markets with a profit motive. As a consequence they are more selective regarding at what price they enter a position and both can be broadly described as liquidity providers.

I make a distinction between outright speculators and market makers because they differ in the way they manage their positions. The former will take outright option positions that expose them to movements in the underlying without necessarily hedging any of the associated risk, while the latter immediately lay off directional risk via offsetting futures or options positions. The motivation of a speculator to enter an option position may be driven by his view that the option is cheap given his outlook on the underlying, while market makers do not take any view of the underlying market and instead attempt to profit from the spread they earn between buying and selling while keeping their inventories balanced.

Critically to the channel I have in mind, market makers, but especially financial speculators, differ from hedgers in one important aspect, namely that the former are much more prone than the latter to external liquidity shocks such as changes in the ease of funding, losses in unrelated positions or changes in margin requirements. This is turn is due to their use of leverage in derivatives markets and their reliance on short-term financing such as repurchase agreements. A recent study in commodity futures markets by Cheng et al. (2012) supports this notion. They find that during the recent financial crisis, increases in equity market volatility led financial traders (hedge funds and commodity index traders) to close out positions while some hedgers increased them. No such effect was observable prior to the crisis.

Further, the recent literature is full of examples where availability of arbitrage capital (Hu et al., 2010), margin requirements (Garleanu and Pedersen, 2011), funding liquidity (Brunnermeier and Pedersen, 2009), bank leverage (Adrian and Shin, 2010) or intermediaries' balance sheets (Adrian et al., 2011a,b) affect asset prices in stock markets, individual equities, currencies or commodity futures. Thus, this study can be seen along those lines as being concerned with the effect of liquidity shocks such as changes in margins or funding access on the prices and returns of options.

In any financial market, financial traders 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 the ability of arbitrageurs and market makers to raise funds and their willingness to add to their existing positions to accommodate that demand given a price. This type of limits to arbitrage is what I call broadly limits to market making.

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 supporting the theory (Bessembinder, 1992; Bodie and Rosansky, 1980; De Roon et al., 2000).

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.[18] For instance, Etula (2010) and Acharya et al. (2011) both assume VaR constraints for speculators/financial intermediaries, which makes them effectively risk averse, giving rise to limits to arbitrage in both models.

The same principle should 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. Financial traders, be they outright speculators or market makers will fill that void, but may only do so under price concessions from hedgers, just like in the case of futures. In this example, this should lead to more expensive puts and/or cheaper calls, or simply more negative implied skewness.

But what exactly determines the extent of the price impact? Is it purely based on the amount of hedging demand or does it depend on the aggregate positions of speculators? Assume there exists currently a fixed level of hedging demand of short x units via options. In one case the group of financial traders collectively has a large number of long and short positions relative to x, while in another case the x units of long option positions constitute a large portion of their overall positions. To evaluate the different pricing implications in those two situations, I investigate below how measures of hedging demand and net speculator exposure relate to various measures of skewness in the cross section of commodities.

2.4.2 Evidence from the Cross-Section

Figure 2.2 contains six scatter plots. The top row shows plots of the time series average of OHP, the net short hedging demand in options, against the average of 3 measures of option-implied skewness, the measure proposed by Bakshi et al. (2003, , BKM), the price of a risk reversal (quoted in % IV) and a scaled version of the measure proposed by Neuberger (2011). The bottom row shows plots against time series averages of measures of ex-post realized skewness (also proposed by Neuberger, 2011), against a scaled version of the skewness risk premium (SRP) (again Neuberger, 2011), and finally a Newey-West adjusted t-statistics of this premium testing the significance of the premium against being zero. All measures except for the risk reversal are as in Table 2.3.

From the top row, 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 commodities in which hedgers generally hold larger short options positions, i.e. have a more positive OHP, tend to have more negative implied skewness and vice versa. OHP seems to be similarly

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

related to ex-post realized skewness as well as the skewness risk premium, albeit in the case of SRP rather weakly. We see that more net short hedging pressure via options is related to a more positive SRP, indicating that implied skewness is on average more negative than realized skewness for those commodities. For those commodities, OTM puts (calls) will have lower (higher) returns on average. OHP explains around one third of total cross-sectional variation of implied skewness, realized skewness and the SRP, respectively. This could be seen as supporting the notion that prices are affected in a way that is consistent with demand pressure from hedgers. Note, however, that had I chosen total open interest in the denominator instead of total hedging positions, the fit in all plots would have been significantly worse.

In Figure 2.3, I show the corresponding plots against OSP, i.e. the net long exposure of speculators scaled by their (speculators) total exposure. We notice a at times dramatic increase in fit across all six plots. An R-squared of 72 percent implies a correlation of over 84 percent in the case of BKM implied skewness. Note that a positive reading of OSP implies that speculators are long calls and short puts. This is on average correlated with negative skewness, i.e. low call prices relative to higher put prices. Thus, it can be ruled out that speculators exert demand pressure. Further, the direct comparison of hedging pressure with speculative exposure lends strong support to the idea that it is the relative net exposure of speculators that matters for skewness rather than pure hedging demand.

To avoid the criticism that the documented correlations are mechanical in some way, I plot the implied skewness measures and premia against ex-post realized skewness in Figure 2.4. Naturally, implied skewness and realized skewness are highly correlated. However, the correlations with the premium measures are much lower than the corresponding values for OSP, which shows that the latter contains important information for option prices as well as returns in commodity futures.

2.4.3 Possible Explanations

The cross-sectional findings presented in Figures 2.2 and 2.3 merely show a correlation and do not prove causality, but they serve as a starting point. The fact that the net long options exposure of speculators (OSP) relates in the observed fashion both to relative option prices (implied skewness) as well as returns is potentially consistent with the notion that market makers are subject to frictions such as capital constraints or risk limits (Bates, 2003; Bollen and Whaley, 2004), i.e. limits to market making. This is under the assumption that market makers constitute a large proportion of the group classified as speculators by the CFTC.

The cross-sectional results represent a level effect, i.e. commodities with on average more positive OSP have permanently more negative skewness (and vice versa). If this was due to lack of available capital alone, it raises the question why does new capital not flow into these markets over time, for example a hedge fund could easily devote part of its capital to 'help out' existing market makers and profit from the premium for skewness in options.

One potential explanation could be barriers to entry into market making. The Chicago Mercantile Exchange (CME), for example, does indeed give preferential rates to the owners of its exchange seats, making it all but impossible to compute with seat owners at making markets as an outsider. However, these seats are freely and publicly traded, so that one would expect that if capital was all that was needed to remove these 'mis-pricings', over time well-capitalized traders would move in to take advantage. A risk-based explanation seems much more plausible, which would cause the inflow of new capital to stop at the point where risk-adjusted returns from market making in general and the skewness premium in particular go to zero.

Garleanu et al. (2009) suggest the risk of imperfect hedging as culprit. They formalize this notion in a model that incorporates specific frictions that market makers face when trying to optimally hedge their book over time. Most notably, in the case of skewness, which is what we are concerned with here, their model finds that it is jump risk in the underlying asset that makes hedging imperfect and thus costly and risky for market makers. Empirically, they show that market maker positions weighted by their exposure to jumps do indeed explain some of the variation in implied skewness. However, they only conduct this test for the equity market index on a relatively short sample period.

Further note that under delta-hedging, jumps in the underlying inflict losses only on written options, while jumps always lead to gains on bought options regardless of direction. Thus, market makers being long delta-hedged calls and short puts face a reduced jump risk on their portfolio, as most jump sizes lead to small gains or losses and large upward jumps even to large gains, while only very large negative jumps result in large losses. Secondly, for an observed level of OSP, market makers may at different times tilt more heavily towards long calls rather than short puts and vice versa, leading to variations in the actual exposure to jumps. This is not to say that jumps do not play a role, but other mechanisms may be better suited to explain what we observe in the present context.

2.4.4 A Theory of Liquidation Risk

I propose an explanation based on the idea that financial traders, especially outright speculators tend to be disproportionately affected by exogenous liquidity shocks such as deterioration in funding liquidity or increase in margin requirements. This is largely owed to their use of leverage and their preference for short-term financing. Commercial hedgers face this problem much less due to their natural hedge position. As a consequence, financial traders occasionally are forced out of their positions in a hurry following a liquidity shock while hedgers are not.

Initially a given shock may only affect a small proportion of financial traders, but the fear of additional selling may lead to preemptive selling by other traders much like in the model of Bernardo and Welch (2004), where the fear of a liquidity run alone can cause one. In addition, the first wave of selling will affect prices adversely, so that other traders with similar exposures that initially were not affected may now face losses and/or margin calls and be forced to cover as well in a reinforcing loss spiral.

The severity of the liquidation within a given (options) market will be a function of the size of the initial net positions of speculators relative to overall open interest, and more importantly relative to the capacity of (yet) unaffected liquidity providers. As speculators reduce their positions, it generally falls to market makers and other arbitrageurs to pick up the slack. But also the likelihood that a forced liquidation will occur increases with the initial level of the positions of financial traders. More unbalanced inventories of speculators, or liquidity providers more generally, may indicate that they are close to hitting inventory, margin and/or risk constraints even in the absence of an external shock. This increases the chance that some are forced to liquidate their positions for a shock of a given size.

The finance literature knows of a number of instances where the presence of speculators has the ability to destabilize markets, for example, via the carry trade in currency markets. Brunnermeier et al. (2008) find that currencies that speculators tend to herd into are subject to crash risk, i.e. a forced liquidation out of their positions due to some external shock, e.g. losses in other markets. The small size and generally much lower liquidity of option markets does not mix well with one group of traders having to exit in a hurry. Price effects will be magnified in this case. More specifically, assume that initially, speculators were net long in options, having bought calls and sold puts. Upon the arrival of some shock, speculators are forced to drastically reduce their positions, buying back puts and selling calls. Other liquidity providers would need to sell puts and buy the calls. Given their aforementioned constraints, this will lead to a price impact pushing up put prices and reducing call prices. This price pressure, however, should vanish with time as liquidity providers re-balance their inventories and additional capital moves in to take advantage of prices being out of line.

Market makers face some additional complications when deciding whether to step in to absorb inventory during liquidations, especially in option markets. First of all, it is not always immediately clear to market participants what the cause of some order imbalance may be; the order flow may be informed instead of liquidity driven. Second, the perceived dislocation in prices, e.g. of OTM puts relative to OTM calls, depends on an option pricing model exposing option market makers to model risk. Third, given the relatively wide bid/ask spreads in options it is almost impossible to shut down an options book without incurring prohibitive costs, should the market maker himself become subject to funding problems. Lastly, and interesting from a strategic point of view, the question at what price to absorb the order flow will depend on expectations regarding the total liquidation volume and the number and size of other market makers available as well as their expectations and so on. Somewhat related, Cheng et al. (2012) find that during the recent financial crisis, it was in fact large commercial traders that assumed the roles of liquidity providers and added positions whenever the volatility index VIX was rising while hedge funds were doing the opposite.

Note that forced selling can potentially take place exclusively within the options market without greatly affecting the price of the underlying futures contracts. However, Table 2.1 suggests that speculative positions between options and underlying futures are positively correlated for most commodities in our sample. Thus a liquidity shock may also cause forced selling of long futures contracts which could very well result in a downward jump. Alternatively, the jump may occur exogenously and cause forced selling by speculators in both markets. Thus, the jump risk of Garleanu et al. (2009) can be a cause as well as consequence of liquidation risk and generally only high-frequency data would allow the researcher disentangle the two.

Another interesting question concerns the number of traders that are leaning heavily to, say, the long side, holding the amount of total net long exposure constant. If outright speculators are highly concentrated, this represents another dimension of the risk of liquidations occurring. Fewer, but larger traders may be especially desperate to liquidate given a liquidity shock, increasing the price impact. When it is already mostly market makers that hold those concentrated positions, it may affect their ability to take on additional exposure. Further, allowing for strategic considerations, those knowing that they face less competition in a market may strategically offer to trade at less favorable prices (Oehmke, 2009; Fardeau, 2011). Both channels affect prices in the same way.

In equilibrium, the consequences of liquidation events, in particular the potentially large temporary price impact, will be anticipated and priced in by rational traders. Therefore, option prices as well as returns should respond to the probability and the potential severity of a forced liquidation in the options market. Below, I suggest a number of hypotheses that attempt to channel the main points of the discussion into empirically verifiable statements.

By itself, the finding that OSP relates to implied skewness and the skewness risk premium in the cross-section does not provide evidence that is exclusive to the notion of liquidation risk. For instance, it is possible that some unknown factor causes both the observed cross-sectional patterns in the skewness measures as well as different propensities by commercial traders to use options to hedge different commodities or by speculators to implement their views. Differences in individual industry or market structure are potential candidates. In order to build additional support for limits to market making in general and the relevance of liquidation risk in particular, I investigate in Section 2.5 how the price of skewness (i.e. option-implied skewness) as well as the return on skewness (i.e. the skewness risk premium) is affected along the following dimensions *over time*:

H1: The current net long position of financial traders in options (OSP) is negatively related to current option-implied skewness, but has little impact on future realized skewness, giving rise to a premium.

H2: Under normal market conditions, (short-term) price pressure does not affect the price of options.

H3: During a liquidation event, it is financial traders, not hegders, that face price pressure.

H4: Deteriorations in market-wide funding liquidity conditions magnify the effects of OSP in **H1**, as the likelihood of and price impact given forced selling increases.

H5: If the net exposure on the long (short) side of the market is concentrated among fewer traders, implied skewness will be more negative, and the premium more positive.

In summary, the notion of liquidation risk driving the skewness risk premium still makes use of the assumption that the market making sector cannot absorb infinite amounts of demand in the short run, as such it should be seen as an extension of a more general view that limits to market making affect option price and returns (Bollen and Whaley, 2004). Hypotheses (H1) and (H5) could be as much a consequence of the risk of liquidation as well as of a market making sector that is generally constrained with regards to inventory capacity, capital and internal position limits. However, as mentioned in the discussion of the cross-sectional evidence, the notion of purely capital-constrained market makers does not explain why those price effects should be permanent, absent insurmountably large barriers to entry. Garleanu et al. (2009) provide a risk-based explanation that focuses on jumps in the underlying asset, while I suggest another risk that is independent of such jumps occurring. Finally, the idea of liquidation risk differs also in the differential treatment of hedgers' vs. speculators' exposure due to their different sensitivity towards market-wide liquidity conditions. Thus, previous explanations do not directly speak to (H3).

2.5 Time Series Results

In Section 2.3.3, 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 price of options as well as their returns. This is visible in Figure 2.3 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. The skewness measures as well as the independent variables are de-meaned separately for each commodity in order to filter out any level effects that may obfuscate the results.

2.5.1 Design of the Time Series Analysis

With the exception of Table 2.4, 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 (2.10) and (2.12). By taking the difference between the coefficients within each column, the corresponding coefficient for the realized skewness can be recovered. This can be seen in Table 2.4, which for the purpose of demonstration contains the results for realized skewness as well. Minor differences are due to the fact that I winsorize each variable including the skewness measures separately (at the 5th and 95th percentile).

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 and option returns are not affected. If, however, 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 that is independent of the physical return distribution of the underlying asset. In the latter case, the effect would thus seem to be driven by financial frictions in the options market.

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, all regressions have predictive character and, just like forecasting regressions for equity returns, generally low R^2. The choice of 3-month horizons constitutes a compromise between having timely predictors and a large enough sample on the one hand and less noise in the realized measure on the other.

Table 2.4 depicts the coefficients for a number of control variables that are included in each regression that follow, 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), the current convenience yield that is implied by the front futures contract and the futures contract that is closest to a 6-month maturity, and the most recently available realized skewness measured over 3 months. 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 the next two variables is motivated by the existing literature. Brunnermeier et al. (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 two variables add some explanatory power to the regression, but are mostly insignificant. Growth in open interest increases both implied as well as realized skewness by a similar extent, leaving the SRP unchanged. Finally, lagged realized skewness has a very strongly positive effect on both implied and realized skewness, with implied skewness being affected even more. The end result is a significantly negative effect on the premium.

Further, both the speculator net exposure in options (OSP) and the concentration ratio of net long traders (CRL) have strong effects on the relative pricing of options. The meaning of these results will be discussed below in separate tables. The regressions shown in the following tables that discuss different aspects of OSP do include CRL as a control, but OSP is only included when specifically indicated (and vice versa).

2.5.2 The Effect of Demand Pressure

The literature on commodity futures (e.g. Bessembinder and Chan, 1992) finds the futures risk premia to be predictable by hedging pressure. 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. That is, the demand by hedgers not speculators drives returns in futures. Here, I test if demand pressure variables affect any of the skewness measures. My analysis includes the hedging pressure and speculators' net exposure present in the underlying futures market as well as the equivalent measures in the futures options market.

Table 2.5 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) have any affect on skewness. More interesting are the following 2 columns containing the corresponding measures based on the option market, OHP and OSP. 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 of Figures 2.2 and 2.3. Not only do commodities with on average high levels of OSP have more negative skewness on average, this table provides evidence that commodity skewness reacts to variations in the level of OSP over time for a given commodity. Further, because the effect is absent for realized skewness it points towards frictions in the options market. As a consequence, OTM puts will exhibit lower returns and calls relatively higher returns in times when OSP is high, i.e. when speculators as a group have written a lot of puts and bought a lot of calls as a share of their total overall positions.

A key result is visible in Column 5, containing both demand pressure variables for the options market. Notably, OSP dominates OHP both for implied skewness as well as the risk premium, indicating that it matters who and how the hedging pressure is absorbed. Column 6 conducts a robustness check splitting OSP conditional on its sign. While for implied skewness, a difference in slope is discernible, this effect vanishes for the premium. Table 2.5 thus supports the findings depicted in the scatter plots and alleviates concerns about an omitted variable that drives both average trader positions and skewness. The findings also confirm hypothesis (H1).

In unreported results, I also test a slightly modified version of OSP, which is scaled not by total speculative positions, but by total open interest. This version does generally perform equally well, but importantly it is dominated by the original OSP when explaining the skewness risk premium.

2.5.3 Unconditional Price Pressure

The literature on hedging pressure in futures (De Roon et al., 2000; Szymanowska et al., 2011) is careful to delineate permanent effects of the level of hedging demand from temporary effects from changes in hedging demand, called price pressure. Equal care should be taken in the present context to understand if temporary effect from changes in the positions of traders drive option prices. Table 2.6 directly compares the effects of the level with those from changes in demand for options. $\Delta\text{OHP}_{t-m,t}$ ($\Delta\text{OSP}_{t-m,t}$) represent the change in each measure over the last $m \in \{1, 3\}$ months.

Controlling for the level, changes in hedging demand affect neither the price of skewness nor the premium. Surprisingly, changes in the net exposure of speculators does have a significant effect, at least for 1-month changes. As with levels, the effect due to changes in OSP dominates the OHP-related measure and the sign is notably in the opposite direction as for levels.

The results in Table 2.6 suggest that as speculators as a group buy call options and/or sell put options the price of former temporarily increases and the price of the latter decreases. The comparison of the two time horizons shows that the effect does fade with time. This seems to suggest that shifts in demand from speculators have a significant impact and that market makers appear to be slow in accommodating these longer-term order imbalances, which in turn provides additional evidence for the existence of limits to market making.

2.5.4 Reversal vs. Continuation of Demand

Section 2.5.3 does not directly test hypotheses H2 and H3, because the price pressure variables are computed unconditionally. Instead, we need to ask whether a liquidation event is currently taking place. Given the low frequency of the position data (weekly) and end-of-day mid-quotes in options, this is not directly possible.

However, one may be able to identify forced selling indirectly by conditioning on the previously existing sign and level of demand. A group of traders is almost certainly not in liquidation mode if they are observed to add to their positions. Seeing a reversal in the size of positions, however, can occasionally be due to forced selling. This is motivated by the literature on currency carry trades (e.g. Brunnermeier et al., 2008) where speculators add to their positions in a high-yielding currency in a slow and orderly fashion to avoid price impact, until their positions become rather one-sided. At that point any small external shock (bad macroeconomic news, central bank intervention, margin calls) sends them collectively scrambling for the exits. As in the model of Bernardo and Welch (2004), the fear of such a liquidity driven run alone may be cause enough and this fear is largely a function of the imbalance in speculators positions. To test this in the present context, I split the price pressure variables into 2 parts, conditional on whether the change in positions adds to existing exposure by that trader group, i.e. represents a continuation of a trend, or if the change subtracts from existing positions, i.e. represents a reversal.

Table 2.7 presents those results. Quite remarkably, in the case of changes in OSP, I find that the price pressure effect documented in Table 2.6 is rather asymmetric and entirely concentrated in reversions of speculative demand rather than continuations. This nicely fits the aforementioned analogy with carry trades and is direct evidence that at least occasionally liquidations of speculative positions are accompanied by a significant price impact, which in turn once again suggests that market makers face limitations in the size of order imbalance they are willing or able to accommodate in the short run. At the 3-month horizon, the effect has vanished. The effects due to the level of speculative exposure and those due to a reversal in this exposure push prices the same way. As speculators attempt to cover short positions in puts and sell calls prices move against them making puts even more expensive and calls even cheaper.

Note that these results do not suggest that every time speculators exit existing positions, it happens in a disorderly fashion with price impact. Most of the time, it is in the best interest of rational financial traders to keep the price impact to a minimum. But in order to get to the observed *average* effect in Column 2, the less often forced liquidations occur, the more powerful they must be when they happen. This in a nutshell is what the risk emanating from forced liquidations of speculative positions is about. Given a certain level of speculative net long exposure, holder of net long positions know that with a certain probability they may be forced to sell at temporarily depressed prices and incorporate this probability into the price they are willing to enter the position to begin with.

2.5.5 The Effect of Financial Constraints

The recent literature has shown that financial intermediaries, and as a result asset prices, are sensitive to changes in their ease to access funding, i.e. credit (Adrian and Shin, 2010; Adrian et al., 2011b). 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 or market making in the context of commodity futures. To test, financial constraints and their effect on option skewness, I use two of the measures they propose, namely the year-over-year change in the effective risk aversion $d\hat{\Phi}$ of broker-dealers 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), both on a 12-month rolling basis. The data are extracted from the quarterly flow of funds database of the Federal Reserve Board. For both measures, higher values mean a deterioration in financial conditions. In a related study, Adrian et al. (2011a) relate the risk premium of the U.S. dollar against other currencies to yet another set of market-wide proxies for the ease of funding of financial intermediaries. Of those I pick two, namely the year-over-year change in the net issuance of commercial paper and (the 3-month moving average of) the log ratio of bond issuance of financial firms relative to non-financial firms. Adrian et al. (2011a) interpret the former as a measure of short-term funding liquidity and the latter as one of medium and longer term funding liquidity. For these two, higher values represent improvement in general financial conditions.

The first 4 columns of Table 2.8 show how skewness and premium are affected by these four proxies for the funding ability of financial intermediaries. At least in the case of first two, I find that deteriorating balance sheets of broker-dealers seem to affect implied skewness negatively, i.e. puts rise in price relative to calls, and the skewness risk premium positively by about the same magnitude. 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 has historically not been the case for commodities. Almost all of the price effect translates into a premium.

Rather than in the pure effects from changes in funding liquidity, I am more interested whether the effect of OSP is amplified under deteriorating conditions testing hypothesis (H4). This is what the next 4 columns in Table 2.8 try to answer by including the interaction terms between proxies of financial constraints and OSP.

I find that recent increases in financial constraints for broker-dealers and financial institutions more generally amplify the effect of OSP in the case of implied skewness. This makes intuitively sense, as one would expect that market makers price in the increased likelihood of liquidation events in the options market for given level of speculative net exposure since precisely deteriorations in the ease of funding are prime reasons for speculators, mostly financial institutions and hedge funds themselves, to be forced to liquidate positions. When it comes to the effect on options net of the realized skewness, i.e. the skewness premium, evidence for a conditional increase in premium are hard to find. It appears that half or more of the magnifying effect present in implied skewness also affects the realized skewness, leaving only insignificant amounts to the premium. Thus, while prices are affected in intuitive ways, the results for the skewness risk premium and thus option returns are negative. I will return to this question in Section 2.6, where I do obtain evidence that option returns are indeed significantly impacted by the the interaction between OSP and funding constraints.

2.5.6 The Effect of Trader Concentration

Finally, in order to test the last hypothesis (H5) is to consider the proportion of net long (short) open interest that is concentrated in the hands of the largest N traders in that market after netting across all futures and options positions individually for each trader. When much of net open interest is concentrated among the largest traders it suggests that those large traders may be near or at their limits with regards to the exposure they are allowed or willing to have and that they are more prone to changes in funding liquidity.

Table 2.9 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 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, puts are particularly expensive relative to calls reflecting the higher likelihood of forced selling coming from that side, or the lack of additional capacity by market makers holding concentrated and unbalanced positions on that side.

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 4 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. Given that the effect from changes in short concentration is still about as strong for 3-month changes as for 1-month changes would indicate that it takes a significant amount of time for prices to adjust back to normal.

Overall, I find some support for a permanent effect for concentration on the long side that supports a risk-based explanation such as liquidation risk. On the short side, the effect seems to be slowly decaying after a change in trader concentration, which would indicate that a slow adjustment of capital is taking place. Both results, however, clearly show that the concentration of traders matters for both the pricing and returns of options.

2.6 Returns to Portfolio Strategies

The results in Section 2.5 provide evidence that liquidation risk and other limits to market making affect both the price of skewness as well the skewness risk premium over time in commodities. In this section, I want to focus on two of the strongest effects. Speculative option net exposure (OSP) and the long net concentration ratio (CRL) have predictive power for implied skewness (negatively) and the skewness risk premium (positively), but seem not to affect realized skewness. This suggests that in times when current exposure in options by speculators is positive and concentration of long futures positions is high, OTM put options will yield lower returns and OTM calls relatively higher returns on average.

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 market making? 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 'OSP' 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 in options. Since his positions would align with those of the existing speculators and/or market makers, he essentially represents the marginal investor on the speculators' side. As such, he should be particularly vulnerable to shocks to the funding constraints of financial traders. If those investors on the same side of the trade as himself are forced to liquidate, his position will suffer accordingly and may be subject to margin calls. Thus, the return from such a strategy exhibit features of carry trade returns in that the traders wins if no liquidation or reversal of positions occurs. Alternatively, one can consider such a strategy as one of liquidity provision, as the ultimate cause of the current price of skewness is that a given hedging demand is met by insufficient liquidity on the market makers'/speculators' side.

2.6.1 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 in practice and further, given the large 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 re-balancing in the underlying futures contract. If one chooses instead to re-balance daily, the floating leg, i.e. the realized skewness, becomes very noisy for short periods of time. After all, in Section 2.5 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.[19]

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 (see e.g. Brunnermeier et al., 2008). 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.[20]

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 one 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 the criterion, OSP or CRL respectively. The bottom third receives a negative, the top a positive and the middle a zero weight. Obviously, more intricate weighting schemes are possible, but I will focus in the most simple ones here.

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.[21] 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 requirements[22].

On the part of the option positions, I assume that the required margin is 300 percent of the option that is written for the .25-delta risk reversal. 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.[23]

[19] In unreported tests, I do form portfolios of skewness swaps, albeit without intermittent re-balancing. The signs and significances do line up with the results reported below.

[20] Bali and Murray (2012) 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.

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

[22] http://www.cmegroup.com/clearing/margins/ provides current information on futures margins on the CME.

[23] I used http://www.cboe.com/tradtool/mcalc/default.aspx for the example calculations.

2.6.2 Construction of Factor-Mimicking Portfolios

One of the goals of the portfolio formation is to gage potential returns in excess of known risk factors. As mentioned previously, it is likely that the strategies are sensitive to the availability of capital to arbitrageurs or the financial system in general and subsequent effects on funding liquidity and asset liquidity (Brunnermeier and Pedersen, 2009). In this context, I investigate if a number of proxies known to be related to capital abundance in the financial system can partially explain the strategies' returns. I use monthly changes in the U.S. Treasury-derived measure of arbitrage capital by Hu et al. (2010) as well two of the balance sheet measures used in Adrian et al. (2011a) to explain the risk premium of the U.S. dollar, namely the (year-over-year) change in U.S. commercial paper outstanding[24] and the log ratio of financial bond issuances relative to non-financial bond issuances at the monthly frequency.

To keep the meaning of the intercept term intact as a measure of performance, I report results based on factor-mimicking portfolios, i.e. I replace the factors with achievable returns from tradeable assets. These are constructed as follows. Every month t, I regress the returns of each common stock in the CRSP universe over the previous 36 months on the 4 Fama-French factors and one of the liquidity factors F_t.

$$R_{i,\tau} = \gamma_0 + \beta_{i,t}^F F_t + \beta_{i,t}^M R_\tau^M + \beta_{i,t}^{SMB} R_\tau^{SMB} + \beta_{i,t}^{HML} R_\tau^{HML} + \beta_{i,t}^{UMD} R_\tau^{UMD} + \epsilon_{i,\tau} \tag{2.15}$$

where $\tau \in t-36, \ldots, t-1$. This yields a factor sensitivity for each stock, i.e. its pre-ranking beta $\hat{\beta}_{i,t}^F$. In month t, stocks are sorted into ten deciles according to their pre-formation beta and returns are value-weighted to yield a continuous time series of returns for each decile. For simplicity, I choose the factor-mimicking portfolio as the strategy that is long the stocks in decile 10 and short those in decile 1. Alternatively, one could let the data decide how much weight to put on each of the deciles, performing a regression that uses the factor on the left and the 10 decile returns on the right.

For robustness, I conduct two tests for each factor to make sure the portfolio formed in the above manner does indeed pick up some essential part of the risk factor. First, I confirm that post-ranking betas of the 10 deciles are meaningfully different from each other and monotonic. Second, I regress the factor itself on its mimicking portfolio, once with and once without the control factors. I find the R-squares of these time series regressions to be remarkable high, on average on the order of 50 percent. This gives me confidence that the factor-mimicking portfolios contain much more than just noise.

2.6.3 Portfolio Results

Using the method described above, I construct a time series of portfolio returns based on the key variables that measure some aspect of limits to market making. Table 2.10 shows those results for

[24] Adrian et al. (2011a) use commercial paper of financial firms only, while I used all commercial paper due to data limitation. For the time period where both were available to me, both time series tracked each other closely.

the net exposure of speculators in options (OSP); table 2.11 shows the same results for the trader concentration on the long side based on combined net positions (CRL). The analysis is done for 2 levels of delta, $\Delta = .25$ and $\Delta = .10$. Using monthly re-balancing there are 2 overlapping portfolios at any point in time. For each commodity, the risk reversals are based on an expiration that lies between 2 and 4 months after first formation.

A few things jump out. First, for both conditioning variables, the raw returns are higher when using options that are farther out of the money, presumably owing to their increased leverage and thus exposure to the difference in skewness. Second, the standard equity risk factor do not contribute much to explaining the portfolio returns, on the order of 2 to 3 percent in each case. The OSP-based strategy loads weakly on the market return and possibly HML, while CRL-based returns appear to be related to momentum in equity returns and, again weakly, the market.

Table 2.10 shows that the raw returns for the OSP-based strategies amount to up to 2.5 percent per month for lower delta options and about half that for the more conservative choice of option delta. Risk-adjusted returns are slightly lower, but still strongly significant even when controlling for the funding liquidity-related risk factors. Where those factor loadings are significant they all have the intuitively correct sign, pointing towards a positive exposure to aspects of funding liquidity. Note that by construction, higher values of HPW (or its factor-mimicking portfolio) signify times of lower funding liquidity. A positive loading thus means that the return on the strategy is higher during times of lower arbitrage capital abundance. The other 2 factors are concerned with the amount of funding available to (financial) firms and higher values mean better funding liquidity for those institutions. Thus, a negative loading in these cases reinforces the result for the HPW factor. Maximum R-square reaches a respectable 11 percent.

Table 2.11 reveals raw returns of between 1.3 and 1.6 percent for strategies based on CRL. As with OSP, the liquidity factors all receive loadings with the correct sign and are almost always significant. Risk-adjusted returns fall to around 1 percent per month and R-square reaches a maximum of almost 14 percent.

Taken together, we see strong evidence that the skewness risk premium is in part 'earned' by being exposed to liquidity-related risk factors. 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 by arbitrageurs. It works in favor of a non-zero α, however, that in a delta-hedged risk reversal, the proportion of capital actually 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 only. Secondly, effective spreads in option markets tend to much lower than posted spreads seem to indicate.

Further, the strategy does not fully reflect the situation that market makers are in. They enter positions by earning half the spread and can either hold it, properly hedged as part of their larger book until maturity, or exit the position again earning half the spread. The main point of this analysis, however, is not to prove that these strategies maintain a significant α in a highly stylized simulation, rather that they appear to load on risk factors that are associated with the general health of the financial intermediary sector. While I was unable to find robust evidence in Section 2.5.5 that the skewness premium, and thus option returns, are sensitive to the interaction between OSP and funding constraints, the results here do support this hypothesis.

2.7 Conclusion

Asset pricing models have some difficulty to rationalize the abnormal returns that have been found for a number of strategies involving equity 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 biased beliefs, a small strand of the literature searches for alternative explanations in the micro structure of the options market, more specifically, in the process of financial intermediation fulfilled by market makers and arbitrageurs more generally (Bates, 2003). Those studies provide evidence that the market making sector faces capital constraints and as a consequence of upward sloping supply curves, option prices diverge from frictionless no-arbitrage prices (Bollen and Whaley, 2004).

Using a newly proposed measure of realized skewness (Kozhan et al., 2011), unlike the previous literature I am able to a) exclude informed demand and adverse selection as alternate cause of the price effect and b) link *option returns* directly to demand effects for the first time. Further, analyzing the market for commodity futures options, I find persistent price effects of hedging demand on options that depend not purely on the size of hedging demand, but instead on the positions of the financial traders that accommodate this demand. A larger net long exposure of speculators in options leads to more negative implied skewness, i.e. puts being more expensive and calls being less expensive, and a positive skewness risk premium, i.e. lower returns to puts and higher returns to calls. A higher concentration of net holdings among traders on the long side has the same effect.

I attribute these effects to an increase in liquidation risk, i.e. the possibility that financial traders are forced to liquidate their holdings following an external liquidity-related shock such as changes in margins, losses in other markets or worsening funding conditions, causing temporarily adverse price movements comparable to liquidity runs (Bernardo and Welch, 2004) and carry trade unwinds (Brunnermeier et al., 2008). I also discover evidence of the impact of episodes of forced selling in option prices and returns.

In the final part of the paper, I construct portfolios that are aimed at theoretically capturing the skewness risk premium and find monthly raw returns of 2.5 and 1.3 percent, respectively. Lastly, using factor-mimicking portfolios I show that parts of the returns are compensation for exposure to market-wide funding liquidity.

Table 2.1: Sample Overview

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 90 days. The next 3 columns depict the interquartile range of the main variable of interest taken from the COT reports, OSP, the long net exposure of speculators (delta-weighted and scaled by speculators' total open interest). The last column shows the time series correlation of OSP with the corresponding measure for the underlying futures market.

Commodity	**Futures Options Data**			**OSP Measure**			
	Begin	**End**	**#Obs**	**Q1**	**Q2**	**Q3**	ρ_{SP}
Agricultural							
Soybean Oil	1989/07	2010/09	139	-39%	-21%	-9%	-3%
Corn	1989/07	2010/09	108	-27%	-12%	-4%	2%
Oats	1991/03	2010/07	78	-45%	-15%	7%	11%
Rough Rice	1992/09	2010/09	108	-18%	-3%	15%	9%
Soybeans	1989/07	2010/09	150	-13%	-3%	2%	-44%
Soybean Meal	1989/07	2010/09	164	-25%	-10%	5%	-6%
Wheat (CBOT)	1989/07	2010/09	107	-9%	2%	10%	-25%
Energy							
Crude Oil (WTI)	1990/01	2010/09	248	5%	9%	15%	3%
Heating Oil No. 2	1990/01	2010/09	219	-8%	-3%	1%	-20%
Unl. Gasoline	1990/04	2006/12	192	-5%	1%	10%	-22%
Natural Gas	1993/02	2010/09	204	-4%	0%	5%	4%
RBOB Gasoline	2007/03	2010/09	42	-2%	0%	3%	34%
Meat							
Feeder Cattle	1987/04	2010/08	186	23%	37%	49%	29%
Live Cattle	1991/04	2010/08	117	28%	39%	56%	37%
Lean Hogs	1997/02	2010/08	104	17%	25%	41%	43%
Pork Bellies	1987/02	2008/02	79	12%	25%	30%	68%
Metal							
Gold (NYMEX)	1989/06	2010/08	128	1%	7%	12%	12%
Copper (HG)	1990/05	2007/08	157	-19%	-5%	13%	24%
Silver (NYMEX)	2002/03	2010/09	44	-5%	1%	7%	50%
Soft							
Cocoa	1990/09	2010/09	147	-17%	-8%	-1%	19%
Cotton No. 2	1990/07	2010/09	112	0%	7%	16%	-26%
Orange Juice	1990/09	2010/09	144	-22%	-3%	13%	2%
Coffee C	1990/09	2010/09	162	-7%	-3%	3%	46%
Lumber	1987/11	2010/09	117	-14%	15%	35%	51%
Sugar No. 11	1990/07	2010/09	160	-1%	11%	23%	0%

Table 2.2: Volatility and Variance Risk Premium

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.

Commodity	# Obs.	Implied Vol.	Realized Vol.	Variance Premium avg.	Variance 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]

Table 2.3: Implied vs. Realized Skewness

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.

Commodity	**#**	**Implied**		**Realized**		**Skewness Premium**	
	Obs.	**BKM**	**MFIS**	r^3	**RSkew**	**avg.**	**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]

Table 2.4: Control Variables and Skewness

This table shows the results from a panel regression of implied skewness, realized skewness and the skewness risk premium on two key variables and some controls. 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 $\text{RSkew}^{1/3}$, the signed third root of realized skewness (as in Neuberger, 2011). In panel C, the dependent variable is the difference between $\text{RSkew}^{1/3}$ and $\text{MFIS}^{1/3}$ averaged over the same range of maturities. Conv. Yield(6m), Δ OI(6m), $\text{Ret}_{(t-3,t)}$ and $\text{RSkew}^{1/3}_{t-m}$ are the 6-month convenience yield, the change in futures open interest (relative to 6 months ago), the recent 3-month return of the underlying future and the most recent realized skewness based on a 3-month maturity, respectively. OSP is the scaled, delta-weighted net long exposure of speculators in options. logCRL is the log of the net concentration ratio of the largest traders being net long in both futures and options. Standard errors are clustered by month and commodity following Thompson (2011).

	Panel A: Implied Skewness				Panel B: Realized Skewness				Panel C: Skewness Risk Premium			
Conv. Yield (6m)	-0.122	-0.099	-0.161	-0.133	0.024	0.027	0.021	0.026	0.156	0.137	0.192	0.169
	[-0.70]	[-0.61]	[-0.90]	[-0.79]	[0.19]	[0.22]	[0.17]	[0.21]	[1.09]	[1.01]	[1.27]	[1.19]
Δ OI (6m)	0.133	0.132	0.117	0.118	0.217	0.216	0.216	0.216	0.070	0.072	0.086	0.084
	*[1.76]	*[1.82]	[1.66]	*[1.76]	***[2.81]	***[2.80]	***[2.84]	***[2.84]	[0.72]	[0.75]	[0.91]	[0.91]
$\text{Ret}_{(t-3,t)}$	-0.081	0.001	-0.109	-0.028	0.250	0.263	0.249	0.262	0.309	0.240	0.335	0.268
	[-0.37]	[0.01]	[-0.50]	[-0.13]	[1.51]	[1.56]	[1.45]	[1.48]	[1.64]	[1.30]	*[1.73]	[1.40]
$\text{RSkew}^{1/3}_{t-3}$	0.279	0.274	0.277	0.274	0.160	0.159	0.160	0.159	-0.114	-0.111	-0.113	-0.110
	***[5.14]	***[5.10]	***[5.00]	***[4.99]	***[3.33]	***[3.31]	***[3.31]	***[3.30]	***[-3.01]	***[-2.95]	***[-3.02]	***[-2.96]
OSP		-0.507		-0.468		-0.079		-0.077		0.423		0.387
		***[-3.98]		***[-3.88]		[-0.88]		[-0.81]		***[3.20]		***[2.93]
logCRL			-0.243	-0.197			-0.016	-0.008			0.222	0.184
			***[-3.04]	**[-2.30]			[-0.16]	[-0.08]			***[4.10]	***[3.29]
R^2	8.7%	11.5%	10.0%	12.3%	4.0%	4.0%	4.0%	4.0%	2.1%	3.6%	3.0%	4.2%
# commodities	25	25	25	25	25	25	25	25	25	25	25	25
# months	183	183	183	183	183	183	183	183	183	183	183	183
Total # Obs.	2,570	2,570	2,570	2,570	2,570	2,570	2,570	2,570	2,570	2,570	2,570	2,570

Table 2.5: Level of Demand Pressure and Skewness

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 $\text{MFIS}^{1/3}$, 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). HP (SP) is the scaled net short (long) exposure of hedgers (speculators) in futures. OHP (OSP) is the scaled net short (long) exposure of hedgers (speculators) in options, delta-weighted. OSP+ (OSP−) is OSP conditional on it being positive (negative). Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Panel A: Implied skewness						
HP	0.096					
	[0.59]					
SP		0.068				
		[0.61]				
OHP			-0.249		-0.003	
			**[-2.28]		[-0.04]	
OSP				-0.473	-0.470	
				***[-3.92]	***[-4.61]	
OSP+						-0.716
						**[-2.76]
OSP-						-0.236
						[-1.42]
R^2	10.0%	10.0%	11.1%	12.2%	12.2%	12.5%
Panel B: Skewness risk premium						
HP	-0.104					
	[-1.14]					
SP		-0.012				
		[-0.15]				
OHP			0.175		-0.056	
			**[2.11]		[-0.71]	
OSP				0.388	0.442	
				***[2.93]	**[2.72]	
OSP+						0.424
						**[2.56]
OSP-						0.352
						**[2.52]
R^2	3.0%	3.0%	3.4%	4.2%	4.2%	4.2%
# commodities	25	25	25	25	25	25
# months	183	183	183	183	183	183
Total # obs	2,587	2,587	2,587	2,587	2,587	2,587

Table 2.6: Price Pressure and Skewness

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 $\text{MFIS}^{1/3}$, 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). OHP (OSP) is the scaled net short (long) exposure of hedgers (speculators) in options, delta-weighted. $\Delta\text{OHP}_{t-m,t}$ ($\Delta\text{OSP}_{t-m,t}$) is the change in OHP (OSP) over the last m months. Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Panel A: Implied skewness						
	m = 1 month			**m = 3 months**		
OHP	-0.259		0.059	-0.259		0.063
	**[-2.09]		[0.57]	*[-1.95]		[0.57]
OSP		-0.515	-0.572		-0.509	-0.572
		***[-3.99]	***[-5.02]		***[-3.56]	***[-4.44]
$\Delta\mathbf{OHP}_{t-m,t}$	0.046		-0.112	0.024		-0.078
	[0.72]		[-1.57]	[0.44]		[-1.29]
$\Delta\mathbf{OSP}_{t-m,t}$		0.157	0.26		0.066	0.139
		**[2.10]	***[3.03]		[0.83]	[1.50]
R^2	11.0%	12.3%	12.3%	11.0%	12.2%	12.3%
Panel B: Skewness risk premium						
	m = 1 month			**m = 3 months**		
OHP	0.207		-0.074	0.251		-0.009
	**[2.29]		[-0.80]	**[2.51]		[-0.08]
OSP		0.435	0.506		0.451	0.458
		***[2.93]	**[2.58]		**[2.65]	*[2.05]
$\Delta\mathbf{OHP}_{t-m,t}$	-0.073		0.082	-0.103		-0.033
	[-1.25]		[0.86]	*[-1.95]		[-0.46]
$\Delta\mathbf{OSP}_{t-m,t}$		-0.19	-0.266		-0.119	-0.085
		**[-2.07]	*[-1.75]		[-1.54]	[-0.78]
R^2	3.6%	4.4%	4.4%	3.7%	4.3%	4.4%
# commodities	25	25	25	25	25	25
# months	179	179	179	179	179	179
Total # obs	2,538	2,539	2,538	2,538	2,539	2,538

Table 2.7: Reversal and Continuation of Demand and Skewness

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 $\text{MFIS}^{1/3}$, 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). OHP (OSP) is the scaled net short (long) exposure of hedgers (speculators) in options, delta-weighted. $\Delta\text{OHP}_{t-m,t}$ ($\Delta\text{OSP}_{t-m,t}$) is the change in OHP (OSP) over the last m months; changes are split in two variables conditional on whether the change adds to existing net positions or reverses from previous net positions. Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Panel A: Implied skewness						
	m = 1 month			**m = 3 months**		
OHP	-0.269		0.039	-0.257		0.047
	**[-2.16]		[0.41]	*[-1.97]		[0.43]
OSP		-0.484	-0.534		-0.493	-0.548
		***[-3.03]	***[-3.84]		***[-3.08]	***[-3.63]
ΔOHP$_{t-m,t}$ (rev.)	0.039		-0.117	0.025		-0.078
	[0.69]		*[-1.73]	[0.44]		[-1.28]
ΔOHP$_{t-m,t}$ (cont.)	0.090		-0.046	0.016		-0.019
	[0.66]		[-0.43]	[0.18]		[-0.21]
ΔOSP$_{t-m,t}$ (rev.)		0.184	0.281		0.067	0.140
		***[3.27]	***[3.64]		[0.85]	[1.52]
ΔOSP$_{t-m,t}$ (cont.)		0.035	0.128		0.008	0.056
		[0.15]	[0.60]		[0.04]	[0.29]
R^2	11.0%	12.3%	12.4%	11.0%	12.2%	12.3%
Panel B: Skewness risk premium						
	m = 1 month			**m = 3 months**		
OHP	0.230		-0.020	0.245		0.018
	**[2.21]		[-0.20]	**[2.74]		[0.21]
OSP		0.354	0.399		0.385	0.379
		**[2.16]	*[1.88]		**[2.36]	*[1.81]
ΔOHP$_{t-m,t}$ (rev.)	-0.057		0.093	-0.103		-0.034
	[-1.16]		[1.00]	*[-1.95]		[-0.48]
ΔOHP$_{t-m,t}$ (cont.)	-0.169		-0.090	-0.082		-0.120
	[-1.10]		[-0.45]	[-0.61]		[-0.69]
ΔOSP$_{t-m,t}$ (rev.)		-0.260	-0.327		-0.124	-0.090
		**[-2.69]	**[-2.12]		[-1.56]	[-0.81]
ΔOSP$_{t-m,t}$ (cont.)		0.126	0.108		0.111	0.182
		[0.66]	[0.38]		[0.68]	[0.83]
R^2	3.6%	4.5%	4.5%	3.7%	4.4%	4.4%
# commodities	25	25	25	25	25	25
# months	179	179	179	179	179	179
Total # obs	2,516	2,517	2,516	2,516	2,517	2,516

Table 2.8: Financial Constraints and Skewness

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 financial constraints affecting financial intermediaries. 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). OSP is the scaled, delta-weighted net long exposure of speculators in options. $d\hat{\Phi}$ is the year-over-year change in the effective risk-aversion of broker-dealers as in Etula (2010). 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). CP(yoy) is the year-over-year change in the net issuance of commercial paper. Fin. Bond (MA) is the 3 month moving average of the ratio between bond issuances of financial firms relative to non-financial firms. X-term is the interaction term between OSP 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).

				Panel A: Implied Skewness				
OSP	-0.440	-0.438	-0.432	-0.472	-0.431	-0.418	-0.451	-0.482
	***[-3.78]	***[-3.65]	***[-3.67]	***[-3.84]	***[-3.89]	***[-3.70]	***[-4.43]	***[-4.24]
$d\hat{\Phi}$	-0.124				-0.114			
	**[-2.15]				**[-2.11]			
Rel. AG		-0.509				-0.480		
		*[-1.96]				*[-1.96]		
CP (yoy)			0.319				0.284	
			[1.63]				[1.43]	
Fin Bond (MA)				0.028				0.017
				[0.72]				[0.48]
X-term with OSP					-0.304	-1.299	2.031	0.232
					*[-1.83]	*[-1.86]	**[2.74]	*[1.82]
R^2	13.6%	13.4%	12.9%	12.4%	13.9%	13.7%	13.9%	12.6%
				Panel B: Skewness Risk Premium				
OSP	0.362	0.354	0.377	0.376	0.358	0.344	0.384	0.378
	***[2.85]	**[2.73]	***[2.87]	**[2.75]	***[2.93]	**[2.78]	***[2.96]	**[2.68]
$d\hat{\Phi}$	0.114				0.110			
	***[3.06]				***[2.97]			
Rel. AG		0.551				0.535		
		***[3.13]				***[3.12]		
CP (yoy)			-0.088				-0.075	
			[-0.54]				[-0.46]	
Fin Bond (MA)				0.076				0.078
				[1.66]				[1.70]
X-term with OSP					0.122	0.686	-0.721	-0.047
					[0.86]	[1.10]	[-1.07]	[-0.24]
R^2	5.1%	5.3%	4.2%	4.8%	5.1%	5.3%	4.3%	4.8%
# comm.	25	25	25	25	25	25	25	25
# months	183	183	183	183	183	183	183	183
Total # Obs.	2,570	2,570	2,570	2,570	2,570	2,570	2,570	2,570

Table 2.9: Trader Concentration and Skewness

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). logCRL (logCRS) is the log of the share of the largest 8 traders on the long (short) side based on their netted positions in both the futures and options markets among all netted long (short) positions. dCRL (dCRS) is the net change in the long (short) concentration ratio relative to either 1 or 3 months earlier. Coefficients from control variables are omitted. Standard errors are clustered by month and commodity following Thompson (2011).

Panel A: Implied skewness							
logCRL	-0.211		-0.207		-0.210		-0.182
	*[-1.77]		**[-2.10]		*[-2.06]		[-1.71]
logCRS		-0.097	-0.011		-0.031		-0.049
		[-0.68]	[-0.09]		[-0.28]		[-0.40]
dCRL(1m)				-0.469	-0.018		
				[-1.56]	[-0.05]		
dCRS(1m)				0.416	0.506		
				*[1.90]	**[2.08]		
dCRL(3m)						-0.698	-0.337
						**[-2.37]	[-0.97]
dCRS(3m)						0.365	0.448
						[1.67]	[1.69]
$\boldsymbol{R^2}$	12.6%	11.7%	12.6%	11.5%	12.7%	11.9%	12.8%
Panel B: Skewness Risk Premium							
logCRL	0.220		0.197		0.195		0.191
	***[4.65]		***[4.40]		***[4.26]		***[3.63]
logCRS		0.144	0.062		0.103		0.117
		*[1.90]	[0.89]		[1.33]		[1.38]
dCRL(1m)				0.494	0.070		
				[1.69]	[0.24]		
dCRS(1m)				-0.662	-0.874		
				**[-2.31]	**[-2.40]		
dCRL(3m)						0.614	0.232
						**[2.21]	[0.80]
dCRS(3m)						-0.455	-0.639
						*[-1.83]	*[-2.05]
$\boldsymbol{R^2}$	4.6%	4.0%	4.7%	3.9%	5.1%	4.0%	5.2%
# comm.	25	25	25	25	25	25	25
# months	183	183	183	183	183	183	183
Total # Obs.	2,570	2,570	2,570	2,547	2,547	2,521	2,521

Table 2.10: Portfolio Returns based on Net Long Speculator Exposure in Options

This table shows raw and risk-adjusted returns to strategies of going long/short delta-neutral risk reversals (DNRR) in commodity options based on the net long exposure of speculators in commodity futures options. Overlapping portfolios are held for 2 months and rebalanced monthly. The risk reversals are based on options with a delta of .25 or .10, respectively, and delta-hedged at each re-balancing date. The columns β, SMB, HML, UMD contain the coefficients to the Fama-French 4-factor model. The columns HPW, CP(yoy) and Fin. Bond contain the loadings on the corresponding factor-mimicking portfolio returns. All percentage returns reported are scaled to a monthly horizon.

α	β	SMB	HML	UMD	HPW	CP(yoy)	Fin.Bond	N	R^2
2 Monthly Turnovers, Delta = .25									
1.27%								175	0.00%
***[3.44]									
1.13%	0.17	0.02	0.14	0.01				175	3.31%
***[3.00]	*[1.90]	[0.16]	[1.24]	[0.16]					
1.05%	0.19	-0.04	0.11	0.02	0.09			175	5.97%
***[2.79]	**[2.14]	[-0.36]	[1.01]	[0.34]	**[2.18]				
1.10%	0.21	-0.01	0.14	0.03		-0.06		175	4.70%
***[2.92]	**[2.23]	[-0.09]	[1.22]	[0.42]		[-1.57]			
1.03%	0.16	0.08	0.24	0.03			-0.11	175	7.67%
***[2.76]	*[1.79]	[0.75]	**[2.09]	[0.47]			***[-2.82]		
0.92%	0.20	0.00	0.21	0.06	0.10	-0.03	-0.11	175	11.31%
**[2.49]	**[2.19]	[0.03]	*[1.82]	[0.81]	**[2.52]	[-0.75]	***[-2.68]		
2 Monthly Turnovers, Delta = .10									
2.56%								175	0.00%
***[4.40]									
2.38%	0.23	0.05	0.20	0.00				175	2.80%
***[3.98]	[1.64]	[0.25]	[1.11]	[-0.01]					
2.32%	0.25	0.00	0.18	0.01	0.06			175	3.34%
***[3.86]	*[1.73]	[0.01]	[0.99]	[0.07]	[0.98]				
2.32%	0.30	-0.01	0.19	0.04		-0.11		175	4.90%
***[3.90]	**[2.07]	[-0.06]	[1.09]	[0.31]		*[-1.93]			
2.22%	0.21	0.15	0.35	0.03			-0.16	175	6.91%
***[3.77]	[1.53]	[0.82]	*[1.93]	[0.29]			***[-2.73]		
2.12%	0.28	0.05	0.31	0.06	0.08	-0.07	-0.15	175	8.49%
***[3.59]	*[1.91]	[0.26]	*[1.69]	[0.57]	[1.26]	[-1.14]	**[-2.32]		

Table 2.11: Portfolio Returns based on the Long-Side Trader Concentration Ratio

This table shows raw and risk-adjusted returns to strategies of going long/short delta-neutral risk reversals (DNRR) in commodity options based on the long-side concentration ratio in futures and options combined (CRL) based on netted positions. Overlapping portfolios are held for 2 months and rebalanced monthly. The risk reversals are based on options with a delta of .25 or .10, respectively, and delta-hedged at each re-balancing date. The columns β, SMB, HML, UMD contain the coefficients to the Fama-French 4-factor model. The columns HPW, CP(yoy) and Fin. Bond contain the loadings on the corresponding factor-mimicking portfolio returns. All percentage returns reported are scaled to a monthly horizon.

α	β	SMB	HML	UMD	HPW	CP(yoy)	Fin.Bond	N	R^2
2 Monthly Turnovers, Delta = .25									
1.33%								179	0.00%
**[2.44]									
1.18%	0.06	-0.01	-0.02	0.20				179	2.19%
**[2.09]	[0.46]	[-0.06]	[-0.13]	*[1.81]					
1.05%	0.09	-0.10	-0.06	0.22	0.13			179	4.77%
*[1.87]	[0.69]	[-0.58]	[-0.37]	**[2.00]	**[2.17]				
1.16%	0.13	-0.07	-0.02	0.24		-0.11		179	4.44%
**[2.06]	[0.96]	[-0.40]	[-0.15]	**[2.15]		**[-2.02]			
1.08%	0.05	0.07	0.11	0.23			-0.13	179	5.16%
*[1.92]	[0.37]	[0.40]	[0.60]	**[2.09]			**[-2.33]		
0.92%	0.14	-0.08	0.04	0.27	0.15	-0.07	-0.12	179	9.33%
*[1.66]	[0.99]	[-0.46]	[0.26]	**[2.52]	**[2.47]	[-1.37]	**[-1.98]		
2 Monthly Turnovers, Delta = .10									
1.63%								179	0.00%
***[3.33]									
1.45%	0.17	0.10	0.02	0.13				179	2.17%
***[2.85]	[1.41]	[0.63]	[0.11]	[1.35]					
1.25%	0.22	-0.04	-0.04	0.16	0.20			179	9.55%
**[2.54]	*[1.85]	[-0.26]	[-0.30]	*[1.70]	***[3.76]				
1.43%	0.22	0.05	0.01	0.16		-0.08		179	3.70%
***[2.83]	*[1.79]	[0.34]	[0.10]	[1.63]		*[-1.66]			
1.35%	0.16	0.16	0.13	0.16			-0.12	179	4.95%
***[2.69]	[1.34]	[1.08]	[0.81]	[1.61]			**[-2.25]		
1.13%	0.25	-0.01	0.06	0.21	0.22	-0.05	-0.12	179	13.87%
**[2.33]	**[2.03]	[-0.05]	[0.40]	**[2.19]	***[4.09]	[-1.07]	**[-2.21]		

Figure 2.1: Exchange-Traded vs. Over-The-Counter Derivatives

This figure compares the exchange-traded open interest and over-the-counter (OTC) notional amounts for commodity derivatives. The left graph shows the aggregate open interest (in $B) of the 25 exchange-traded commodities in my sample. The data is derived from the CFTC commmitment 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 graph 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.

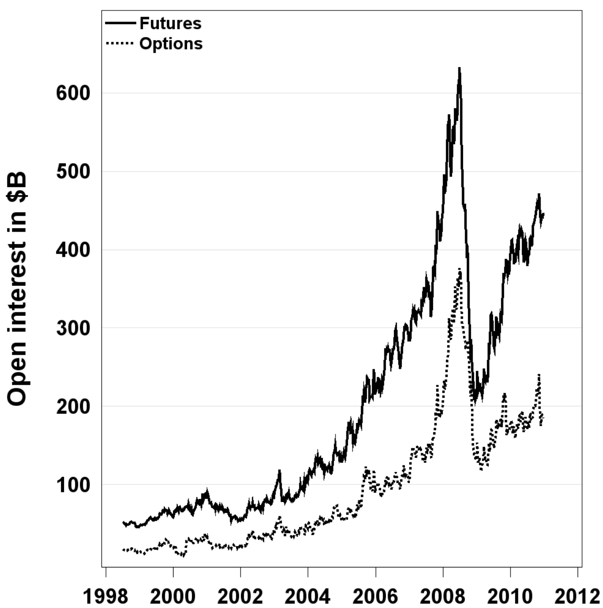

(a) *Exchange-traded Commodity Derivatives*

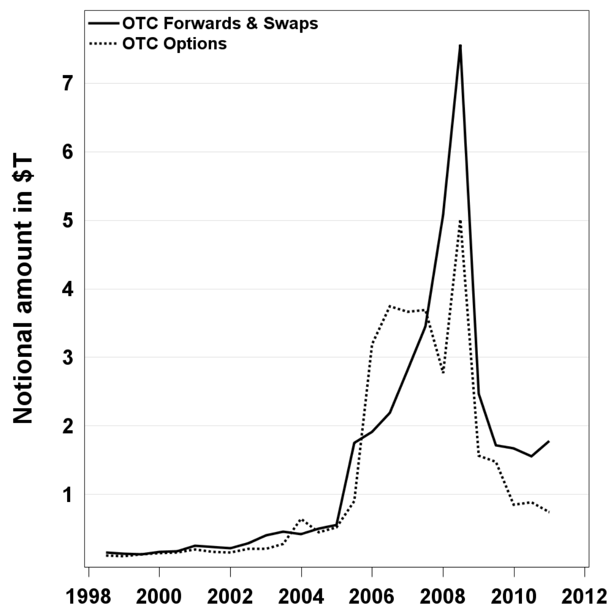

(b) *OTC Commodity Derivatives*

Figure 2.2: Skewness Measures vs. Hedging Pressure

This figure contains six scatter plots that plot the time series average of one of six skewness measures on the Y-axis against the time series average of commercial net options positions over total commercial positions on the X-axis for the cross-section of the commodities in the sample. The first two measures are implied skewness (as defined by Bakshi et al. (2003)) and the price of a risk reversal (in percentage of implied volatility). The next 3 measures are implied skewness, realized skewness and the skewness risk premium (all as defined by Kozhan et al. (2011)). The bottom right graph contains the (Newey-West adjusted) t-statistic of the skewness risk premium (bottom center). The data is based on options data with a remaining maturity of between 90 to 99 days.

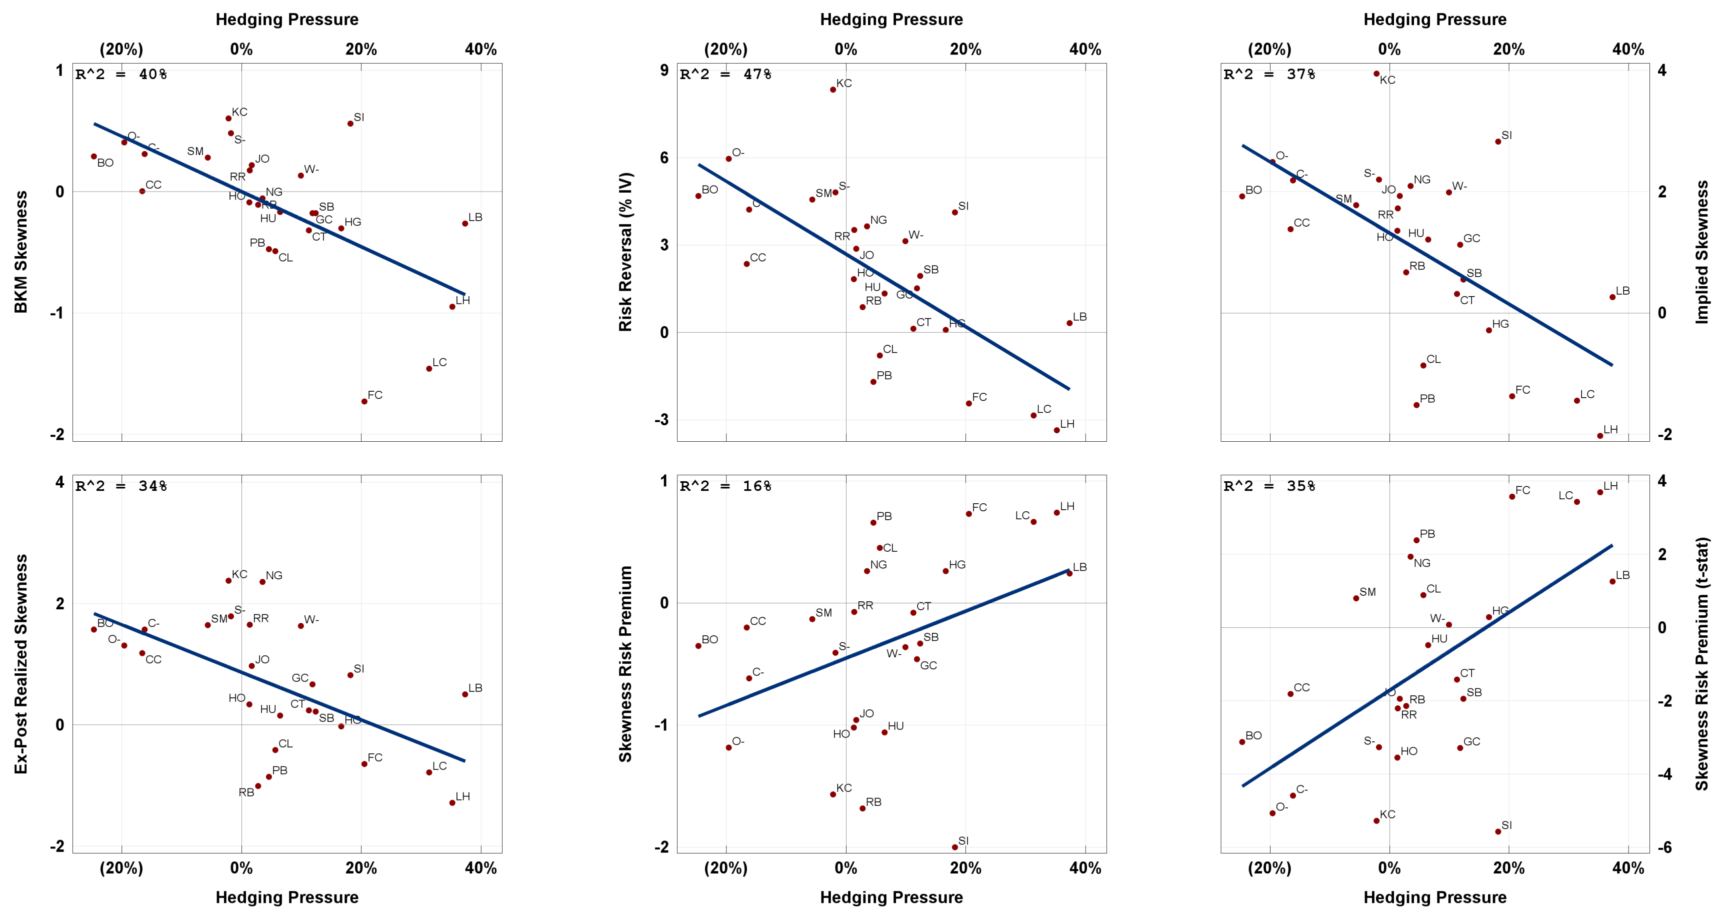

Figure 2.3: Skewness Measures vs. Speculative Positions

This figure contains six scatter plots that plot the time series average of one of six skewness measures on the Y-axis against the time series average of speculative net long options positions over total speculative positions on the X-axis for the cross-section of the commodities in the sample. The first two measures are implied skewness (as defined by Bakshi et al. (2003)) and the price of a risk reversal (in percentage of implied volatility). The next 3 measures are implied skewness, realized skewness and the skewness risk premium (all as defined by Kozhan et al. (2011)). The bottom right graph contains the (Newey-West adjusted) t-statistic of the skewness risk premium (bottom center). The data is based on options data with a remaining maturity of between 90 to 99 days.

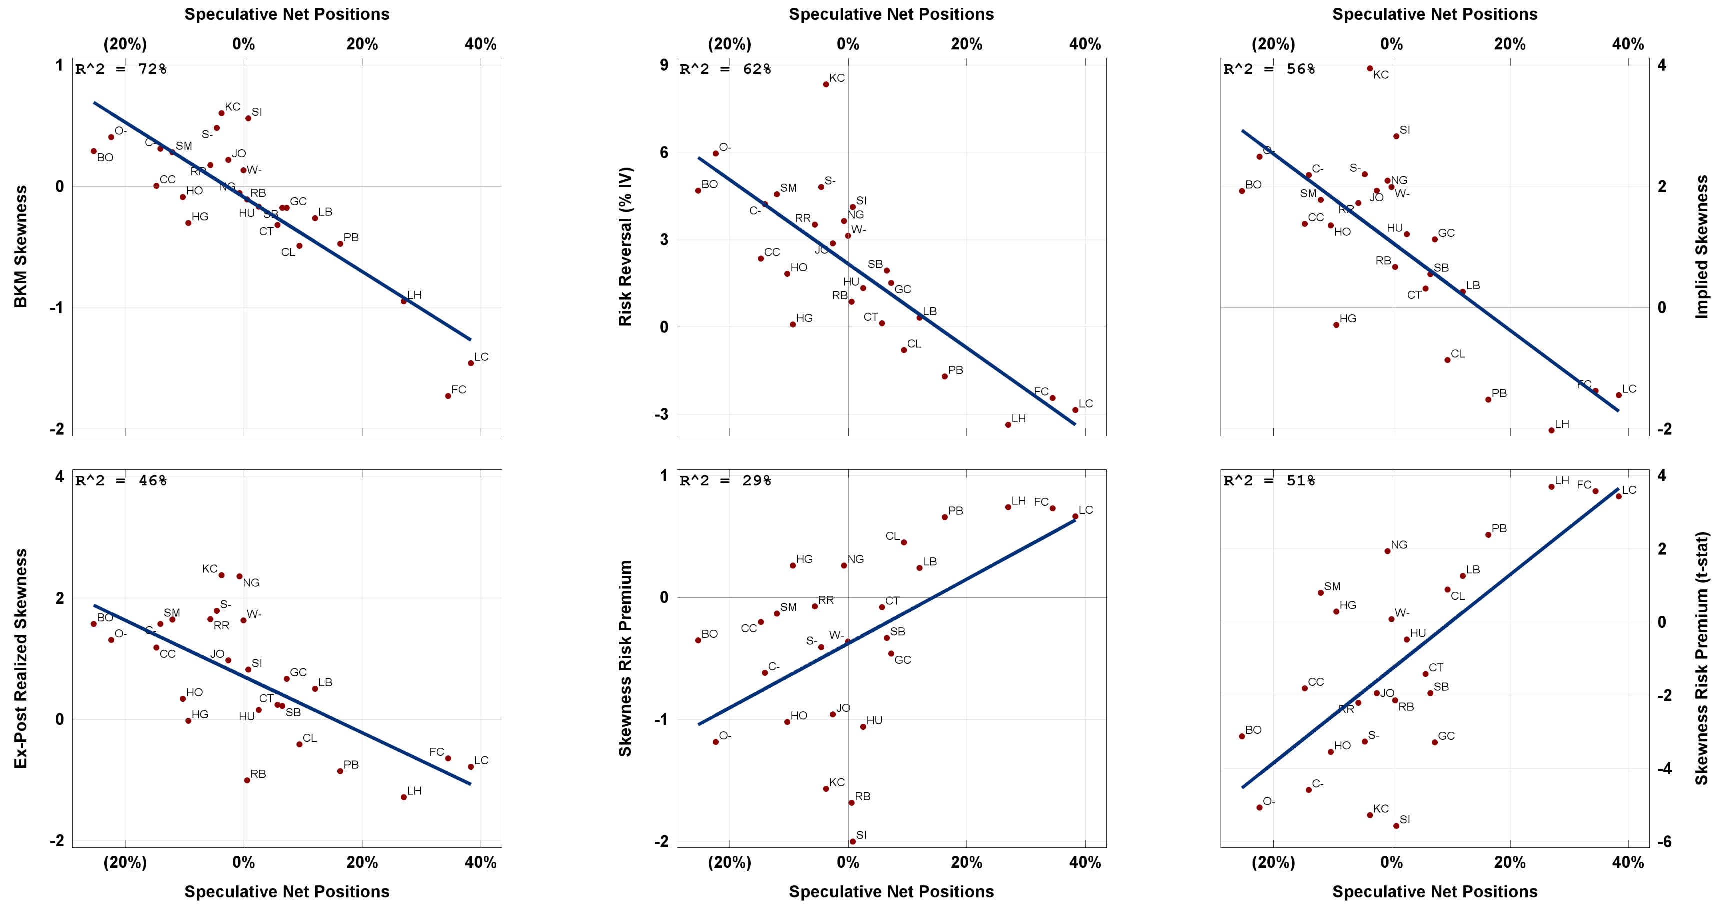

Figure 2.4: Realized Skewness vs. Other Skewness Measures

This figure contains five scatter plots that plot the time series average of one of five skewness measures on the Y-axis against the time series average of realized skewness on the X-axis for the cross-section of the commodities in the sample. The first two measures are implied skewness (as defined by Bakshi et al. (2003)), the price of a risk reversal (in percentage of implied volatility). The next 2 measures are implied skewness and the skewness risk premium (both as defined by Kozhan et al. (2011)). The bottom right contains the (Newey-West adjusted) t-statistics of the skewness risk premium (bottom center). The data is based on options data with a remaining maturity of between 90 to 99 days.

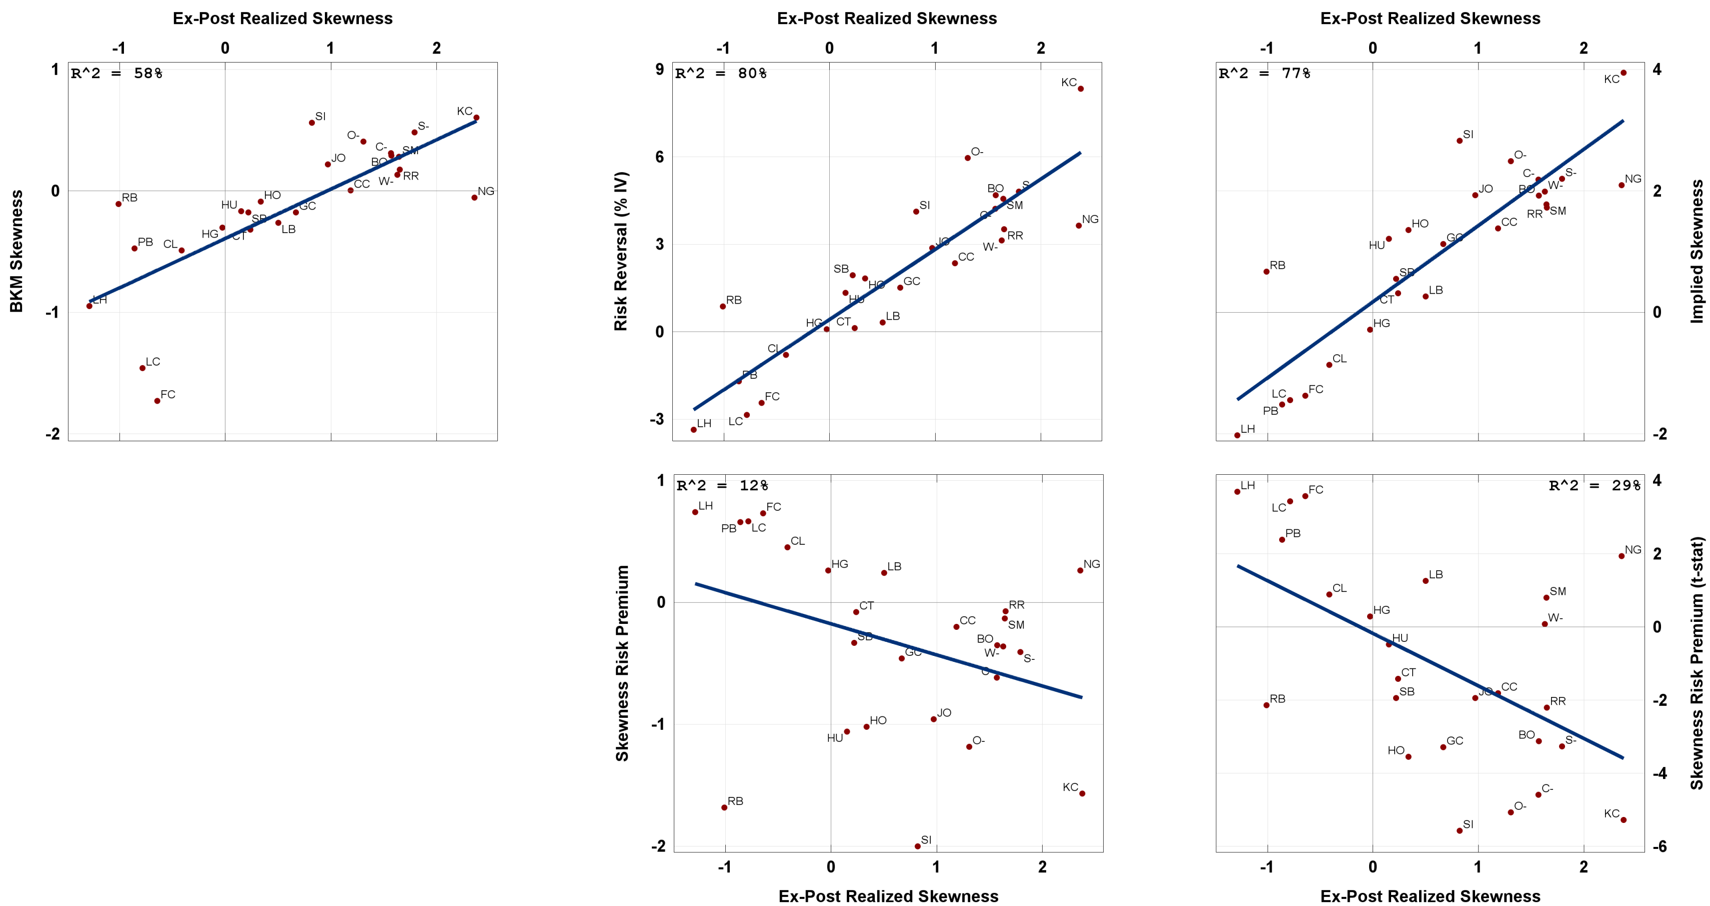

Chapter 3

The Dynamics of Overpricing in Structured Products[25]

3.1 Introduction

The retail market for so-called structured products (SPs) has been growing rapidly around the globe over the last 20 years.[26] Although they come in many forms, structured products exhibit a number of key commonalities. Their payoffs are combinations of several primary securities that may include options, equities, equity indices and fixed income securities. The primary clientele for structured products are small investors who cannot replicate the payoff by themselves. In addition, SPs are issued by financial institutions that stand by to redeem the securities over their lifetimes. This is typically necessary, because the issuance volume of these instruments is too small to create a continuous market unless the issuer provides liquidity. Short-selling, however, is not possible. Lastly, because structured products are traded outside of derivatives exchanges with clearinghouse, they carry the credit risk of the issuing party.

The existing empirical literature commonly finds significant overpricing relative to the primary components (Bergstresser, 2009). This is sometimes attributed to a number of features beneficial to small investors, for instance small spreads (Bartram et al., 2008) and access to complex payoffs (Wilkens et al., 2003). More recently, however, a second strand of research has begun to focus on the negative aspects of SPs arguing that issuers market and sell complex products with low expected returns (Henderson and Pearson, 2011) to retail investors by exploiting their behavioral biases and lack of financial literacy (Bernard et al., 2011). Recent articles in the financial press provide ample evidence that some of the risks involved in structured products are ill understood by small investors[27]. In particular, the bankruptcy of Lehman Brothers caused unexpected losses to investors of Lehman structured products that were marketed as being safe[28].

In this paper, we present convincing evidence consistent with the more pessimistic view on SPs. We find that issuers dynamically exploit their position as monopolistic liquidity suppliers to extract

[25] A version of this chapter will be submitted for publication. Ruf, Thomas (2012) The Dynamics of Overpricing in Structured Products.

[26] In line with the existing literature, this study focuses on publicly available structured products designed and marketed to retail investors as opposed to products that banks tailor individually to the needs of large investors (e.g. over-the-counter swaps). Bergstresser (2009) refers to them as structured notes. They were developed in the U.S. in the late 1980s and early 1990s (Jarrow and O'Hara, 1989; Chen and Kensinger, 1990) and have spread to Europe, in particular Germany (Wilkens et al., 2003). More recently, they are experiencing rapid growth in Asian markets.

[27] Wall Street Journal, May 28th, 2009, 'Twice Shy On Structured Products?'

[28] Wall Street Journal, October 27th, 2009, 'FSA to Clean Up Structured-Products Market'

gains from retail investors that go beyond the static overpricing previously documented.

Our analysis is based on a very large dataset of high-frequency trade and quote data of German bank-issued warrants. These warrants are the most basic SP because their payoff is simply that of an option[29]. We think that the insights that we gain from looking at warrants apply equally to more complex structured products, because they share most, and perhaps even all institutional features with bank-issued warrants in countries where both exist. In fact, they typically trade in the same segment of exchanges and many issuers of one type of instrument are active participants in the issuance of the other type of instrument as well. In addition, results based on warrants are not confounded by effects that arise from bundling several securities into one. A further advantage is the ease with which we can compare and match them to options on regular derivatives exchanges.

First, we investigate which types of warrants retail investors trade and how their preference affects overpricing. We document that retail investors, compared with professional investors, have a preference for far out-of-the-money (OTM) warrants offering high leverage as well as some far in-the-money (ITM) warrants. We argue that far OTM warrants are the most overpriced because unsophisticated investors find them difficult to evaluate, and no alternative instrument is available to them. Among the far ITM warrants, only puts are significantly overpriced because investors have few substitutes for short positions.

Second, we explore how issuers adjust prices facing demand in the secondary market. In particular, can issuers anticipate demand and exploit the liquidity needs of investors? Or do prices increase only after a positive demand shock consistent with a demand pressure explanation in the spirit of Garleanu et al. (2009)? Our results suggest prices for warrants are systematically higher (lower) on days when investors are net buyers (sellers). We show that it is not realized demand by investors which subsequently drives prices higher; rather, issuers are able to anticipate future net demand and opportunistically adjust prices in advance. Thus, the quoted bid/ask spread is not representative of the round-trip transaction costs that most investors face and returns are systematically lower, benefiting the issuer.

Third, we explore the 'life cycle hypothesis' (Wilkens et al., 2003) which suggests a declining pattern of overpricing over the lifetime of SPs. Previous studies use relatively small datasets and the methodologies of computing premiums are relatively crude due to lack of data. We revisit this question with our expansive dataset by applying a number of more refined matching techniques. We do find some evidence of a declining premium, but the decline depends on the warrant's moneyness and time to expiration. In particular, close-to-expiry OTM puts do not conform to the hypothesis and display an increasing premium. We argue that in both cases, the issuer acts rationally and exploits investors' demand, albeit in different ways. Further, we suggest that the large decline in premium over the lifetime of warrants that was found in some previous studies may be due

[29] It is important to clarify that bank-issued warrants are not warrants in the usual sense, i.e. warrants issued by firms on their own stock that dilute existing shares upon exercise. Rather they are option-like instruments issued by banks on equity, equity indices or any other underlying and settled in cash only. They are virtually unknown in the U.S. because options exchanges have a long history there and are readily accessible by retail investors. In many other countries, however, centralized derivatives exchanges are a relatively recent development or have been out of reach for small investors. In those countries, bank-issued warrants can fill part of this gap.

to improper matching along the maturity dimension between warrants and similar options on derivatives exchanges. We suggest several ways in which to adjust this mismatch.

Last, we investigate if an increase (decrease) in issuer credit risk leads to a decrease (increase) in the price of the structured product. Since SPs are unsecured debt obligations to the issuer, in an efficient market, their prices ought to rise and fall with the credit quality of the issuer. To the best of our knowledge, we are the first to measure this effect empirically.[30] We do find a negative effect of issuer credit risk on prices in our sample, but only in the aftermath of the bankruptcy of Lehman Brothers. However, the sensitivity seems generally too small and more specific predictions of vulnerable options models (Klein, 1996) are not borne out in the data. We would for instance expect that put warrant prices should be more sensitive to credit risk than calls. If anything, we find the opposite. These results suggest that investors are essentially providing issuers with cheap financing that goes beyond the notion of credit enhancement (Chidambaran et al., 2001; Benet et al., 2006).

The remainder of the paper is organized as follows. Section 3.2 discusses the existing literature. Section 3.3 details the data used and the methodology employed. Section 3.4 contains the empirical analysis. Section 3.6 concludes.

3.2 Literature Review and Institutional Background

3.2.1 Structured Products

While structured products (SP) may differ in many ways across borders they share a number of similarities. Common to all of them is that their payoff is a combination of several primary securities that may include options, equities, equity indices and fixed income securities. Since large, sophisticated investors can build these combinations easily by themselves, most structured products can be thought of as being exclusively designed and marketed to the wants of smaller retail investors.

SPs are most commonly issued by financial institutions such as large investment banks who typically also act as a market maker or liquidity provider in a secondary market. However, there is anecdotal evidence (see e.g. Pratt, 1995) that the early SPs issued in the U.S. were troubled by low volumes in the secondary market. On the other hand, SP issuers in German exchanges are obligated by the exchange to provide liquidity in narrowly defined terms (see e.g. Stuttgart Stock Exchange, 2010a, p.11). Among other things, issuers have to continuously provide binding quotes and keep the bid/ask spread within tight bounds. In all other ways, issuers are essentially free to set ask and offer prices.

While there may not be enough liquidity to allow for investors to trade structured products among themselves, the issuers stands by to sell and redeem its products at all times. Further,

[30] Recently, there have been a number of studies that discuss credit risk in the context of structured products. However, their starting point is always a model of vulnerable options from which fair values are computed. See for instance, Baule et al. (2008).

short-selling is not permitted by the issuer or by the exchange (see e.g. (Stuttgart Stock Exchange, 2010b), Section 49). Because SPs are traded outside of regular derivatives exchanges there is no central clearing house guaranteeing both sides of each trade and as a consequence all SPs carry the credit risk of the issuing party.

By now there exists a sizable literature on structured products from a number of international markets. Studies for the U.S. market have discussed Primes and Scores (Jarrow and O'Hara, 1989), S&P indexed notes or SPINs (Chen and Sears, 1990) and market-index certificates of deposits or MICDs (Chen and Kensinger, 1990), foreign currency exchange warrants (Rogalski and Seward, 1991) and more recently, reverse exchangeables (Benet, Giannetti, and Pissaris, 2006) and SPARQs (Henderson and Pearson, 2011). International studies have covered markets in Switzerland (Wasserfallen and Schenk, 1996; Burth et al., 2001; Grünbichler and Wohlwend, 2005), in Australia (Brown and Davis, 2004) and Germany (Wilkens et al., 2003; Stoimenov and Wilkens, 2005; Wilkens and Stoimenov, 2007). More or less all of them report that SPs contain a premium when compared to their individual components. The premium seems to be particularly large around issuance; e.g. Horst and Veld (2008) report premia of over 25% during the first week of trading.

The size of the premium is generally hard to justify but some studies suggest beneficial properties like guaranteed liquidity (Chan and Pinder, 2000) and smaller bid/ask spreads (Bartram et al., 2008) as causes. In addition, trading frictions like access to derivative markets or scalability may prevent retail investors from building these payoffs by themselves. In that sense, structured products may offer value by bundling securities, or 'packaging' (Stoimenov and Wilkens, 2005), for which investors should be willing to pay a premium. These benefits are largely absent for retail investors in the U.S. though; in particular, they already enjoy easy access to options exchanges and even more sophisticated strategies like the writing of options or spread trading can be implemented with little difficulty.

It is therefore almost surprising that only recently the literature takes a more negative view of structured products. Henderson and Pearson (2011) call SPs the 'dark side of financial innovation' because investors would be better off in the money market than buying SPARQs, the particular SP they analyze. Bernard et al. (2011) argue that issuers emphasize outcomes with high payoffs and low probabilities in their marketing materials leading retail investors to over-weigh those states in their expected return calculation. Bethel and Ferrell (2007) discuss legal and policy implications of the explicit targeting of unsophisticated investors with offers of complex financial securities. In the model of Carlin (2009) issuers faced with increasing competition increase the complexity of their products to make comparisons for investors more costly and maintain overpricing. Dorn (2010) documents that investors regularly fail to identify the cheapest security among in principle identical options. According to this view, issuers market and sell complex products to investors by exploiting their lack of financial literacy and behavioral biases (Shefrin and Statman, 1985, 1993).

3.2.2 Bank-Issued Warrants

Bank issued-warrants can be thought of as simple structured products as their payoff structure is just that of a put or a call option. Their existence is mainly due to the difficulty with which retail investors can access derivatives markets in a number of countries. They are unknown in the U.S. precisely because well-regulated and easily accessible derivatives exchanges have developed as early as the 1970s. By stark contrast, Germany did not have a derivatives exchange until 1990 (formerly called DTB, now EUREX). Even at the time of writing, according to the EUREX website (*http://www.eurex.de*), there are a mere 2 German brokers that offer retail investors access to EUREX products. The fact that warrants share most of the institutional features with structured products makes them closely related and any insight that we gain on the price dynamics of warrants should hold true for structured products as well.

Warrants enjoy the aforementioned benefits of SPs regarding liquidity and binding quotes, but they also offer some advantages over regular options. While stock options listed on the CBOE or EUREX trade in lots representing 100 shares of the underlying, warrants can be traded in much smaller increments; warrants representing one tenth of a share are common. For indices, the contrast is similarly large as EUREX contracts trade in lots of 5 while the typical ratio for warrants is 1 to 100. Being able to trade in small increments versus a rather large contract size at the option exchange again seems to be tailored to the small sums with which retail investor invest.

The literature on warrants is quite small and so far only covers markets in Europe and Australia. In line with the literature on structured products, empirical studies on bank-issued warrants find a pattern of overpricing relative to identical options from derivative exchanges that is particularly strong around issuance; see e.g. Chan and Pinder (2000) for Australia; Horst and Veld (2008) for the Netherlands; Abad and Nieto (2011) for the Spanish market.

Other studies focus on the bid/ask spread of warrants. Bartram and Fehle (2007) and Bartram, Fehle, and Shrider (2008) investigate the effects of competition and adverse selection on bid/ask spreads between German bank-issued warrants and options on the European derivatives exchange EUREX. The first study finds that the bid/ask spread for both options and warrants is lowered if a comparable instrument exists on the other market, thereby documenting that competition between these markets lowers transaction costs in both even though instruments traded in one market are not fungible in the other. The second study relates the much higher bid/ask spread for options to the potential for adverse selection that the market makers face as informed investors are more likely to be encountered on the options exchange EUREX. They find the ask price for warrants to be slightly higher than for options, but the bid price to be much higher. As a consequence the round trip costs are lower for warrants, even though the cost of entering the position and thus the value at risk is larger. The authors suggest that investors who are planning on holding the instrument for only a short period should be willing to pay more initially if they can sell it back at a better bid price relative to EUREX.

3.3 Description of Data and Methodology

We combine data from a number of sources. All datasets cover at least the period from June 2007 through May 2009. This time period contains a large fraction of the credit crisis that started in August of 2007. Therefore, the sample is well-suited to investigate the impact of credit risk on structured products.

For simplicity, we consider warrants and options on a single underlying, the German stock index DAX. The DAX consists of the 30 biggest German firms and is a performance index, i.e. dividends are assumed to be reinvested in the index. Our sample consists of 8,750 warrants with expiration dates between April 2008 and December 2012, while the options in the sample expire between June 2007 and December 2012.

To compute moneyness and deltas we acquire second-by-second tick values of the DAX index from KKMDB database at the University of Karlsruhe. Supplemental data such as issuer CDS spreads and the closing value of the German volatility index VDAX (the equivalent of the CBOE VIX) were taken from DATASTREAM.

We acquired all DAX warrant transactions taking place on Scoach from Deutsche Börse AG, while transaction data from EUWAX comes from the KKMDB at the University Karlsruhe. Bid/ask quotes for the warrants are from Börse Stuttgart AG. EUREX transaction data was supplied by Deutsche Börse AG in Frankfurt; it covers all trades that took place in options on the German stock index DAX during the sample period.

3.3.1 Warrant Data

In Germany, warrants are traded in one of two ways. Each issuer offers an OTC platform, in which investors can trade directly with the issuer via the online interface of their broker. Within seconds of the investor requesting a quote, the issuer transmits a binding offer for selling or buying back warrants which is valid for the next 2-3 seconds. Investors thus have the opportunity to quickly trade for prices known in advance. Conversations with issuers revealed that more than half of all trading in warrants and structured products takes place through this channel. In a sample of actual retail investor transactions used in Dorn (2010), over 80% of warrant trades are executed this way. The remaining trade takes place in special segments of the regular stock exchanges: in Frankfurt this segment is called Scoach; in Stuttgart it is called EUWAX (European warrants exchange); the other regional exchanges play virtually no role in the trading of these instruments.

In their study, Garleanu, Pedersen, and Poteshman (2009), henceforth GPP, make use of a particular dataset that reports daily open interest by investor group (public customers and proprietary traders) and derive the level of net demand facing market makers. The unique features of the market for warrants allow us to do something similar. Every transaction takes place with an investor on one side and the issuing bank on the other. Therefore, a transaction with a price above the mid quote likely constitutes a buy by an investor increasing the total of outstanding product; a price below the mid quote constitutes the opposite.

Unfortunately, we only observe warrants transactions that take place on the regular exchanges; data on OTC transactions with the issuer is not available. Representing less than half of all transactions, the sum of buys minus the sum of sells from the exchanges yields only a very noisy picture of the true number of warrants currently in the hands of investors; in fact, summing up buys and sells over time leads to negative numbers for numerous warrants in our sample, which is obviously not possible. Instead of investigating how the *level of demand* impacts the *price level* of warrants like GPP, we focus on investigating if *changes in net demand* have an influence on *changes in prices*.

We are confident though that over the course of a trading day, the net change in demand derived from observed transactions is a suitable proxy of the true net change in demand from all transactions. This is valid as long as the observed trades are an unbiased subsample of all trades.

3.3.1.1 Identification of Transaction Type

We acquired quote data for 8,750 individual warrants totaling more than 4 billion quotes over 2 years. To identify if a transaction was a buy or a sell by an investor, each trade is matched with the currently valid bid/ask quote. Unfortunately, the time stamp of the quote data is only given up the full second for all but the last day in the sample, i.e. the quote could have been updated any time during that 1 second interval. In contrast, transaction data is given with milliseconds. Thus, if the second of the quote and the transaction coincide we do not know for certain which came first, i.e. if the quote was already in place at the time of the trade, or rather if the trade came first and triggered an update of the quote from the issuer. This occurs for 8.5% of all trades. We proceed as follows: If the transaction happens in the first half-second, we assume that it arrived before the quote, therefore the previous quote is assumed to be valid. Naturally, this is what we do with transactions that do not coincide with a quote. In the other case, we assume that the same-second quote came first and use it to evaluate the type of the trade. If this does not bring a decision because the transaction price is equal to the mid quote, we use the previous quote.

If the type of the trade could not be identified up to this point, we consider the following quote as long as it occurs within 60 seconds of the transaction. If the mid quote increased we consider the transaction a buy, and a sell if otherwise. Trade for which no decision could be made are not considered in the aggregation of net demand. All in all, this algorithm fails to identify 2,506 transactions, or 0.27 percent, out of a total of 925,000 transactions for which we had quote data. We identify 54.36% of all transactions as buys and 45.37% as sells. This seems reasonable as a certain share of warrants will likely remain on the books of investors until they expire.

3.3.1.2 Construction of European Warrant Prices

Warrant bid/ask quotes are updated frequently over the course of the day. In most of our analysis we therefore make use of multiple quotes per day. Once every hour from 9:30CET to 17:30CET we extract the latest quote for each warrant from the data. We eliminate quotes that are older than 1 hour so as to have no overlap between measurements.

Most warrants are of American exercise type and are thus not immediately comparable to European EUREX options. Since the underlying does not distribute dividends, only put options potentially incorporate an early exercise premium (EEP). Since our study deals with warrants, many of which are longer dated, this premium is not negligible even for contracts that are at-the-money. We find the EEP to be 5 percent on average for a one-year at-the-money (ATM) option/warrant.

We proceed as follows: First, we back out the implied volatility of the American warrant ask quote via the BBSR algorithm[31] algorithm developed by Broadie and Detemple (1996). We choose the ask price because it is a frequent occurrence for options in all markets that mid-quotes of far-in-the-money or close-to-expiry contracts violate the no-arbitrage bounds. We would like to retain as much of the sample as possible and not lose valid quotes because the mid-quote is too low for the Black-Scholes model. The error that we incur is small because in-the-money (ITM) contracts have very small bid/ask spreads. In a second step, we use the implied volatility (IV) of the ask quote as input into the Black-Scholes formula holding all other parameters constant to get a European ask quote. We deduct the originally observed spread to back out the new bid quote. The parts of the analysis that genuinely require IV estimates as inputs are based on Black-Scholes IVs from mid quotes.

3.3.2 Option Data

3.3.2.1 Option Implied Volatility Functions (IVF)

As we lack access to bid/ask quotes on EUREX options (the dataset has a price tag of over EUR 10,000), we focus on transactions data. The goal is to compute the premium that a warrant demands over an identical EUREX option at a particular point in time. It is generally not very likely to find a transaction for a EUREX option in close temporal proximity to each warrant quote with identical or very similar features regarding strike, expiration date and type.

Instead we opt for a different route: Over a small time interval (e.g., a trading day or less), we collect all option trades for a given maturity and type, compute their moneyness at the time of the transaction and back out Black-Scholes implied volatilities from prices; this step is straightforward as all DAX options are of European exercise style and there are no dividends to consider. Bartram and Fehle (2007) report that the difference between bid and ask prices can be relatively large on EUREX (8-9 percent in their year 2000 data); in practice, however, trades take place within a much smaller effective bid/ask spread, which should be somewhere in the vicinity of the mid-quote that one would otherwise use.

Then we fit a 3rd degree polynomial through all IVs as functions of moneyness to get an implied volatility function (IVF) for each maturity and type, separately for each time interval. We require a minimum of 12 observations per interval to include the IVF in the sample and we save the moneyness levels of the most extreme observations that went into estimating the IVF. When matching this

[31]BBSR is a binomial tree algorithm that uses the Black-Scholes formula in the ultimate step of the tree as well as a two-point Richardson extrapolation.

IVF with warrant quotes we only allow quotes whose moneyness lies within these bounds to avoid any issues resulting from extrapolation. We stop creating IVFs if the time to maturity is 5 days or less, because the tails of IVFs tend to go vertical as maturity approaches. In addition, if the range of moneyness levels observed in trades falls below a certain width, the estimate is dropped from the sample regardless of the number of observations. A too narrow moneyness range may result in unrealistic IVFs. On average, the R-squared from the fit of the polynomial is a very high 95 percent.

As for the choice of the length of the interval, we face a trade-off between higher accuracy of the IVF and smaller number of trades as we decrease the time window over which to estimate the IVF. Even though the sample is substantially reduced as we go from daily IVFs to hourly IVFs, we still opt for hourly IVF estimates for most of the analysis given the rather volatile sample period.[32] To derive hourly IVFs, we split the trading day (from 9:00CET to 17:30CET) into 9 time intervals, the first half-hour from 9:00CET to 9:30CET (having the highest trading volume of the day), and the following 8 one-hour intervals up to 17:30CET.

Implicit in this estimation method is the assumption that the IVF stays more or less constant in terms of moneyness over the course of the respective intra-day interval. Obviously, on volatile trading days this may not hold, but this is purely an issue of noise, not a systematic bias.

3.3.2.2 Premium Computation

Parts of the empirical analysis require the computation of a premium of the warrant price over the price of a EUREX option with identical features. We use the mid quote from the warrant and compute its moneyness at the time of the quote extraction, e.g. 10:30CET. We require an hourly estimate of the EUREX IVF with expiration in the same month as the warrant and for the same type. We match the option IVF estimated from trades taking place between e.g. 9:30CET and 10:30CET to the warrant quotes that were valid at 10:30CET (i.e. issued sometime between 9:30CET and 10:30CET). The moneyness of the warrant quote is plugged into the option IVF to compute the IV for a virtual EUREX option, which in turn is used to compute a European option price. The simplest version of the premium is the percentage difference between warrant mid quote and this artificial EUREX option price.

In each calendar month, almost all warrants in the sample expire 2-4 days prior to the EUREX expiration date depending on the issuer. We exclude warrants with other expiration dates from the analysis. However, we still face a slight time to maturity mismatch. One way to adjust for this is to use the time to maturity of the warrant instead of that of the option in the Black-Scholes formula when computing an artificial option price. This implicitly assumes that the shape of the IVF does not systematically change when shifting the maturity by a few days. We know that this assumption will be violated for options close to maturity. Therefore, our empirical analysis will exclude observations with less than 20 days to maturity. We call the second premium the adjusted

[32] We find that results using the less exact matching technique are generally similar, but due to the greater noise, levels of significance are somewhat reduced.

premium and the first one unadjusted.

We develop two more versions of the warrant premium. Implied volatilities do not only exhibit a smile shape across the strike price dimension, they also exhibit a specific term structure along the maturity dimension. For instance, during volatile times short-term options have a much higher implied volatility than long-dated options, while the opposite holds during calm periods. The reason for this tapering off is the mean-reverting behavior of volatility.

We can therefore use the information between two adjacent expiration dates to infer implied volatilities for options with expiration dates in between. In order to avoid erroneous interpolation, we require that both IVFs used in the process extend to the moneyness of the warrant for which we intend to compute a premium. The interpolated IV is simply the linearly weighted average of the IVs derived by plugging in the warrant moneyness into both IVFs. Since 2-4 days is relatively close to one of the end points, the error from using a linear fit remains small, even if the IV terms structure is strongly concave or convex.

One problem of this method is that it requires an IVF for the expiration date of the previous month. Yet the previous month IVF is no longer available close to its expiration. Furthermore, the range of traded strikes contracts rapidly as expiration approaches. To mitigate this effect, we compute one more premium version. Here, the estimated IVF is based on the EUREX IVF with expiration in the same month as the warrant and the one closest to expiration after that. This means that we extrapolate 2-4 days outside of the two observation points. The requirement that both IVFs extend to the desired moneyness level is maintained, however.

This solves the issue of losing a large fraction of the sample on the short end, but it leads to larger losses at the long end. For instance, to compute the premium of a December 2010 warrant, the IVFs for December 2010 and the next available expiration are required. During the year 2008, this was likely December 2011 as additional expirations are only filled in much later. As EUREX options do not trade much this far into the future, the likelihood of having enough trades to estimate both IVFs on the same day are very small.

Both interpolated and extrapolated premium have higher data requirements and samples are therefore smaller. On the other hand, the first premium based on a mismatch is inaccurate and in fact biased. We use the interpolated version of the premium for a large fraction of the empirical analysis.

3.4 Empirical Analysis

3.4.1 Preference and Overpricing

The comparison of warrants and options is an almost ideal laboratory set up where two distinct groups of investors trade in two separate market segments. Warrants and other structured products are most successful in countries, where retail investors are subject to high barriers to entry into derivative exchanges. Germany is such a case among developed markets, where a derivatives exchange was only established in 1990 and to the present is difficult to access by small investors.

The lack of access is likely the main reason why those countries have flourishing markets for structured products. The EUWAX in Germany, for instance, is the world's largest exchange for SPs by number of products with more than 500,000 securities outstanding.[33]

In principle, there are no barriers that keep large investors from trading warrants or any other structured product. In practice, however, the design of the market place is geared towards small transactions and the needs of retail investors, because 1) bid/ask spreads and the stand-by liquidity are only binding up to a certain order size (EUR 3,000-10,000) and 2) the amount of units outstanding per product is generally not very large (warrants, e.g., are typically issued in batches of 1-10 million units; on the CBOE this would equate to a mere 100-1000 index option contracts). In addition, sophisticated investors can easily replicate all payoffs by themselves and will certainly want to avoid the overpricing and the credit risk involved.

Thus, we argue the segmentation is strong and comparing trading pattern between the two markets offers insight into the different motivations of private investors relative to institutional investors. In particular, we can see where the demand for warrants differs from the demand for options on the EUREX.

Table 3.1 depicts the relative frequency of transactions across different ranges of time to maturity and moneyness. We use the frequency of trades as a proxy for net demand, because the data does not allow accurate measurement of true net demand. Notable is the extremely strong concentration of EUREX option transactions close to expiration as well as at-the-money (ATM). More than 80% of all transactions happen in contracts with less than 3 months to expiration and over half of all trading occurs in a narrow 5% band around the money; there is barely any trading in in-the-money (ITM) options. In contrast, warrants transactions are much more dispersed across maturities and moneyness. Far-in-the-money and far-out-of-the-money warrants experience larger demand than comparable options. Particularly, the demand for out-of-the-money (OTM) call warrants greatly exceeds the trading activity in OTM call options. In addition, there is asymmetry in the demand for OTM call vs. OTM put warrants. More than 36% of all trades in long-term warrants happen in excess of 25% OTM; for puts this number is just 15%.

In buying warrants that are even more OTM than what option traders buy, warrant traders exhibit a greater preference for high-leverage securities paying off in states with low probability. This supports recent findings that warrant traders are motivated by speculation (Glaser and Schmitz, 2007) and entertainment (Bauer et al., 2009), rather than hedging. The activity in far-ITM warrants on the other hand is quite puzzling at first. Leverage is rather low and warrant unit prices are high in this region. Further, given that Glaser and Schmitz (2007) report the median holding period of warrants to be 3 days, it is hard to understand why more than one third of trading is concentrated in securities with a remaining life time in excess of 6 months.

In summary, we recognize that there are clear differences in the demand pattern between warrant and option market participants. Given these differences in demand it would be interesting to see

[33] See https://www.boerse-stuttgart.de/en/marketandprices/marketandprices.html. As of February 15th, 2011, Stuttgart stock exchange listed more than 230,000 warrants, 100,000 knock-out products and over 250,000 certificates.

if the warrant premium shows a pattern that reflects the pattern in demand. Does the premium differ across regions of moneyness or maturity?

Because the moneyness range of actively traded warrants declines for closer-to-expiration contracts, we use the (absolute of the) warrant's delta rather than its raw moneyness to make the pattern in overpricing more comparable across different maturities. When computing delta we replace the warrant's IV with the IV of the matching option. Otherwise the degree of overpricing of the warrant would directly impact its delta. To further avoid influence of outliers we first cut off the top and bottom 1% of premiums, then we cut off the lowest and largest 1% delta values for each issuer.

Figure 3.1 shows the median premium as a function of the warrant's delta for the 6 largest issuers in the sample based on quote data, separately for puts and calls. The sample is split into the maturity bins that correspond to the maturity bins in Table 3.1. All graphs exhibit the same striking pattern where low-$|\Delta|$, far-OTM warrants command much higher premiums than ATM or ITM contracts. OTM calls are generally more overpriced than puts by the same issuer with the exception of some issuers over some moneyness regions for close-to-expiration warrants. This pattern corresponds to the asymmetry of demand for OTM contracts in Table 3.1.

One possible explanation is that retail investors find it much harder to determine the fair price of an OTM warrant, which consists of time value only, relative to ITM warrants, which are mostly intrinsic value. Carlin (2009) suggests an equilibrium model in which more complex products can be overpriced more heavily. If we consider OTM warrants as being more complex than ITM warrants, the overpricing is in line with his model. However, this does not directly speak to the large demand for high-leverage warrants.

Alternatively, investors with different beliefs may choose securities with different degrees of leverage. It seems intuitive that a very optimistic investor would choose warrants with higher leverage anticipating higher returns relative to a less optimistic investor who may choose lower-leverage warrants or the underlying. The issuer is then able to extract some portion of the consumer surplus (with regards to the beliefs of the investor) by charging a relatively higher premium. Assuming that retail investors that trade warrants tend to be more risk-seeking (in the spirit of Kumar, 2009), this explanation fits well with a recently emerging strand in the asset pricing literature that focuses on the preference for positive skewness in stocks and other assets (e.g. Brunnermeier et al., 2007). Boyer et al. (2010) and Bali et al. (2011) empirically document a negative correlation between stock returns and expected skewness. Barberis and Huang (2008) motivate this behavior within prospect theory. Finally and most related, Boyer and Vorkink (2011) show that even option returns and option return skewness are negatively correlated and ascribe this finding again to a preference for positive skewness by end-users of options (i.e. retail investors).

Both explanations for the observed cross-sectional pattern, complexity and skewness preference, are behavioral in nature and fit well with the related literature. By contrast, the commonly cited reasons for the *average* overpricing in structured products do not immediately have cross-sectional pricing implications. Liquidity, for instance, is mandated by exchange regulation and posted quotes are binding equally for all instruments. Packaging services Stoimenov and Wilkens (as suggested by 2005) are not relevant in the case of warrants, but would not differ as a function of strike either. To our knowledge, taxation does also not affect the cross-section as gains and losses are taxed symmetrically on the part of the retail investor.

If the demand pattern in Table 3.1 drives prices, one might have expected to see some overpricing for far-ITM warrants as well. We notice a kink in the premium for far-ITM puts but none for calls. This suggests that the elasticity of demand differs in those areas. The charged premium and the demand for warrants by investors are equilibrium outcomes. One possible explanation is that issuers exploit the lack of good alternatives for investors who would like to express negative views with low leverage securities, while there are plenty of instruments for going long with low leverage.[34] Hence, issuers can charge a premium for ITM puts, but cannot demand a premium for ITM calls. Thus, this section provides some basic evidence for the hypothesis that issuers take advantage of investors' demand for certain payoffs through overpricing. We turn to a more systematic investigation of the dynamics of the overpricing in the following section.

3.4.2 Demand Pressure vs. Demand Anticipation

Issuers of structured products in Europe maintain binding quotes for all of their products and stand by to buy and sell if investors want to trade. From conversations with issuers, we understand that it is literally unheard of that a trade is executed between two private investors. Without the issuer there would be no liquidity. The flip side of this coin is that the issuer determines the price and investors have no choice but to accept that price if they want to trade.

Given this monopoly power, the issuer has some incentive to skew the price in his favor. One way to extract profits from trading in excess of the clearly defined bid/ask spread would be to offer higher than usual prices on days when investors predominantly buy and offer lower than usual prices on days investors are mainly selling. This requires some degree of predictability for the change in net demand, i.e. the order flow of structured products. As we will show below, we find evidence for such predictability. The frequency at which order flows change are quite high and depend mostly on the returns of the underlying in the immediate past (yesterday's return and the overnight return) and as a consequence, predictability as well as price impacts are also limited to very short horizons.

Our time frame is thus distinctively different from the literature on option demand pressure (Bollen and Whaley, 2004; Garleanu, Pedersen, and Poteshman, 2009). Both studies relate private investor demand to relative prices of options, i.e. the skew of the option smile, at a monthly frequency. Similarly, Amin et al. (2004) condition on large returns over the past 60 days to explain changes in the implied volatility of options. Lemmon and Ni (2011) combine those findings and posit that market sentiment along with lagged market returns drive option demand which in turn

[34] For regulatory and historical reasons, retail investors cannot short stocks in their accounts at Germany-based brokerages. One would have to open an account with a foreign broker. Also, during the sample period there were no so-called inverse ETFs listed on the German market.

impacts option prices. Their analysis is also at the monthly frequency.

Our analysis is similar to Lemmon and Ni (2011) in that we show that returns drive demand which in turn drives warrant prices.[35] However, the channel that we propose is quite different from the limits-to-arbitrage explanation suggested by Bollen and Whaley (2004) (BW) and formalized by Garleanu et al. (2009) (GPP).

BW argue that market makers in the options market face limits to arbitrage because their access to capital and thus, their tolerance for intermediate losses is limited. In this case a growing net position causes increasing 'hedging costs and/or volatility-risk exposure' for market makers. Consequently, market makers are willing to supply additional options only at increasingly higher prices. For instance, institutional investors have a large demand for index puts for which there is no natural counter-party in the market. Market makers absorb this demand imbalance at a premium, which according to BW can partially explain the volatility smile observed in index options.

GPP suggest discontinuous trading, price jumps and/or stochastic volatility as causes for the inability of market makers to hedge perfectly and derive analytically how demand pressure increases option prices in the presence of these frictions.

By contrast, we suggest that issuers use their position and knowledge of future order flow to adjust prices at a relatively high frequency to extract additional profits from investors. This is not to say that issuers in the warrants market could not be subject to demand pressures at the monthly frequency as well. However, with our quite short time series we must leave this question to future research.

Finally, we should point that e.g. Lemmon and Ni (2011) find sentiment and market returns to have strong effects on equity options but less so on index options because small investors account for only 3 percent of trading in index options, but 18 percent in stock options in the U.S. options market. In our case, the separation between small and large investors happens already with the choice of exchange. As we pointed out earlier, both index and stock warrants are overwhelmingly the domain of retail investors trading on EUWAX and Scoach, while large investors trade in the derivatives exchange EUREX.

3.4.2.1 Aggregate Demand

We start by investigating the dynamics of daily aggregate demand for warrants, separated by type of warrant and type of transaction. First, we allow contemporaneous variables to explain demand. In a second step, we predict aggregate demand using only variables that are known prior to the arrival of demand.

Using a rare dataset of actual transactions by retail investors in the German warrants market, Glaser and Schmitz (2007) study the motives of retail warrant investors and find that hedging considerations play virtually no role as the median holding period of both put and call index warrants is a remarkably short 3 days. This indicates that most retail investors use warrants to speculate on very short-term movements in the market. When partitioning sells into profitable vs.

[35] Since our analysis is at the daily or even intra-daily frequency we are precluded from using measures of sentiment.

losing trades, they find that investors tend to hold warrants twice as long when they are trading at a loss vs. at a gain: median holding periods are 4 vs. 2 days, and average holding periods are 24 vs. 12 days. This shows that warrant traders are also affected by the disposition effect documented for stocks by Shefrin and Statman (1985) and Odean (1998).

These observations make it evident that contemporaneous and immediate past returns will play a prominent role in determining selling decisions in particular. The buyer of a index call warrant appears to be very likely to sell within 1 or 2 days if he experiences a positive return over this period. Likewise, the buyer of a index put will sell within the next day or two if the market declines over that period. Further, since the holding period is quite short regardless of gains, we think that recent buying activity should foretell selling activity as well.

First, we regress daily aggregate demand for warrants on lagged demand and lagged and contemporaneous returns. Returns are measured from yesterday's closing of the regular market to today's closing (1730CET). We measure today's demand as aggregate Euro volume per category (Calls vs. puts, buys vs. sells). Warrant trading continues until 2200CET each day, but at much lower volumes than during regular trading hours. We would like to measure returns and order flow over identical time periods, therefore we assign any trading activity that occurs after the official closing to the next trading day. Lagged demand is defined as the sum of daily demand per category over the preceding 3 trading days. In addition, we include total unsigned trading volume over the previous 2 weeks as well as the lagged change and the level in market volatility as control variables. Total volume should help us to distinguish the impact of generally higher trading activity from short-term buying and selling pressure. Including volatility variables will help to measure the effect of market returns on demand more precisely because of the well-known negative correlation between returns and volatility.

Explanatory power in Table 3.2 is very high. In line with our predictions we find a high propensity to sell after positive returns for calls and after negative returns for puts. Lagged buys and sells positively impact today's buys and sells. Somewhat surprising, both puts and calls are more likely to be bought following negative returns. Total lagged unsigned demand is strongly significant and positive indicating more trading activity in the present given higher trading activity in the recent past. Interestingly, the coefficients on volatility levels and changes are of opposite sign. Times of higher volatility are associated with generally lower trading volumes, while a positive shock to volatility during the previous day causes more selling activity today. Presumably, an increase in volatility makes all warrants worth more and thus the chances of a position being profitable increase, leading to faster selling by investors.

In order to avoid undue influence from extreme returns over the rather volatile sample period, we repeat the analysis with all return variables winsorized at 3 per cent. Results are qualitatively the same, albeit R-squared are slightly lower.

Having found a contemporaneous connection between demand and returns, we turn our attention to the predictive power of lagged demand and returns with regards to future demand. This is important because we would like to know if issuers can anticipate order flows, for instance at the beginning of a trading day. If this is the case then issuers are in a position to adjust prices to take advantage of investors' demand before it arrives.

To this end, we will form expectations of intraday demand pressure that occurs between 930CET and 1730CET. As explanatory variables we use data that is known to the issuers at 930CET, i.e. yesterday's return and net demand as well as return and demand occurring between 1730CET yesterday and 930CET today. We call this the overnight return and the overnight demand.

Compared to the results earlier R-squared in Table 3.3 are naturally lower but still very high. The results from the previous regression seem to carry through with regards to the signs and importance of returns on demand. As a robustness check, we repeat the analysis with return variabels winsorized at 3 percent and find results essentially unchanged.

It is remarkable that simply by using returns and order flows from the immediate past, we can explain over 60 percent of total variation in demand for puts over the course of the day. Calls are somewhat more difficult to predict, and we might thus expect that prices are more responsive to demand for puts than for calls.

We are aware that our predicted demand is an in-sample prediction. We choose this path because of the relatively short sample period and the fact that we only observe a subset of true demand, both of which might make out-of-sample predictions extremely noisy. Further, all we have to assume for this prediction to be attainable by issuers at the time is that coefficients of lagged demand and return stay constant over time. We have no reason to believe that e.g. investors switch their propensity to sell calls after markets went up from one year to the next.

3.4.2.2 The Impact of Demand on Warrant Prices

We now turn to investigating the connection between warrant prices and contemporaneous demand in two steps. First, we check if demand and prices are correlated contemporaneously. Second, we want to know if prices adjust before demand occurs or after.

One possible methodology would be similar to what existing literature has done: match each warrant with an option individually and estimate the relative premium of warrants over options. We follow this path in all other parts of the empirical analysis but for the effect of demand we choose a different route for several reasons.

First, our study is at a disadvantage to others because we do not have bid/ask quotes of options, only transactions data. We computed IVFs for options based on transactions but found the imputed option prices too noisy to be of use even when pooling transaction over the entire trading day. We would require hourly IVFs.

Secondly, assume we were to observe bid/ask quotes for options and were thus able to compute reliable estimates for warrant premiums. We try to identify the impact of warrant demand on warrant prices relative to option prices, which according to the findings of GPP also face price impact from demand pressure. Unlike GPP, we do not have access to a dataset that shows outstanding net option positions by investor group. We therefore have no way of knowing who bought and who sold a particular option contract as not all trades have to occur between a market maker and an investor. Then, given estimated net demand for options we would have to estimate the impact of option demand on the premium at the same time that we estimate the impact of warrant demand. This procedure appears to be dependent on too many moving parts that we know too little about.

Instead we opt for a different route that is able to circumvent our data limitations. We use option implied skewness proposed by Bakshi, Kapadia, and Madan (2003) (BKM) as our measure of choice. BKM use prices of OTM puts and calls to derive non-parametric moments of the option-implied expectations of the return distribution of the underlying. Relatively higher prices in some range of moneyness imply that investors assign a higher probability to the underlying being in that range at maturity. Intuitively, if prices of OTM puts are higher than prices of equally OTM calls, investors assign higher probabilities to negative outcomes, which leads to negative implied skewness. Additional details of constructing BKM measures as well as minimum number of options required are discussed in Dennis and Mayhew (2002).[36]

Given the large quote dataset we have too many rather than too few warrant quotes available at each point in time. We have seen in Figure 3.1 that far-OTM warrants are subject to extreme overpricing. To guard against the results being driven by the large premia in low moneyness warrants we exclude put warrants that are more than 20% OTM and call warrants that are more than 25% OTM. The reason for the slight asymmetry lies in the weighting scheme of the BKM skewness which is based on the log of moneyness.

We compute BKM skewness for each warrant chain of each issuer once every hour. A warrant chain consists of all warrants by the same issuer that have a common expiration date. To compute skewness at e.g. 1730CET, we select the last mid quote of each OTM warrant issued prior to this time. If the quote is older than an hour it is discarded. Because the derivation of the BKM measures is based on European-type options, we transform quotes of American-type warrants into European-type prices via the binomial model described earlier.

In Table 3.4, we regress *daily changes in skewness* (1730CET to 1730CET) on several versions of net demand for puts and calls that are measured over the same time period. We include the lagged level of skewness and yesterday's return (split up into up and down part) as controls.

We find that total net demand for both calls and puts enters highly significantly. The signs of the coefficients indicate that demand is positively correlated with warrant prices: higher demand for calls coincides with higher skewness, which means that the right tail, i.e. calls, becomes more expensive than the left tail; higher demand for puts coincides with lower skewness, which means that the left tail becomes more expensive relative to the right tail. The issuer-specific net demands by itself have the same effect, albeit weaker. Because issuer-specific variables generally turn out to be quite weak, we will omit them from the following analysis. Using net demand by expiration instead of total demand has again similar effects but is also weaker. If we use it in addition to total demand, only the put demand remains significant.

In summary, we find that, on a daily horizon, demand is positively related to contemporaneous

[36]BKM as well as Dennis and Mayhew (2002) suggest to use at least 3 option prices per side. In the present context of warrant prices, this suggestion is never binding. 10 to 20 observations per side are typical.

changes in skewness. The question is: Does demand cause prices to rise, in which case we would be in the world of Garleanu et al. (2009)? Or do prices adjust preemptively to expected demand? This would support the case for opportunistic price setting by issuing banks.

To distinguish between these two explanations, we compute the overnight change in skewness (measured from 1730CET of the previous day to 930CET of the next day). We choose 930CET, because the regular market opens at 900CET and we want to make sure that orders entered overnight are not counted towards the intraday demand.

We then split daily demand into two parts. Overnight demand consists of all transactions that occurred between 1730CET and 930CET, while intraday demand consists of all transactions that occur after 930CET until that day's closing at 1730CET. Table 3.5 shows results for both realized intraday demand as well as predicted demand. The latter is based on the fitted values from the regressions on total demand shown in Table 3.3 and identical regressions for demand by expiration not reported.

Note from column (1) in Table 3.5 that lagged and overnight demand alone only explain a small part of skewness changes, as some controls (not shown) are already highly significant. Further their signs are different from the previous table. In columns (2a-c) actual net demands that occur during the day after 930CET have been added. Explanatory power of the models is still relatively low and adds at most 3.4 percent over column (1). Compare that to columns (3a-c) where actual demand is replaced by predicted demand. Significances are generally higher and the explanatory power is raised substantially. With one exception, all predicted net demands enter with the right sign and are significant.

The results lend strong support to the view that price changes preempt demand. Future expected net demand for calls is met by increases in skewness, i.e. higher call prices, while future expected net demand for puts is met by a decrease in skewness, which indicates higher put prices relative to calls. Thus, issuers systematically short-change investors by overpricing warrants that are in net positive demand over the following hours, while underpricing warrants that will be redeemed on a net basis.

3.4.3 The Life Cycle Hypothesis

Wilkens, Erner, and Röder (2003) analyze two types of SPs in the German market, reverse convertibles and discount certificates, and find that the overpricing present in these products declines as expiration comes closer, which they term the 'life cycle hypothesis' (LCH). Some subsequent studies find supporting evidence for a number of other structured products (Grünbichler and Wohlwend (2005), Stoimenov and Wilkens (2005), and Entrop, Scholz, and Wilkens (2009)), but Abad and Nieto (2011) fail to find such a pattern for warrants in the Spanish market. Wilkens and Stoimenov (2007) argue against using the LCH for products with knock-out feature[37] because expiration is a random event in that case.

[37] A knock-out feature causes the instrument to expire worthless as soon as the underlying hits a pre-determined level for the first time.

The idea behind LCH is as follows: At issuance, most SPs are in possession of the issuer (although active marketing of upcoming IPOs is meant to place a portion of the product with investors ahead of issuance). Over time, as investors buy the product, the chances of some of them wanting to redeem securities from the issuer increase. Since bid and ask prices are bound together rather closely (by exchange regulation), an asking price far above fair value would imply a bid price that is also too high. Thus, by keeping the bid price too high for long, the issuer risks being sold back some of her product at inflated prices. Wilkens et al. (2003) argue that investor buying activity should generally decline as maturity comes closer, while selling activity will likely increase. To optimally profit from the life cycle of demand, the issuers should gradually reduce the overpricing over the life time of the product.

Neither Wilkens et al. (2003) nor any of the subsequent studies are able to test this hypothesis directly because of the lack of demand data. In this study, we only have access to a fraction of total demand. Thus, estimates of net demand over periods longer than a few days are likely too noisy. Nevertheless, using daily net demand, we documented in Section 3.4.2 how issuers adjust prices on a daily basis to exploit high-frequency changes in net demand. It would not be surprising to find such a pattern at longer horizons as well. However, the unavailability of OTC transaction data on warrants and other structured products makes testing this hypothesis directly rather difficult.

In the following, we revert to proxying for the effect of life time net demand by using the time to maturity just like previous studies did. Where we differ from previous research, however, is how we compute the warrant premium. Wilkens et al. (2003) and Stoimenov and Wilkens (2005) use a simple hierarchical matching, which with slight deviations is employed by related studies as well. For each transaction (or quote) of a structured product, all transactions (quotes) of EUREX options are considered that minimize the difference in strike prices; in the next step, among all matches of the first step, the difference in time to maturity in minimized. In a third step, the difference in time stamps is either minimized or used as a filter criterion. Wilkens et al. (2003) are able to match strike prices quite well, but generally match EUREX options with a maturity half as long as that of the SP. Similarly, Stoimenov and Wilkens (2005) report average deviations in maturities of 5-7 months.

To evaluate the effect of maturity on the premium, it seems crucial to compute premia from options that are very close precisely in the dimension of maturity. Section 3.3.2 describes how premia are constructed for the purpose of this study. In contrast to previous studies, we match warrant quotes with implied volatility functions constructed from option trades that take place in the same hour using options that expire in the same month as the warrant. Further, we impose a minimum requirement on the number of options, the range of moneyness of the option transactions and we require that the moneyness of the warrant at the time of the quote is covered by that range. Daily averages of premia are calculated and admitted if there are more than 2 premiums observed for a warrant on a given day.

In the absence of option quotes (which could be easily matched one-to-one with warrant quotes), we feel this is a robust way to compute premia. As mentioned in Section 3.3.2, the method is not free of bias because issuers choose expiration dates in a systematic fashion: most warrants expire between 2 and 4 days prior to the option expiration date.[38] To see if this mismatch impacts any conclusions drawn with regards to the life cycle hypothesis, we compare the raw unadjusted premium with the three other versions of the warrant premium developed in Section 3.3.2.

Table 3.6 depicts the results of a regression of warrant premium on the time to maturity TTM (in years). The sample data is split by warrant type and further into 3 maturity bins to see if the effect changes over the life time of the warrant. In addition within each subsample, the coefficient of TTM is allowed to vary depending on the delta of the warrant. Panels A-D repeat the analysis using a different version of the warrant premium. To conserve space, we omit the coefficients of the control variables that are included. Because we explicitly use them in the next section as well, we will discuss them there in detail. Standard errors of all coefficients are based on two-way clustering by date and issuer following Thompson (2011).[39]

The difference between the unadjusted premium and the remaining three versions is quite striking. Practically all coefficients in Panel A are highly significant and positive, indicating a premium that decreases as maturity comes closer. The fact that t-statistics reach levels of 20 and more, should be reason for concern. This is clearly the manifestation of the mismatch that we described in much detail above. If the warrant expires 2-4 days earlier its time value sinks at a increasing rate, ahead of the time value of the option. This is what the TTM coefficient is picking up. Thus, it appears as if the overpricing is declining strongly.

By contrast, Panels B-D use warrant premia that are adjusted for this mismatch in one of three ways. Even though the adjustment used in Panel B is quite crude, its results are already much in line with the more data-intensive versions using intrapolation (Panel C) or extrapolation (Panel D). Our preferred method is the intrapolated premium, because the sample size is not much smaller than for the first two methods and the long-term maturity bin contains several times more observations than in the extrapolated case. All three panels show similar results across moneyness and maturity. It appears that the premia of OTM calls and more significantly, OTM puts tend to appreciate during the last 3 months of their lives at a rate of 2 and 4% annually. ITM warrants as well as OTM warrants with more than 3 months till expiration generally experience a decline in premium on the order of 2% per year.

The reversal in premium decline for close-to-maturity OTM warrants is somewhat puzzling. We can think of 2 potential explanations. The first is based on the disposition effect, i.e. behavioral. Glaser and Schmitz (2007) find that warrant investors are prone to this effect. The second is based on the transaction cost structure that prevails in most retail financial markets.

First, in a Black-Scholes framework, the change in option price with respect to time as a fraction of the option price, i.e. θ_t/P_t, is stronger for OTM options than for ITM options. This difference in

[38] We randomly compared expiration dates of warrants and other structured products in the German market and found that they cluster at precisely the same dates.

[39] Tables 3.6 and following are based on a large number of individual warrant observations; adjusting standard errors along two dimensions, especially by time, should ensure that the commonly used significance levels remain valid.

decline becomes much wider as maturity approaches. This intuitively makes sense, as ITM options have some intrinsic value, which makes up an increasing part of the total option price, while OTM options are time value only. To give one stylized example, assume $\sigma = .25$, $r_f = 0$, $S = 100$, $K_1 = 107$ and $K_2 = 94$. At one year to maturity, the daily loss of time value is 0.31% (0.25%) for a call with strike K_1 (K_2), but at two weeks to maturity the daily rate of price decline has increased to 11.0% (3.6%).

On average, it seems plausible then that contracts that are close to maturity and far-OTM have the highest likelihood of being a losing position to investors. If the current warrant holders resist redeeming those warrants because they have an aversion to realize losses, the issuer is free to charge higher premiums to newly arriving investors without having to fear large redemptions at high premiums.

The second explanation is based on the fact that transaction costs are typically a percentage of transaction volume above some minimum amount. An investor currently holding warrants has the choice between holding on until expiration, at which point he does not incur any transaction costs[40], or to sell prior to maturity incurring the cost. As in the previous explanation, on average, OTM warrants close to maturity are the most likely to be worth less than when they were originally purchased. Thus, the minimum transaction cost becomes large as a percentage of the transaction amount and can tilt the investor towards holding on to the warrants until maturity. This again leads to less selling pressure and the opportunity for the issuer to extract higher premiums from new investors.[41]

How do our results relate to other classes of SPs? Based on our findings we suggest to divide SPs into two categories, one in which time value plays a subordinate role vs. one where the price is essentially all time value. SPs with principles that are paid back at expiration fall into the first category, i.e. discount certificates, as do ITM warrants. In line with previous research mentioned at the beginning of this section, these products exhibit a declining pattern in overpricing.

In contrast, far-OTM warrants and SPs based on exotic options fall into the second category. In particular if they are close to the knock-out barrier or far from the knock-in barrier etc., most of the price consists of time value, i.e. of moving into the money. Again, in line with previous research mentioned above (Wilkens and Stoimenov (2007), Abad and Nieto (2011)), our findings suggest that the degree of overpricing for these products is not driven by the LCH.

3.4.4 The Effect of Credit Risk

Credit risk has been a long overlooked issue in the literature on structured products, at least empirically. Structured products in general, and warrants in particular, are unsecured debt obligations by the issuing bank and as such they are likely to receive a low recovery value in the case of bankruptcy. Most studies, however, ignore credit risk in their analysis. A common misconception

[40]At maturity, contracts are cash-settled and if in the money, the value is credited to the investor's account without any additional fees.

[41]We plan on formalizing both channels in a simple model in future versions of the paper.

is expressed in Bartram and Fehle (2007): 'the issuer is obligated ... to hedge all options sold. Thus, bank-issued options [i.e. warrants] are generally considered to be free of default risk.' Some studies explicitly incorporate default risk into the fair value computation. Stoimenov and Wilkens (2005) and Wilkens and Stoimenov (2007) use discount rates derived from issuer bonds instead of a default free rate to discount cash flows. Baule, Entrop, and Wilkens (2008) explicitly starts in the vulnerable options framework of Hull and White (1995) and Klein (1996). These studies decrease the fair value of structured products, but do not investigate if observed prices react to changes in credit risk.

Until recently, the default risk of large banks has not been a major concern. The credit crisis that started in 2007 will likely have changed that perception. German retail investors have become acutely aware of the risk involved in structured products after the collapse of Lehman Brothers caused a total loss in high-yield Lehman certificates. These were previously thought of and marketed to small investors as riskless[42]. Similar products underwritten by Lehman caused small investors severe losses in the U.S.[43], in Great Britain[44] and in Hong Kong[45]. Due to recent worries about the default of sovereign debt in the Euro-zone and the consequences this may have for European banks in particular, we are seeing yet another spike in the CDS spreads for European commercial banks, many of which are active participants in the issuance of warrants and structured products. Issuer default risk therefore seems to remain an important topic to investors.

Figure 3.2 shows the evolution of unsecured 1-year CDS premiums for the issuers in our sample that do have traded CDS contracts. For Citigroup we use the senior secured contract, because the unsecured contract contains too many stale quotes. Issuer default risk is clearly non-negligible for a large fraction of the time period that is covered by our sample. It seems therefore worthwhile to ask if the market for structured products (and warrants) is efficient in the sense that observed prices incorporate the effects of issuer credit risk.

It is entirely possible that retail investors are unable to properly incorporate credit risk in their demand function for a particular product and, as a consequence, issuers enjoy a form of cheap borrowing from unknowing retail clients. The results of Baule, Entrop, and Wilkens (2008) could be seen as supportive of this view. Based on 2004 data they find that average overpricing has declined relative to the data of Wilkens, Erner, and Röder (2003) from the year 2000, but that imputed credit risk constitutes a larger share of the total premium in their sample.

Instead of computing fair values of warrants that implicitly incorporate credit risk, we take a different route by investigating if observed premiums are sensitive to credit risk, i.e. if the overpricing of warrants diminishes if issuer credit risk increases. Credit risk is measured as last trading day's premium on a 1-year CDS for unsecured debt of the issuer (ScaledCDS). We describe the construction of premiums of warrants over options in Section 3.3.2 of the paper. In our analysis of the effect of credit risk and other factors on the individual warrant premium, we focus on the

[42] New York Times, October 14th, 2008, 'Lehmans Certificates Proved Risky in Germany'
[43] Wall Street Journal, December 5th, 2009, 'Investor Wins Lehman Note Arbitration'
[44] Wall Street Journal, October 27th, 2009, 'FSA to Clean Up Structured-Products Market'
[45] Wall Street Journal, May 31st, 2010, 'The Fine Print is Enlarged, But Will Investors Read It?'

interpolated version of the warrant premium because it strikes a good balance between sample size and accuracy.

Warrants in the sample can differ along several dimensions: issuer, expiration date, strike, and type. Due to this multi-dimensionality we are careful to include a number of control variables. Issuer dummies are a natural candidate, because they take care of persistent issuer fixed effects. Previous studies typically find strong differences in the average premium between issuers and our study is no exception. Presumably, issuers differ by their fixed costs and margins as well as by bid/ask spreads which may affect premia. We also include the warrant's $|\Delta|$ as well as Δ^2 (computed with the IV of the option) motivated by the observation that the warrant premium decreases with delta in a convex fashion (as documented in Figure 3.1). We exclude warrants with extreme levels of moneyness. Specifically, for each issuer, we cut off the 5% of warrants with the lowest and the highest values of delta. For the remaining warrants, we get an approximate range of $0.1 \leq |\Delta| \leq 0.8$.

Motivated by Bartram, Fehle, and Shrider (2008), we include 2 competition dummies. D_{Comp} is set to one on a given day for a particular warrant, if at least another warrant exists that expires in the same month with the same strike and type. D_{EUREX} is set to one for a particular warrant, if a EUREX option with identical features existed on a given day. Time to maturity (TTM) is meant to catch a general change in premium over the life time of the warrant. Lastly, LagWarVol is the warrant-specific trailing total volume (buys plus sells) over the previous 14 days. A large fraction of warrants outstanding do not trade at all or only very sparsely, while the most traded warrants achieve daily turnover in excess of one million EUR. It is conceivable that premia differ between warrants with low and high volume.

Table 3.7 contains the basic result with regards to the effect of credit risk on warrant premium based on the full sample. Results are reported separately by warrant type, i.e. put vs. call and by the method of computing premiums, in this case the adjusted, the interpolated and the extrapolated version. We exclude the unadjusted premium from the remaining analysis because of the biased results discovered in the previous section with regards to time to maturity.

Since Figure 3.2 shows that our sample contains a relatively calm early period and a period of heightened credit fears caused by the sudden collapse of Lehman Brothers, it seems plausible that credit risk did not play a role in the first period because investors in general and retail investors in particular were not concerned with the possibility of bank failures. Table 3.8 presents those results. As a minor difference to the previous table, the variable for credit risk is split up into 4 parts to allow different effects of credit risk on premium depending on the delta of the warrant.

In both tables, R-squared are generally quite high at between $40-50$ percent. Credit risk has a strongly significant negative effect on the warrant premium for the full sample shown in Table 3.7. The sample split in Table 3.8 reveals, however, that the effect is entirely due to the post-Lehman period. This is indicative of a structural break, where investors learned from the Lehman event and started to price in the possibility of an issuing bank defaulting.

It is possible that results in the early period are partially affected by a lack of signal in the credit risk variable rather than lacking attention by investors. However, as can be seen from Figure 3.2, while there was little cross-sectional variation of perceived credit risk among issuers in the pre-Lehman part of the sample, the time series still shows some variation from essentially no credit risk to a peak at around 200 basis points (bps). This compares to peaks of 800 bps for Citigroup and between 200 and 400 bps for other banks in the post-Lehman part. Thus, credit risk was not negligible entirely pre-Lehman and as a consequence should show up in results if it had been priced fully.

Another question is whether given the very short holding periods exhibited by retail investors in these products we should expect to see any effect at all. Obviously, we cannot exclude the possibility that many retail investors think this way, maybe due to overconfidence in their own ability to foresee the default of the issuer. However, this view is not grounded in theory as it implies that credit risk priced into an asset should depend on its turnover or trading volume. Instead, a rational investor should properly discount the credit risk of the asset over its lifetime and expect to be able to sell it at any point in the future at a price that also properly discounts the then present credit risk.

With regards to the control variables, the two delta terms pick up the declining and convex pattern visible in Figure 3.1 and the issuer dummies pick up differences in the level of overpricing that were also visible in Figure 3.1. The size of the dummy coefficients is of no importance to the analysis.

There seems to be some evidence that the existence of identical warrants by other issuers decreases the extent of the overpricing. Thus, competition does have an effect on prices. The result is much weaker for competition from an identical option listed on EUREX, which supports the idea of strong segmentation between the two markets. The level of volatility does have a strongly negative effect on premia but again, only after Lehman. The complex pattern of the TTM coefficient was the topic of the previous section, thus we will disregard it here. Lastly, lagged warrant turnover is consistently negative and significant. It is also economically important for high-volume warrants than can have around 1M EUR in daily turnover. Such a warrant may experience a decline in premium of around 2% relative to low-volume or no-volume warrants.

The Correlation of Payoff and Default

The coefficient on credit risk in the previous tables suggested that there is an effect of issuer default risk on overpricing, at least post-Lehman, but is its size reasonable? To make the coefficients comparable across maturities they scaled such that a coefficient of -1 indicates that the premium declines by 1% per year of maturity for each additional 100 bp in issuer credit risk. The coefficient for a 6-month warrant is already multiplied by a factor of two to account for the shorter maturity. A coefficient of -1 is what a naive model of default risk would predict as coefficient for a vulnerable claim. The observed values generally lie around -0.25, thus far away from the prediction.

Models of vulnerable options (Klein, 1996) suggest that options should be impacted to different degrees depending on the correlation between default risk and option payoff. Consider the most extreme case where the issuer was to issue warrants on its own stock price. An OTM call should not be affected by the issuer default risk at all, because all states of the world in which the issuer defaults are states in which the call payoff is zero anyways. An OTM put, on the other hand, potentially pays the most in states where the issuer share price goes to zero, which would likely coincide with bankruptcy. Being an unsecured debt obligation, the warrant would lose its value in line with other debt instruments of equal seniority. Therefore, moneyness should matter for the pricing impact of credit risk on warrants.

In our case, the warrants are issued on the DAX, the German headline market index. It is not difficult to assume that the default risk of issuers, many of which German banks, is negatively related to market returns, in particular, because the time period was already termed a banking crisis at the time. Thus, DAX calls pay out in states where the market has risen, which are most likely states of the world in which banks survived. DAX puts pay out in states of further market declines, a scenario that could have been likely triggered by bad news from the banking sector.

The credit risk variable in Table 3.8 is split into moneyness regions. Contrary to theoretical predictions, far-OTM puts are the least sensitive to credit risk; in fact, the post-Lehman sample exhibits declining coefficients for credit risk from ITM to OTM. Calls on the other hand exhibit an increasing for the full sample and a flat pattern for the post Lehman sample. Next we split the sample into maturity bins to test a second natural prediction of a default risk model, namely that the sensitivity to default risk should increase with the time to maturity. We maintain the division of the credit risk variable according to the delta of the warrant.

Table 3.9 shows the results. While the patterns are rarely monotonic we would argue that it appears as if for both types sensitivity to credit risk is more pronounced for OTM warrants. So, while we find evidence of credit risk being a factor across all maturities, we cannot argue that this is because of investors rationally relating the payoff of warrants with the likelihood of the issuer's default. Further, with maybe the exception of OTM calls, it is hard to argue that there is an increase in default risk sensitivity for longer maturity contracts.

We conclude that investors do not fully incorporate the effects of credit risk into their demand. The Lehman bankruptcy may have served as a wake-up call to retail investors that issuer credit risk is not negligible but the size of the effect suggests that issuers can continue to use structured products as a source of cheap financing.

The issuance of SPs should not be underestimated in this regard. According to the monthly market volume statistics issued by the industry group 'Deutscher Derivate Verband' (DDV) [46], open interest in structured products in the German market amounted to 76 billion EUR in the month of May 2009, the last month of the sample period. The number is based on 14 reporting member banks. The DDV estimates the total open interest at around 90 billion EUR for that time. In April 2010, estimates of total open interest were in excess of 100 billion EUR. Curiously, about 60% of the total amount outstanding is invested in capital protection certificates that are marketed as conservative and safe investments with the primary goal of principal protection. Similarly, retail structured products in the U.S. reported sales volumes of over $100 billion USD in 2007[47].

[46] http://www.derivateverband.de/DEU/Presse

[47] Wall Street Journal, May 28th, 2009, 'Twice Shy On Structured Products?'

3.5 Potential for Policy?

Bethel and Ferrell (2007) focus on policy in the U.S. market for structured products and advocate restricting the pool of eligible investors for which structured products are suitable and improving transparency and disclosure requirements to avoid uninformed retail investors from making costly mistakes. However, they are quick to point out that regulation must also be mindful of the costs imposed on issuers and retail investors and their response to circumvent new restrictions.

The focus of the present paper are novel empirical findings in the market for structured products. Nevertheless, we shall quickly comment on whether policy implications arise from our findings specifically. With regards to credit risk, the implication is relatively straightforward. Had retail investors that were seeking safety known that Lehman certificates carry the full credit risk of the issuer, it seems likely that fewer had chosen those particular instruments for this purpose. The issue is thus one of disclosure of all risks in clear terms and can be improved by increasing fines and accountability on the part of the financial advisor selling these products. Forcing the seller to confirm the suitability of the product for the individual investor at the point of sale is another potential avenue.

As for the general and time-varying overpricing, the issue is much less obvious. As the previous literature has pointed out, some degree of overpricing maybe justified due to the services provided and the associated costs that arise to the issuer, i.e. packaging services, binding liquidity provision, marketing and the fixed setup costs in creating the technology, products and know-how. In this context it is interesting that there are no discernible barriers to entry (other than fixed entry costs) that would keep competitors from entering this market. Thus, it is a distinct possibility that to a large extent the profits generated by the issuers are compensation for their investment and business risk rather than rents.

It is also difficult to argue that the removal of the short-selling constraint would lead to overall improvements. Given that short-selling requires someone to lend out these products, who would be a willing party? Retail investors with their very short holdings periods are not suitable candidates and the issuers themselves cannot be forced to lend out shares. Further, the infrastructure of the market would have to change significantly. One of the main advantages of this retail-driven market is the absence of complex margin accounts, which short-selling would require. Such a step would move SP markets already very close to regular options exchanges and possible lead to their demise.

3.6 Conclusion and Outlook

This study investigates the price setting behavior of issuers of structured products in the secondary market. Issuers are the sole liquidity providers for their products and do not allow short-selling. In particular, for the securities that we study, bank-issued warrants in the German market, trading is only possible with the issuer being on one side of the transaction. This allows the issuer to exert great power over pricing.

Previous literature has consistently reported overpricing in structured products in general, for the U.S. (Bernard et al., 2011), for Germany (Wilkens et al., 2003) and for Switzerland (Grünbichler and Wohlwend, 2005), and for warrants in particular (Horst and Veld, 2008; Abad and Nieto, 2011) when compared to products on derivatives exchanges.

The extent of overpricing can be sufficiently large to ensure that expected returns lie below the risk-free rate (Henderson and Pearson, 2011) and thus there is no reason for rational investors to buy some of these products. Other studies point out how issuers take advantage of investors' susceptibility to certain mental errors like over-optimism (Bernard et al., 2011), how they optimally increase complexity to maintain overpricing (Carlin, 2009) and increase search costs (Dorn, 2010).

We add to this literature by pointing out a number of ways in which issuers take advantage of the demand of investors for certain payoffs as well as the lack of financial literacy. Our results are as follows: First, retail investors exhibit a preference for high-leverage, OTM warrants. Issuers exploit this demand by overpricing those securities the most. Issuers seem to benefit either from investors that cannot easily spot the overpricing or from investors with sufficiently optimistic beliefs that makes them willing to pay higher premiums. Among ITM warrants, puts are more overpriced than calls, because of the general lack of ways to express negative views on stock markets in Germany.

Second, issuers are able to predict net demand at the daily horizon, because transactions, especially sales, are heavily influenced by recent returns. Knowing future net demand, issuers opportunistically adjust prices ahead of time and thus are able to extract additional gains in excess of the officially quoted bid/ask spread. In turn, this lowers realized profits for investors.

Third, for warrants other than close-to-expiration OTM warrants we find evidence that the extent of overpricing is decreasing over the life time of the warrant. Thus, issuers extract additional gains from investors holding warrants for longer periods, diminishing investors' returns. In contrast, close-to-expiration OTM warrants appear to decrease less or even increase in overpricing. We suggest that it is the relative importance of time value as a share of the total price that gives rise to this divergent behavior and propose two channels through which the relatively larger rate of time decay keeps existing holders of OTM warrants from selling, thus enabling issuers to decrease prices less.

Last, we investigate the effect of issuer credit risk on the extent of overpricing. We find that only since the collapse of Lehman Brothers, is there a discount in prices due to credit risk. However, the effect is too small relative to vulnerable options models and additional predictions by these models cannot be confirmed in the data. Once again, issuers seem to be able to profit from the lack of financial knowledge of investors. We point out that given the large amount of structured products outstanding (currently 100 billion EUR in Germany alone) underestimation of credit risk amounts to a non-negligible source of cheap financing for issuing banks.

3.6.1 Implications for Other Structured Products

Admittedly, warrants are the most simple structured product available. It is an open question if our results carry through to other classes of these instruments. Given that structured products share most if not all institutional features like trading mechanisms, secondary markets, liquidity provision, marketing channels and issuers, we are confident that the main findings apply to all structured products. In particular, success of the practice of adjusting prices in anticipation of future demand is completely independent of the instrument, rather it depends on two factors. First, that investors are forced to trade with the issuer either by market design (as in Germany) or by lack of liquidity (as for some U.S. products) and second, that daily demand can be anticipated reasonably well.

Our finding that the effect of maturity on overpricing depends on the importance of time value relative to total price is also not tied to warrants alone. We argue that all structured products that pay back the principle will behave similar to ITM warrants because for both the time value of the option part of the product gets smaller as maturity approaches. In contrast, SPs that payoff nothing in some states of the world (knock-out, knock-in features) are more similar to OTM warrants in that their price is mostly driven by time value.

Finally, with regards to credit risk, we see no reason why retail investors would fail to correctly price default risk into warrants, but would correctly discount other structured products carrying exactly the same risk.

We are looking forward to seeing an investigation into the predictability of demand for other structured products as well as the potential exploitation by issuers. We would also be interested in understanding cross-sectional patterns in demand for other structured products. For instance, is there a preferred habitat for retail investors with regards to leverage (in speculative, convex-payoff products) or with regards to the proximity to the participation cap (in concave products) or the knock-out barrier (in more exotic products)?

Table 3.1: Relative Transaction Frequencies

This table shows the relative frequency of transactions (in percent) by moneyness category (out-of-the-money (OTM); at-the-money (ATM); in-the-money (ITM)) for several maturity ranges (1-3 months, 4-6 months and 7 or more months) and for the full sample. Numbers represent percent shares of row totals, i.e. within maturity bin, with the exception of the last column 'All', where they represent the share of each maturity of the total sample. The analysis is done separately for warrant transaction on EUWAX/SCoach and options on EUREX and further split by calls and puts.

	OTM			**ATM**	**ITM**			**All**
	> 25%	15-25%	5-15%	± 5%	5-15%	15-25%	> 25%	
Panel A: Calls								
Warrants								
Maturity (in months)								
1 - 3m	6.3	9.2	28.6	40.0	2.3	0.2	12.9	**37.9**
4 - 6m	15.3	19.4	30.1	20.4	2.1	0.3	12.1	**19.7**
≥ 7m	36.3	14.4	16.6	15.9	2.9	0.7	12.9	**42.3**
All	**20.8**	**13.4**	**23.8**	**25.9**	**2.5**	**0.5**	**12.7**	**100.0**
Options								
Maturity (in months)								
1 - 3m	1.2	4.2	34.0	58.4	1.7	0.1	0.2	**84.3**
4 - 6m	8.6	17.5	42.9	26.7	2.9	0.5	0.5	**7.9**
≥ 7m	28.4	16.3	24.5	22.8	4.6	0.9	2.1	**7.6**
All	**3.8**	**6.2**	**34.0**	**53.1**	**2.0**	**0.2**	**0.3**	**100.0**
Panel B: Puts								
Warrants								
Maturity (in months)								
1 - 3m	3.4	7.6	26.9	39.7	6.8	1.4	13.8	**44.6**
4 - 6m	7.2	9.6	22.2	34.4	10.6	2.7	13.0	**22.4**
≥ 7m	15.1	10.9	21.3	25.9	9.4	3.9	13.2	**32.9**
All	**8.1**	**9.1**	**24.0**	**34.0**	**8.5**	**2.5**	**13.4**	**100.0**
Options								
Maturity (in months)								
1 - 3m	1.4	5.9	34.6	53.9	2.8	0.5	0.5	**84.6**
4 - 6m	7.4	14.4	33.2	36.8	5.5	1.0	1.4	**8.0**
≥ 7m	18.9	14.7	25.9	27.2	8.0	2.4	2.6	**7.3**
All	**3.2**	**7.3**	**33.8**	**50.5**	**3.4**	**0.7**	**0.7**	**100.0**

Table 3.2: Day-over-Day Demand Regression

This table show the results of a regression of aggregate daily demand for warrants on a number of lagged demand and market return and volatility variables. Regressions are done separately by type of warrant (calls and puts) and type of transaction (buys and sells). Aggregate daily demand on day t is defined as the sum of all transactions occurring between 1730CET of the previous day and 1730CET of day t. Explanatory variables are total unsigned warrants volume over the past 2 weeks ($\text{TotWarVol}_{t-14,t-1}$), lagged change in and level of volatility (ΔVDAX_{t-1}, VDAX_{t-1}), lagged demand by type of transaction and type of warrant (e.g. Buys^{C}_{t-1}) as well as lagged and contemporaneous market returns (e.g. Ret^{-}_{t-1}). A negative sign indicates that the variable is equal to the market return of that period if negative and zero otherwise; a positive sign indicates the opposite. Standard errors are computed following Newey and West (1987, 1994).

	Calls		**Puts**	
	Buys	**Sells**	**Buys**	**Sells**
$\text{TotWarVol}_{t-14,t-1}$	0.43	0.41	0.68	0.77
	***[6.07]	***[6.81]	***[4.73]	***[5.00]
ΔVDAX_{t-1}	7.19	8.52	5.50	15.98
	*[1.78]	***[3.43]	[1.21]	**[2.07]
VDAX_{t-1}	-3.67	-4.13	-3.13	-3.63
	***[-4.63]	***[-5.89]	***[-2.89]	**[-2.38]
Buys^{C}_{t-1}	0.12	0.12		
	***[5.81]	***[5.33]		
Sells^{C}_{t-1}	0.04	0.05		
	*[1.87]	**[2.45]		
Buys^{P}_{t-1}			0.11	0.09
			***[4.55]	***[3.22]
Sells^{P}_{t-1}			0.07	0.11
			***[2.91]	***[3.14]
Ret^{-}_{t-1}	-10.21	10.78	-16.87	-38.20
	[-1.09]	[1.33]	[-1.59]	**[-2.57]
Ret^{+}_{t-1}	13.04	36.20	-4.62	-42.11
	[1.38]	***[3.55]	[-0.53]	***[-4.07]
Ret^{-}_{t}	-31.90	-5.51	-40.34	-107.60
	***[-4.93]	[-0.94]	***[-3.56]	***[-5.37]
Ret^{+}_{t}	3.50	63.41	18.12	-32.14
	[0.57]	***[7.00]	[1.28]	***[-2.61]
# Observations	504	504	504	504
R^2	0.391	0.523	0.667	0.695

Table 3.3: Intraday Demand Prediction

This table show the results of a predictive regression of aggregate intraday demand for warrants on a number of lagged and overnight demand, lagged and overnight market return and lagged volatility variables. Regressions are done separately by type of warrant (calls and puts) and type of transaction (buys and sells). Aggregate intraday demand on day t is defined as the sum of all transactions occurring between 930CET and 1730CET of day t. Explanatory variables are total unsigned warrants volume over the past 2 weeks ($\text{TotWarVol}_{t-14,t-1}$), lagged change in and level of volatility (ΔVDAX_{t-1}, VDAX_{t-1}), lagged signed market return (Ret^{-}_{t-1} and Ret^{+}_{t-1}), overnight signed market return (Ret^{-}_{ON} and Ret^{+}_{ON}), cumulative demand over the previous three trading days by type ($\text{Buys}^{C}_{t-3,t-1}$ etc.) as well as overnight demand by type (Buys^{C}_{ON} etc.). The definitions of overnight returns and demand are consistent with the times that define intraday demand, i.e. overnight return is defined as the change in the level of the DAX between 1730CET of trading day $t-1$ and 930CET of day t. A negative sign indicates that the variable is equal to the market return of that period if negative and zero otherwise; a positive sign indicates the opposite. Standard errors are computed following Newey and West (1987, 1994).

	Calls		**Puts**	
	Buys	**Sells**	**Buys**	**Sells**
$\text{TotWarVol}_{t-14,t-1}$	0.3912	0.4517	0.5028	0.5232
	***[5.77]	***[5.90]	***[4.02]	***[3.15]
Δ VDAX_{t-1}	7.2715	7.195	6.6414	5.7717
	***[2.69]	***[2.69]	[1.44]	[0.76]
VDAX_{t-1}	-2.3378	-3.1864	-1.8793	-2.3297
	***[-3.41]	***[-3.98]	*[-1.90]	*[-1.74]
Ret^{-}_{t-1}	11.7337	19.1834	8.0943	14.2664
	*[1.68]	**[2.25]	[0.68]	[1.23]
Ret^{+}_{t-1}	-0.9331	12.1836	-18.3902	-43.4377
	[-0.16]	*[1.71]	**[-2.12]	***[-5.41]
Ret^{-}_{ON}	-16.7453	0.2643	-27.5442	-100.191
	[-1.58]	[0.03]	***[-2.77]	***[-6.51]
Ret^{+}_{ON}	-2.835	58.9149	7.0739	-74.673
	[-0.29]	***[3.90]	[0.37]	***[-4.46]
$\text{Buys}^{C}_{t-3,t-1}$	0.0667	0.0512		
	***[3.28]	**[2.13]		
$\text{Sells}^{C}_{t-3,t-1}$	0.0273	0.0196		
	[1.33]	[1.04]		
$\text{Buys}^{P}_{t-3,t-1}$			0.0664	0.0105
			***[2.84]	[0.39]
$\text{Sells}^{P}_{t-3,t-1}$			0.0601	0.0962
			***[2.83]	***[3.65]
Buys^{C}_{ON}	0.3352	0.3237		
	*[1.85]	[1.53]		
Sells^{C}_{ON}	0.0742	0.2262		
	[0.44]	*[1.65]		
Buys^{P}_{ON}			0.494	0.3644
			**[2.45]	[1.31]
Sells^{P}_{ON}			-0.0217	0.71
			[-0.17]	**[2.17]
# Observations	504	504	504	504
R^2	0.32	0.411	0.603	0.614

Table 3.4: Daily Change in BKM Skewness

This table shows the results of regressions of day-over-day changes in implied skewness of warrant chains on contemporaneous demand measures and control variables. A warrant chain is the set of all warrants by the same issuer with the same expiration date. Implied skewness is defined as in Bakshi et al. (2003) using all OTM warrants per warrant chain. Change in skewness is computed as the difference in the level of implied skewness at 1730CET of day $t-1$ and 1730CET of day t. Control variables are the time to maturity of the warrant chain (TTM), the lagged level of the skewness of the warrant chain (SKEW_{t-1}) and the market return on day $t-1$ (Ret_{t-1}). Explanatory variables are contemporaneous net demands (i.e. buys minus sells) by type of warrant and either aggregated by issuer ($\mathrm{ISS\_NetDem}_t^{C/P}$), by expiration ($\mathrm{EXP\_NetDem}_t^{C/P}$) or summed over all expirations and issuers ($\mathrm{TOT\_NetDem}_t^{C/P}$). Regressions are OLS with 2-way clustered standard errors following Thompson (2011), clustered by issuer and date.

	(1)	(2)	(3)	(4)	(5)
TTM	-0.042	-0.043	-0.046	-0.044	-0.045
	***[-3.52]	***[-3.71]	***[-4.09]	***[-3.84]	***[-4.02]
SKEW_{t-1}	-0.058	-0.059	-0.058	-0.060	-0.059
	***[-6.47]	***[-6.69]	***[-6.61]	***[-6.54]	***[-6.54]
Ret_{t-1}	-0.061	0.104	0.181	0.652	0.644
	[-0.29]	[0.49]	[0.91]	***[3.07]	***[3.00]
$\mathrm{ISS\_NetDem}_t^C$		0.049			0.010
		***[3.44]			*[1.85]
$\mathrm{ISS\_NetDem}_t^P$		-0.027			0.008
		*[-1.87]			[0.88]
$\mathrm{EXP\_NetDem}_t^C$			0.025		0.003
			***[4.04]		[0.56]
$\mathrm{EXP\_NetDem}_t^P$			-0.035		-0.011
			***[-6.86]		***[-3.17]
$\mathrm{TOT\_NetDem}_t^C$				0.010	0.009
				***[3.45]	***[3.07]
$\mathrm{TOT\_NetDem}_t^P$				-0.010	-0.009
				***[-6.70]	***[-5.88]
# Issuers			9		
# Days			488		
# Observations			12,608		
R^2	0.026	0.046	0.066	0.118	0.121

Table 3.5: Overnight Change in BKM Skewness

This table shows the results of predictive regressions of overnight changes in implied skewness of warrant chains on future intraday (actual and predicted) demand measures as well as lagged and overnight demand and control variables. The overnight change in skewness is computed as the difference between the level of implied skewness at 1730CET of day $t-1$ and 930CET of day t. Demand measures are split by type of warrant (i.e. puts vs. calls) and by time period. Lagged demand is the cumulative total net demand (buys minus sells) over the previous three trading days ($\text{TOT\_NetDem}^{C/P}_{t-3,t-1}$). Overnight demand is the sum of all buys minus all sells of warrants between 1730CET of trading day $t-1$ and 930CET of day t ($\text{TOT\_NetDem}^{C/P}_{ON}$). Future intraday demand is either total aggregated net demand by type ($\text{TOT\_NetDem}^{C/P}_{Day}$) or split by expiration date ($\text{EXP\_NetDem}^{C/P}_{Day}$). In columns (2a-c) the actual realized intraday demand is used; in column (3a-c) we use the fitted values derived from the regressions in Table 3.3 for total demand and from unreported, but identical regressions for demand by expiration date. Regressions are OLS with 2-way clustered standard errors following Thompson (2011), clustered by issuer and date. Numbers of clusters and observations are as in Table 3.4, but omitted to conserve space.

	Lagged only	Actual Day Demand			Predicted Day Demand		
	(1)	(2a)	(2b)	(2c)	(3a)	(3b)	(3c)
$\text{EXP\_NetDem}^{C}_{\text{Day}}$			0.007	0.000		0.430	0.242
			**[2.23]	[-0.08]		***[7.53]	***[4.78]
$\text{EXP\_NetDem}^{P}_{\text{Day}}$			-0.014	-0.001		-0.078	-0.036
			***[-3.84]	[-0.45]		***[-6.80]	***[-3.15]
$\text{TOT\_NetDem}^{C}_{\text{Day}}$		0.001		0.001	0.095		0.072
		[0.74]		[0.76]	***[8.93]		***[6.72]
$\text{TOT\_NetDem}^{P}_{\text{Day}}$		-0.006		-0.006	-0.012		0.001
		***[-4.54]		***[-4.57]	**[-2.22]		[0.24]
$\text{TOT\_NetDem}^{C}_{t-3,t-1}$	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002
	**[-2.50]	**[-2.12]	**[-2.42]	**[-2.12]	***[-2.87]	***[-2.75]	***[-2.85]
$\text{TOT\_NetDem}^{P}_{t-3,t-1}$	0.000	0.001	0.001	0.001	0.000	0.000	0.000
	[0.82]	*[1.86]	[1.15]	*[1.86]	[0.90]	[1.39]	[0.57]
$\text{TOT\_NetDem}^{C}_{\text{ON}}$	0.017	0.013	0.016	0.013	0.004	0.000	-0.001
	**[2.49]	**[2.22]	**[2.45]	**[2.22]	[0.91]	[-0.10]	[-0.13]
$\text{TOT\_NetDem}^{P}_{\text{ON}}$	0.008	0.010	0.009	0.010	0.005	0.007	0.004
	**[2.00]	**[2.26]	**[2.10]	**[2.26]	[1.43]	**[2.27]	[1.16]
TTM	-0.010	-0.011	-0.011	-0.011	-0.015	0.004	-0.009
	[-1.16]	[-1.23]	[-1.30]	[-1.23]	*[-1.83]	[0.40]	[-1.00]
SKEW_{t-1}	-0.025	-0.024	-0.024	-0.024	-0.029	-0.015	-0.027
	***[-3.69]	***[-3.78]	***[-3.68]	***[-3.78]	***[-4.43]	***[-2.73]	***[-4.39]
Ret_{t-1}	-0.093	0.012	-0.070	0.012	0.531	0.270	0.507
	[-0.50]	[0.06]	[-0.38]	[0.06]	**[2.50]	[1.46]	**[2.46]
R^2	0.034	0.068	0.043	0.068	0.145	0.125	0.155

Table 3.6: The Effect of Time to Maturity

The table shows the coefficients of the time to maturity variable in a regression on the warrant premium. Panels A-D represent the four different versions of the premium computation. Within each panel, the sample was split by type (i.e. call vs. put) and further by maturity range (short-term, medium-term, long-term). In addition, the coefficients are allowed to differ by the delta of the warrant within each regression. To conserve space, the coefficients of the control variables were omitted. They include issuer dummies, $|\Delta|$ and Δ^2, dummies for competition from other warrants (D_{Comp}) and from EUREX (D_{EUREX}), yesterday's market volatility (VDAX_{t-1}), the warrant's cumulative volume over the previous 2 weeks (LagWarVol) and lagged issuer CDS premium (CDS_{t-1}). Regressions are OLS with 2-way clustered standard errors following Thompson (2011), clustered by issuer and date.

	Call Warrants			**Put Warrants**		
	Short 1-3 months	**Medium** 4-6 months	**Long** > 7 months	**Short** 1-3 months	**Medium** 4-6 months	**Long** > 7 months
Panel A: Unadjusted						
$\lvert\Delta\rvert < .25$	0.297	0.079	0.036	0.248	0.047	0.001
	***[22.08]	***[5.61]	***[4.20]	***[22.68]	***[6.60]	[0.31]
$.25 < \lvert\Delta\rvert < .35$	0.250	0.065	0.029	0.200	0.043	0.005
	***[24.69]	***[5.43]	***[4.65]	***[24.79]	***[6.72]	[1.51]
$.35 < \lvert\Delta\rvert < .50$	0.221	0.063	0.025	0.174	0.046	0.009
	***[27.06]	***[6.04]	***[5.49]	***[24.49]	***[7.50]	***[2.58]
$.50 < \lvert\Delta\rvert$	0.180	0.055	0.016	0.122	0.046	0.009
	***[24.87]	***[5.70]	***[3.80]	***[16.73]	***[7.35]	***[2.62]
Panel B: Adjusted						
$\lvert\Delta\rvert < .25$	-0.047	0.037	0.026	-0.070	0.009	-0.006
	***[-3.83]	***[2.60]	***[3.05]	***[-7.08]	[1.32]	[-1.21]
$.25 < \lvert\Delta\rvert < .35$	-0.041	0.027	0.021	-0.059	0.011	0.000
	***[-4.54]	**[2.29]	***[3.42]	***[-8.25]	*[1.75]	[-0.10]
$.35 < \lvert\Delta\rvert < .50$	-0.007	0.031	0.020	-0.022	0.020	0.004
	[-0.98]	***[2.99]	***[4.28]	***[-3.71]	***[3.34]	[1.20]
$.50 < \lvert\Delta\rvert$	0.006	0.027	0.011	-0.002	0.026	0.006
	[0.87]	***[2.86]	***[2.77]	[-0.25]	***[4.29]	[1.64]
Panel C: Intrapolated						
$\lvert\Delta\rvert < .25$	-0.020	0.029	0.027	-0.043	0.013	-0.003
	[-1.57]	**[2.28]	**[2.49]	***[-4.11]	*[1.83]	[-0.70]
$.25 < \lvert\Delta\rvert < .35$	-0.022	0.023	0.023	-0.037	0.015	0.002
	**[-2.46]	**[2.17]	***[3.27]	***[-4.83]	**[2.27]	[0.48]
$.35 < \lvert\Delta\rvert < .50$	-0.003	0.025	0.018	-0.014	0.021	0.005
	[-0.45]	***[2.69]	***[3.86]	**[-2.20]	***[3.40]	[1.46]
$.50 < \lvert\Delta\rvert$	0.005	0.020	0.011	-0.008	0.024	0.004
	[0.76]	**[2.36]	***[2.94]	[-1.17]	***[3.70]	[1.06]
Panel D: Extrapolated						
$\lvert\Delta\rvert < .25$	-0.036	0.080	0.081	-0.068	0.007	0.030
	**[-2.35]	***[2.64]	[0.48]	***[-6.13]	[0.50]	**[2.32]
$.25 < \lvert\Delta\rvert < .35$	-0.033	0.065	-0.188	-0.056	0.016	0.043
	***[-2.86]	**[2.56]	[-0.90]	***[-6.72]	[1.12]	***[3.32]
$.35 < \lvert\Delta\rvert < .50$	-0.004	0.069	-0.237	-0.021	0.029	0.044
	[-0.44]	***[3.11]	[-0.94]	***[-2.89]	**[2.21]	***[2.93]
$.50 < \lvert\Delta\rvert$	0.004	0.053	-0.139	-0.007	0.036	0.041
	[0.54]	***[2.63]	[-0.91]	[-0.97]	***[2.66]	**[2.38]

Table 3.7: Individual Warrant Premium

Warrant premiums are computed by three different methods: adjusted, intrapolated and extrapolated. Explanatory variables are issuer dummies, $|\Delta|$ and Δ^2, dummies for competition from other warrants (D_{Comp}) and from EUREX (D_{EUREX}), lagged market volatility (VDAX_{t-1}), lagged issuer CDS premium (CDS_{t-1}), the warrant's time until expiration (TTM) and its cumulative volume over the previous 2 weeks (LagWarVol). Regressions are OLS with 2-way clustered standard errors following Thompson (2011), clustered by issuer and date.

	Adjusted		**Intrapolated**		**Extrapolated**	
	Call	**Put**	**Call**	**Put**	**Call**	**Put**
D_{BNP}	0.133	0.102	0.130	0.096	0.148	0.115
	***[37.85]	***[43.58]	***[32.79]	***[39.26]	***[21.64]	***[33.21]
D_{CBK}	0.110	0.089	0.108	0.080	0.124	0.098
	***[32.42]	***[40.78]	***[28.00]	***[35.56]	***[18.37]	***[30.12]
D_{CITI}	0.135	0.114	0.131	0.104	0.156	0.128
	***[38.06]	***[46.65]	***[31.43]	***[40.93]	***[22.41]	***[35.29]
D_{DRBK}	0.125	0.101	0.122	0.092	0.144	0.115
	***[35.43]	***[40.60]	***[29.83]	***[35.43]	***[20.21]	***[31.08]
D_{DTBK}	0.116	0.089	0.113	0.080	0.130	0.098
	***[33.42]	***[39.39]	***[27.97]	***[33.87]	***[18.78]	***[28.58]
D_{SCGN}	0.154	0.112	0.148	0.104	0.162	0.118
	***[30.93]	***[35.80]	***[28.58]	***[32.85]	***[18.68]	***[27.83]
D_{TRBK}	0.104	0.082	0.102	0.073	0.119	0.091
	***[30.01]	***[37.22]	***[25.77]	***[32.52]	***[17.67]	***[27.82]
$\|\Delta\|$	-0.336	-0.211	-0.303	-0.175	-0.366	-0.222
	***[-31.84]	***[-34.57]	***[-26.45]	***[-26.70]	***[-27.87]	***[-26.80]
Δ^2	0.267	0.159	0.235	0.133	0.305	0.171
	***[23.68]	***[25.60]	***[19.05]	***[18.31]	***[20.95]	***[17.87]
D_{Comp}	-0.004	-0.002	-0.005	-0.001	-0.009	-0.004
	**[-2.54]	*[-1.75]	**[-2.50]	[-0.77]	***[-2.77]	**[-2.01]
D_{EUREX}	0.003	-0.002	0.003	-0.003	0.007	-0.001
	*[1.68]	[-1.63]	[1.12]	**[-2.03]	[1.48]	[-0.24]
VDAX_{t-1}	0.006	-0.010	-0.012	-0.024	-0.012	-0.029
	[1.15]	***[-3.82]	***[-2.72]	***[-9.50]	**[-2.22]	***[-7.57]
CDS_{t-1}	-0.234	-0.201	-0.224	-0.197	-0.355	-0.244
	***[-6.03]	***[-6.70]	***[-6.23]	***[-6.42]	***[-7.05]	***[-6.23]
TTM	0.004	-0.010	0.008	-0.006	0.005	-0.015
	**[2.37]	***[-6.84]	***[4.25]	***[-3.94]	[0.75]	***[-2.76]
LagWarVol	-0.003	-0.002	-0.003	-0.002	-0.005	-0.002
	***[-5.88]	***[-5.28]	***[-6.12]	***[-5.17]	***[-5.44]	***[-4.61]
# Obs.	93,957	123,035	80,309	107,837	44,251	57,453
R^2	0.433	0.469	0.406	0.424	0.420	0.492

Table 3.8: Warrant Premium by Delta and Time Period

Results are presented for the full sample, and for two subperiods: One from June 2007 - September 2008, the other from October 2008 - May 2009. This roughly coincides with the bankruptcy of Lehman brothers on September 15th, 2008. Warrant premiums are computed by intrapolation method. Explanatory variables are issuer dummies, $|\Delta|$ and Δ^2, dummies for competition from other warrants (D_{Comp}) and from EUREX (D_{EUREX}), yesterday's market volatility (VDAX_{t-1}), the warrant's time until expiration and its cumulative volume over the previous 2 weeks (LagWarVol). Lagged issuer CDS premium (CDS_{t-1}) is split up into 4 variables to allow for different slopes conditional on the delta of the warrant. Regressions are OLS with 2-way clustered standard errors following Thompson (2011), clustered by issuer and date.

	Call Premium			**Put Premium**		
	Full Sample	**Before Lehman**	**Post Lehman**	**Full Sample**	**Before Lehman**	**Post Lehman**
D_{BNP}	0.132	0.125	0.125	0.097	0.096	0.087
	***[30.72]	***[14.76]	***[16.43]	***[35.57]	***[16.10]	***[22.70]
D_{CBK}	0.111	0.106	0.096	0.081	0.082	0.066
	***[25.87]	***[13.24]	***[13.01]	***[31.75]	***[15.40]	***[17.28]
D_{CITI}	0.134	0.133	0.114	0.105	0.108	0.086
	***[29.20]	***[16.23]	***[13.76]	***[37.43]	***[20.37]	***[18.56]
D_{DRBK}	0.125	0.127	0.104	0.093	0.099	0.073
	***[27.97]	***[15.02]	***[15.13]	***[32.22]	***[17.78]	***[17.97]
D_{DTBK}	0.116	0.122	0.091	0.081	0.086	0.062
	***[26.02]	***[14.89]	***[12.85]	***[30.34]	***[15.59]	***[16.50]
D_{SCGN}	0.151	0.145	0.164	0.105	0.102	0.114
	***[27.56]	***[15.83]	***[15.86]	***[31.13]	***[17.39]	***[16.22]
D_{TRBK}	0.105	0.105	0.087	0.073	0.073	0.061
	***[24.06]	***[12.51]	***[12.66]	***[29.60]	***[12.88]	***[16.28]
$\|\Delta\|$	-0.310	-0.307	-0.252	-0.178	-0.202	-0.127
	***[-22.37]	***[-15.00]	***[-11.10]	***[-22.63]	***[-18.53]	***[-11.68]
Δ^2	0.235	0.231	0.189	0.136	0.153	0.099
	***[16.30]	***[11.76]	***[7.84]	***[16.68]	***[13.29]	***[9.03]
D_{Comp}	-0.004	-0.007	-0.004	-0.001	-0.003	0.001
	**[-2.40]	***[-2.58]	***[-2.76]	[-0.76]	[-1.57]	[0.80]
D_{EUREX}	0.002	-0.002	0.003	-0.003	0.000	-0.004
	[1.05]	[-0.39]	[1.15]	**[-2.03]	[-0.00]	***[-2.97]
VDAX_{t-1}	-0.012	0.020	-0.012	-0.024	-0.011	-0.029
	***[-2.79]	[0.59]	[-1.55]	***[-9.46]	[-0.46]	***[-6.71]
TTM	0.008	0.001	0.015	-0.006	-0.014	0.002
	***[4.31]	[0.62]	***[5.57]	***[-3.93]	***[-6.38]	[1.29]
LagWarVol	-0.003	-0.003	-0.002	-0.002	-0.002	-0.002
	***[-6.16]	***[-4.77]	**[-2.54]	***[-5.16]	***[-4.30]	***[-4.99]
CDS_{t-1}						
$\|\Delta\| < .25$	**-0.302**	**0.078**	**-0.169**	**-0.210**	**0.196**	**-0.048**
	*****[-5.14]**	**[0.18]**	***[-1.92]**	*****[-4.67]**	**[0.92]**	**[-0.84]**
$.25 < \|\Delta\| < .35$	**-0.296**	**-0.129**	**-0.233**	**-0.216**	**0.063**	**-0.136**
	*****[-6.92]**	**[-0.37]**	*****[-3.38]**	*****[-6.74]**	**[0.31]**	*****[-3.08]**
$.35 < \|\Delta\| < .50$	**-0.194**	**-0.083**	**-0.204**	**-0.167**	**0.199**	**-0.165**
	*****[-5.78]**	**[-0.28]**	*****[-3.43]**	*****[-6.04]**	**[0.95]**	*****[-4.17]**
$.50 < \|\Delta\|$	**-0.118**	**-0.133**	**-0.157**	**-0.193**	**0.007**	**-0.209**
	*****[-4.09]**	**[-0.47]**	*****[-3.11]**	*****[-5.99]**	**[0.03]**	*****[-5.12]**
# Observations	80,309	39,957	40,354	107,837	58,723	49,114
R^2	0.407	0.502	0.351	0.425	0.468	0.400

Table 3.9: Warrant Premium by Delta and Time to Maturity

Warrant premiums are computed by intrapolation method. Explanatory variables are issuer dummies, $|\Delta|$ and Δ^2, dummies for competition from other warrants (D_{Comp}) and from EUREX (D_{EUREX}), yesterday's market volatility (VDAX_{t-1}), the warrant's time until expiration and its cumulative volume over the previous 2 weeks (LagWarVol). Lagged issuer CDS premium (CDS_{t-1}) is split up into 4 variables that condition on the delta of the warrant. Regressions are OLS with 2-way clustered standard errors following Thompson (2011), clustered by issuer and date.

	Call Premium			**Put Premium**		
	Short 1-3 months	**Medium** 4-6 months	**Long** > 7 months	**Short** 1-3 months	**Medium** 4-6 months	**Long** > 7 months
D_{BNP}	0.145	0.109	0.151	0.110	0.089	0.077
	***[17.99]	***[14.73]	***[12.85]	***[25.22]	***[20.46]	***[14.76]
D_{CBK}	0.120	0.087	0.135	0.091	0.070	0.068
	***[15.10]	***[11.74]	***[11.68]	***[22.08]	***[17.19]	***[13.43]
D_{CITI}	0.149	0.108	0.147	0.121	0.092	0.085
	***[18.60]	***[14.10]	***[12.47]	***[27.59]	***[21.67]	***[15.72]
D_{DRBK}	0.143	0.104	0.134	0.113	0.084	0.071
	***[17.31]	***[13.44]	***[12.13]	***[25.04]	***[19.69]	***[13.01]
D_{DTBK}	0.127	0.094	0.135	0.091	0.072	0.065
	***[15.58]	***[12.55]	***[12.05]	***[21.22]	***[17.96]	***[12.61]
D_{SCGN}	0.153	0.135	0.177	0.112	0.097	0.092
	***[16.43]	***[14.41]	***[13.98]	***[21.93]	***[18.81]	***[15.98]
D_{TRBK}	0.119	0.079	0.121	0.084	0.062	0.062
	***[14.91]	***[10.47]	***[10.64]	***[20.62]	***[14.76]	***[12.42]
$\|\Delta\|$	-0.271	-0.277	-0.441	-0.195	-0.180	-0.187
	***[-19.64]	***[-12.63]	***[-9.00]	***[-20.51]	***[-14.07]	***[-9.41]
Δ^2	0.191	0.208	0.344	0.139	0.144	0.174
	***[14.08]	***[8.52]	***[6.56]	***[15.29]	***[9.45]	***[6.90]
D_{Comp}	-0.010	-0.002	0.005	-0.002	-0.002	0.004
	**[-2.47]	[-1.13]	*[1.82]	[-0.97]	[-1.10]	**[2.47]
D_{EUREX}	-0.002	0.004	-0.001	0.005	-0.005	-0.003
	[-0.41]	[0.89]	[-0.24]	[1.46]	**[-1.96]	*[-1.86]
VDAX_{t-1}	-0.029	0.000	-0.006	-0.043	-0.024	-0.010
	***[-6.13]	[-0.06]	[-0.63]	***[-12.28]	***[-6.11]	**[-2.08]
TTM	-0.009	0.025	0.019	-0.025	0.018	0.000
	[-1.11]	**[2.47]	***[3.65]	***[-3.72]	***[2.80]	[0.12]
LagWarVol	-0.004	-0.003	-0.002	-0.002	-0.001	-0.002
	***[-5.14]	***[-3.18]	***[-4.05]	***[-4.31]	**[-2.19]	***[-4.07]
$\mathbf{CDS}_{t-1}$						
$\|\Delta\| < .25$	**-0.269**	**-0.213**	**-0.567**	**-0.260**	**-0.207**	**-0.246**
	*****[-4.08]**	****[-2.41]**	*****[-4.03]**	*****[-5.38]**	*****[-2.70]**	*****[-3.48]**
$.25 < \|\Delta\| < .35$	**-0.305**	**-0.310**	**-0.347**	**-0.271**	**-0.178**	**-0.165**
	*****[-6.74]**	*****[-4.73]**	*****[-3.58]**	*****[-7.36]**	*****[-3.68]**	*****[-2.90]**
$.35 < \|\Delta\| < .50$	**-0.196**	**-0.191**	**-0.187**	**-0.195**	**-0.120**	**-0.163**
	*****[-5.66]**	*****[-3.55]**	****[-2.52]**	*****[-6.16]**	*****[-2.79]**	*****[-2.96]**
$.50 < \|\Delta\|$	**-0.125**	**-0.126**	**-0.060**	**-0.193**	**-0.128**	**-0.155**
	*****[-4.20]**	****[-2.31]**	**[-1.00]**	*****[-5.90]**	****[-2.02]**	***[-1.76]**
# Observations	35,724	24,185	20,403	47,162	33,974	26,703
R^2	0.473	0.449	0.307	0.513	0.49	0.271

Figure 3.1: Warrant Premium by Maturity and Type

The figures show the median premium, i.e. percentage overpricing, of warrants relative to matching EUREX options as a function of the warrant's (absolute) delta for the six largest issuers for which we have bid/ask quotes. The sample is split by type (calls vs. puts) and into three maturity ranges matching those used in Table 3.1: Short-term (1-3 months); medium-term (4-6 months); long-term (7 and more months). Premia are based on warrant mid-quotes matched with imputed EUREX option prices.

(a) *Short-term*

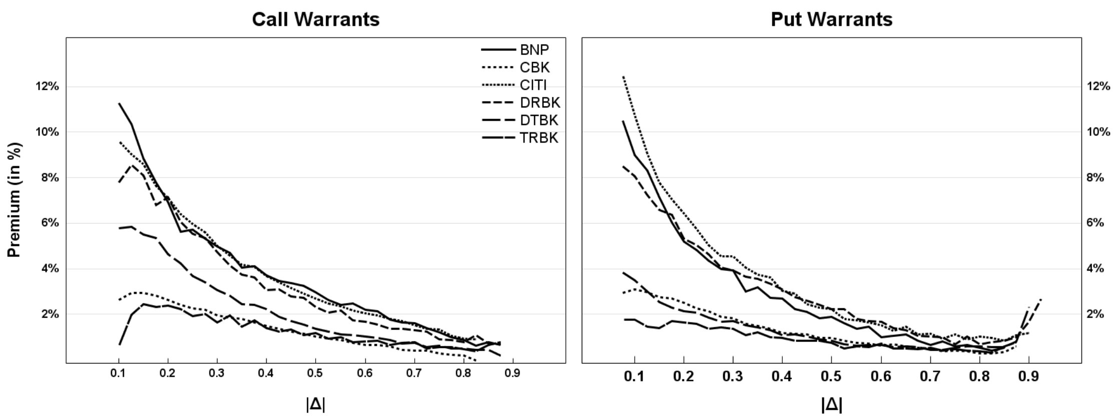

(b) *Medium-term*

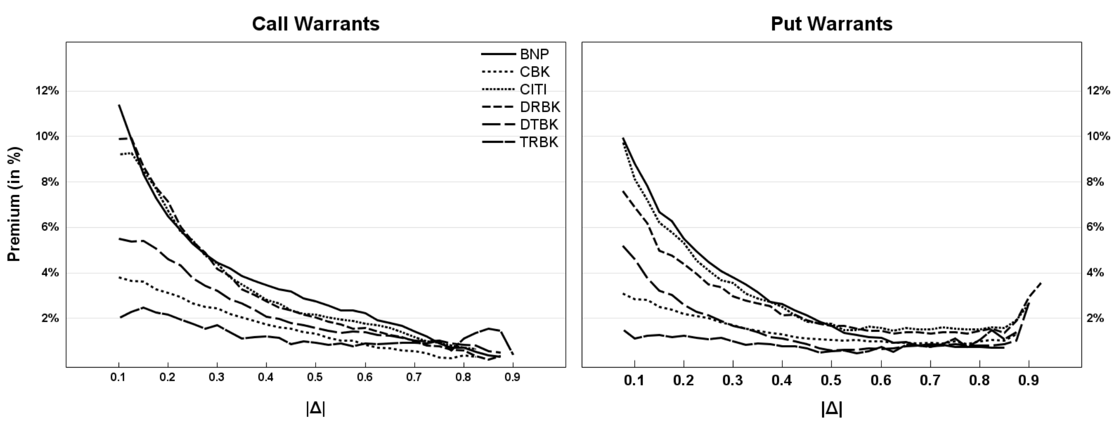

(c) *Long-term*

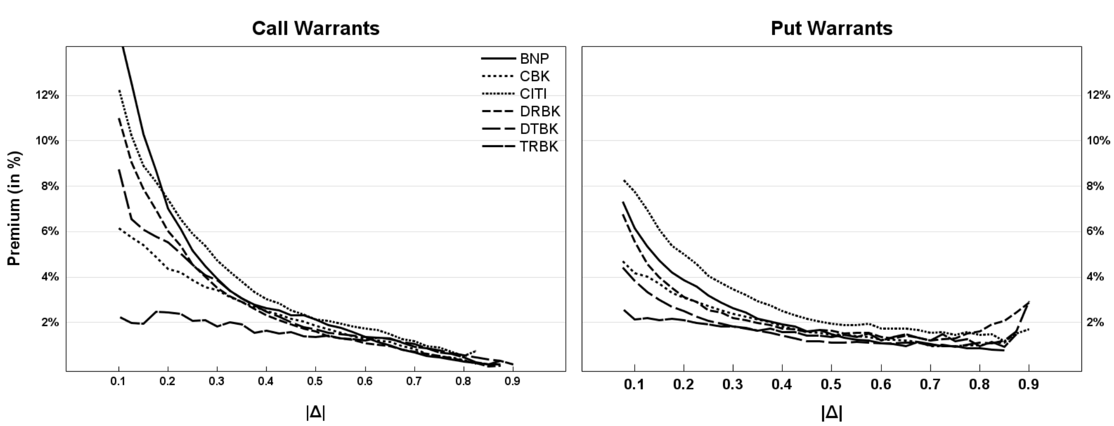

Figure 3.2: CDS Premium by Issuer

The figure shows the evolution of credit risk for the 7 issuers in the sample for which credit default swaps (CDS) are traded from January 2007 through September 2010. The sample period covered by transaction and quote data is June 2007 - May 2009. The bankruptcy of Lehman Brothers took place September 15th, 2008. CDS Premiums are end-of-day mid-quotes of issuer 1-year CDS contracts on unsecured debt (except for Citigroup, where it is for 1-year senior secured debt).

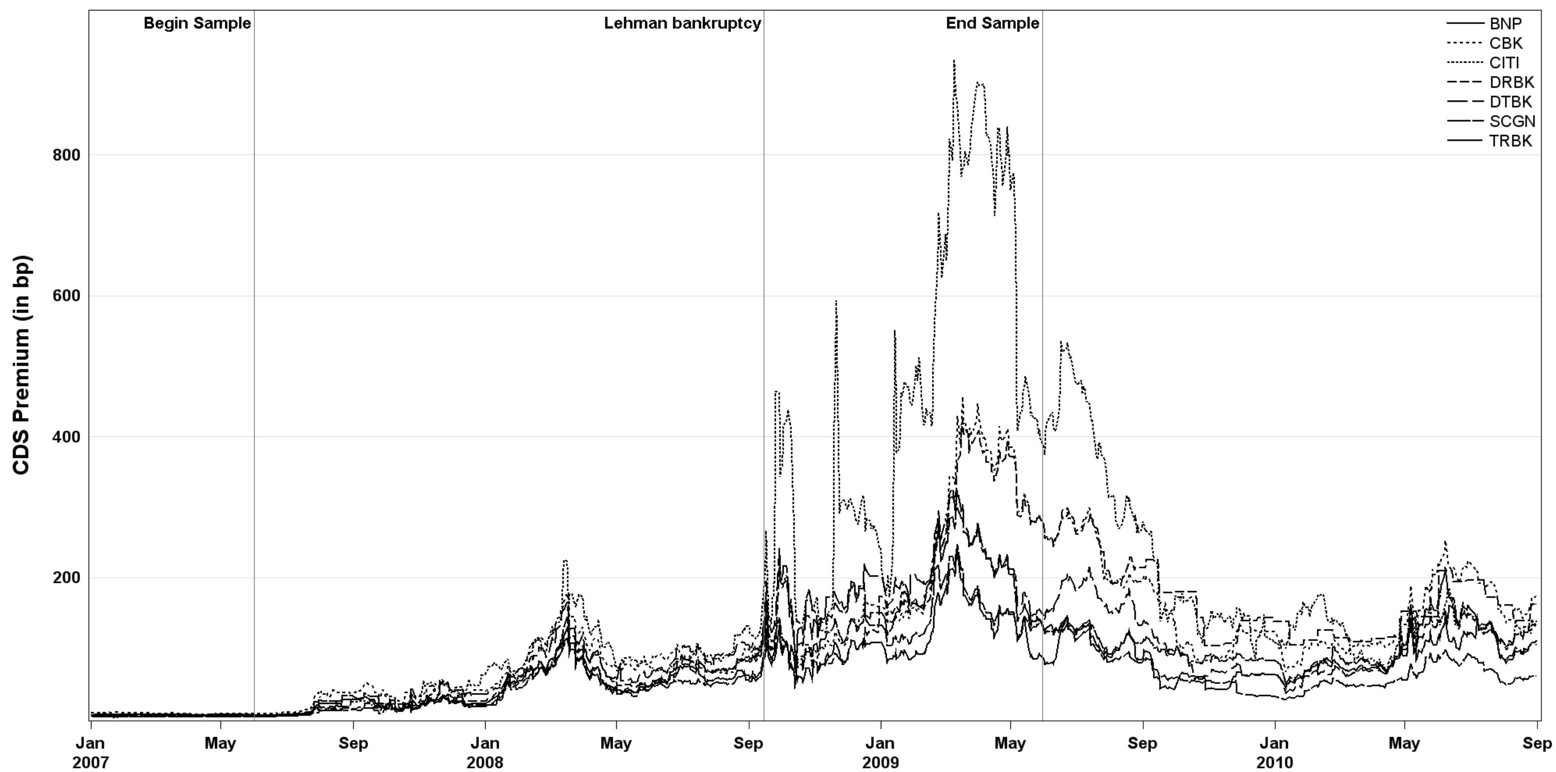

Chapter 4

Conclusion

In this thesis, I empirically investigate the effect of investor demand on option prices in two very different market environment. Chapter 2 expands the literature that considers the constraints and risks faced by a competitive market making sector and their effects on option prices and returns (e.g. Bollen and Whaley, 2004). This work is important because the finance literature at large, particular in macro finance struggles to explain observed puzzles in option prices and returns under no arbitrage assumptions (Bondarenko, 2003). I particularly focus on the risk introduced to options markets through the sensitivity of speculative, leveraged traders towards exogenous liquidity shocks that can result in forced sell-offs of options and temporarily push option prices far away from normal. Market makers and other liquidity providers form expectations about such an event occurring and are only willing to accomodate demand for options at prices that partially insure them against losses from those events. Besides the jump risk-based explanation of Garleanu et al. (2009), this constitutes another risk-based explanation of *permanent* price effects, which does not require jumps in the underlying asset.

Chapter 3 investigates what happens when two major tenets of no-arbitrage option pricing are taken away by market structure. In the market for (retail) structured products, issuing investment banks are monopolists in the provision of liquidity for their own products and further do not allow short-selling. This potentially enables them to sell these option-like securities at inflated prices. By using high-frequency transactions data and techniques previously not employed in this literature, my investigation details several ways in which this appears to happen.

Finding that previous research may have mis-estimated the term structure effect in the over-pricing of structured products, I hope that my work contributes to the prevention of such errors in future research. Finally, the effect of credit risk on warrant prices had not empirically been investigated. The finding that it took an event as disruptive as the Lehman Brothers bankruptcy to (partially) awaken retail investors to the inherent default risk in retail structured products is telling. It should be a wake-up call to regulators alike especially when considering that precisely Lehman certificates had previously been marketed as 'safe'. Thus, this work contributes to and extends an only recently emerging view (e.g. Henderson and Pearson, 2011) that these products expose retail investors to ill-understood risks.

4.1 Limitations

Both chapters have room for improvement, mainly with regards to the data employed. Given the datasets available at the time, some questions that arise naturally from the results presented cannot be answered. The COT data are publicly available only in a number of rather aggregate datasets that do not allow to back out certain pieces of information. Most notably, it is possible to back out the directional net option exposure of trader groups, but not whether this exposure is achieved through writing or buying different options, which in turn would allow us to infer the exposure of traders to jump risk and stochastic volatility risk. This is something that the individual data of Garleanu et al. (2009) allow. A second example is my inability to infer what fraction of, say, long option positions is hedged via futures by each trader. This would allow us to distinguish between hedging activities such as market making and liquidity provision and pure speculative positions within each trading group.

Likewise in the essay on structured products, the quality of the data is hampered by the lack of quotes on EUREX options rather than transaction data. Unfortunately, this has the potential to systematically affect some of the results. It may be the case that not only warrants but also EUREX options are affected by price and demand pressures. This is precisely the focus of the previous essay. For this reason I go to some length to establish that potential demand pressure effects between warrants and options are not highly correlated due to market segmentation and the differential use of instruments across moneyness and maturity.

4.2 Future Work

The measures employed in Chapter 2 (Kozhan et al., 2011) are still evolving and the data requirements are rather high for most markets. A natural extension, which I will tackle in the future, is to employ tick data on commodity futures options. This will allow much more precise estimates of both implied and realized skewness and, much in the way of the high-frequency analysis that I conduct in Chapter 3, this invites an investigation of short-term effect, i.e. the propagation of shocks from one particular option to others or from market-wide indicators such as the VIX on option prices.

As a side product of my work on skewness, 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 market making. Lastly, my findings suggest that limits to arbitrage in the futures market can spill over into the options market and impact prices and returns there. Conversely, the option market may contain information that affect futures returns.

Even 40 years after the ‘invention’ of option pricing, we still have not fully understood their behavior and in particular their deviation from theoretical, no-arbitrage values. Given the renewed focus on financial intermediaries and their effects on asset prices (Adrian and Shin, 2010), options offer a fruitful testing grounds due to their high dimensionality.

Bibliography

D. Abad and B. Nieto. Analysing bank-issued option pricing. *The European Journal of Finance*, 17(1):49–65, 2011.

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.

T. Adrian and H. S. Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3): 418–437, 2010.

T. Adrian, E. Etula, and J. J. Groen. Financial amplification of foreign exchange risk premia. *European Economic Review*, 55(3):354–370, 2011a.

T. Adrian, E. Etula, and T. Muir. Financial intermediaries and the cross-section of asset returns. Staff Report No. 464, Federal Reserve Bank of New York, 2011b.

K. Amin, J. D. Coval, and H. N. Seyhun. Index option prices and stock market momentum. *The Journal of Business*, 77(4):835–874, 2004.

G. Bakshi and N. Kapadia. Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, 16(2):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. Does risk-neutral skewness predict the cross-section of equity option portfolio returns? Working paper, Baruch College, CUNY, 2012.

T. G. Bali, N. Cakici, and R. F. Whitelaw. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2):427–446, 2011.

N. Barberis and M. Huang. Stocks as lotteries: The implications of probability weighting for security prices. *The American Economic Review*, 98(5):2066–2100, 2008.

S. M. Bartram and F. Fehle. Competition without fungibility: Evidence from alternative market structures for derivatives. *Journal of Banking & Finance*, 31(3):659–677, 2007.

S. M. Bartram, F. Fehle, and D. G. Shrider. Does adverse selection affect bid-ask spreads for options? *Journal of the Futures Markets*, 28(5):417–437, May 2008.

D. S. Bates. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1-2):181–238, 2000.

D. S. Bates. Empirical option pricing: A retrospection. *Journal of Econometrics*, 116(1-2):387–404, 2003.

R. Bauer, M. Cosemans, and P. Eichholtz. Option trading and individual investor performance. *Journal of Banking & Finance*, 33(4):731–746, 2009.

R. Baule, O. Entrop, and M. Wilkens. Credit risk and bank margins in structured financial products: Evidence from the german secondary market for discount certificates. *Journal of Futures Markets*, 28(4):376–397, 2008.

B. A. Benet, A. Giannetti, and S. Pissaris. Gains from structured product markets: The case of reverse-exchangeable securities (RES). *Journal of Banking & Finance*, 30(1):111–132, 2006.

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.

D. B. Bergstresser. The retail market for structured notes: Issuance patterns and performance, 1995-2008. Working Paper, Harvard Business School, 2009.

C. Bernard, P. P. Boyle, and W. Gornall. Locally-Capped Investment Products and the Retail Investor. *Journal of Derivatives*, 18(4):72–88, 2011.

A. E. Bernardo and I. Welch. Liquidity and financial market runs. *The Quarterly Journal of Economics*, 119(1):135–158, 2004.

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.

J. Bethel and A. Ferrell. Policy Issues Raised by Structured Products. *Brookings-Nomura Papers On Financial Services, Yasuki Fuchita, Robert E. Litan, Eds., Brookings Institution Press*, 2007.

Z. Bodie and V. I. Rosansky. Risk and return in commodity futures. *Financial Analysts Journal*, 36(3):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):711–753, 2004.

T. Bollerslev and V. Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66(6):21652211, December 2011.

T. Bollerslev, M. Gibson, and H. Zhou. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics*, 160(1): 235–245, 2011.

O. Bondarenko. Why are put options so expensive? Working paper, University of Illinois at Chicago, 2003.

B. Boyer and K. Vorkink. Stock options as lotteries. Working paper, Brigham Young University, 2011.

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.

C. Brown and K. J. Davis. Dividend protection at a price. *The Journal of Derivatives*, 12:62–68, 2004.

M. K. Brunnermeier and L. H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.

M. K. Brunnermeier, C. Gollier, and J. A. Parker. Optimal beliefs, asset prices, and the preference for skewed returns. *The American Economic Review*, 97(2):159–165, 2007.

M. K. Brunnermeier, S. Nagel, and L. H. Pedersen. Carry trades and currency crashes. Working Paper 14473, National Bureau of Economic Research, November 2008.

S. Burth, T. Kraus, and H. Wohlwend. The pricing of structured products in the Swiss market. *Journal of Derivatives*, 9:1074–1240, 2001.

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):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, University of Texas at Austin, 2011.

B. I. Carlin. Strategic price complexity in retail financial markets. *Journal of Financial Economics*, 91(3):278–287, 2009.

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.

H. W.-H. Chan and S. M. Pinder. The value of liquidity: Evidence from the derivatives market. *Pacific-Basin Finance Journal*, 8(3-4):483–503, 2000.

A. H. Chen and J. W. Kensinger. An analysis of market-index certificates of deposit. *Journal of Financial Services Research*, 4:93–110, 1990.

K. C. Chen and R. S. Sears. Pricing the SPIN. *Financial Management*, 19(2):36–47, 1990.

I.-H. Cheng, A. Kirilenko, and W. Xiong. Convective risk flows in commodity futures markets. Working Paper 17921, National Bureau of Economic Research, March 2012.

N. K. Chidambaran, C. S. Fernando, and P. A. Spindt. Credit enhancement through financial engineering: Freeport McMoRan's gold-denominated depositary shares. *Journal of Financial Economics*, 60(2-3):487–528, 2001.

J. S. Conrad, R. F. Dittmar, and E. Ghysels. Ex Ante Skewness and Expected Stock Returns. Working paper, 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. Dorn. Investors with too many options? Working paper, 2010.

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.

O. Entrop, H. Scholz, and M. Wilkens. The price-setting behavior of banks: An analysis of open-end leverage certificates on the german market. *Journal of Banking & Finance*, 33(5):874–882, 2009.

E. Etula. Broker-Dealer Risk Appetite and Commodity Returns. Staff Report No. 406, Federal Reserve Bank of New York, 2010.

V. Fardeau. Strategic arbitrage with entry. Working paper, London School of Economics, 2011.

K. A. Froot. The market for catastrophe risk: a clinical examination. *Journal of Financial Economics*, 60(2-3):529–571, 2001.

N. Garleanu and L. H. Pedersen. Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*, 24(6):1980–2022, 2011.

N. Garleanu, L. H. Pedersen, and A. M. Poteshman. Demand-based option pricing. *Review of Financial Studies*, 22(10):4259–4299, 2009.

M. Glaser and P. Schmitz. Privatinvestoren am Optionsscheinmarkt. *Die Zeitschrift für Bankrecht und Bankwirtschaft*, 19:214–230, 2007.

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.

A. Goyal and A. Saretto. Cross-section of option returns and volatility. *Journal of Financial Economics*, 94(2):310–326, 2009.

A. Grünbichler and H. Wohlwend. The valuation of structured products: Empirical findings for the Swiss market. *Financial Markets and Portfolio Management*, 19:361–380, 2005.

B. J. Henderson and N. D. Pearson. The dark side of financial innovation. *Journal of Financial Economics*, 100(2):227–247, 2011.

D. Hirshleifer. Risk, futures pricing, and the organization of production in commodity markets. *Journal of Political Economy*, 96(6):1206–1220, 1988.

J. t. Horst and C. Veld. An empirical analysis of the pricing of bank issued options versus options exchange options. *European Financial Management*, 14(2):288–314, 2008.

X. Hu, J. Pan, and J. Wang. Noise as information for illiquidity. Working Paper 16468, National Bureau of Economic Research, October 2010.

J. Hull and A. White. The impact of default risk on the prices of options and other derivative securities. *Journal of Banking & Finance*, 19:299–322, 1995.

R. A. Jarrow and M. O'Hara. Primes and Scores: An essay on market imperfections. *The Journal of Finance*, 44(5):1263–1287, 1989.

N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):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.

P. Klein. Pricing Black-Scholes options with correlated credit risk. *Journal of Banking & Finance*, 20(7):1211–1229, 1996.

R. Kozhan, A. Neuberger, and P. G. Schneider. The Skew Risk Premium in Index Option Prices. *SSRN eLibrary*, 2011.

A. Kumar. Who gambles in the stock market? *The Journal of Finance*, 64(4):1889–1933, 2009.

M. L. Lemmon and S. X. Ni. The effects of investor sentiment on speculative trading and prices of stock and index options. Working paper, 2011.

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.

I. Martin. Simple variance swaps. Working Paper 16884, National Bureau of Economic Research, March 2011.

A. Neuberger. Realized Skewness. *SSRN eLibrary*, 2011.

W. K. Newey and K. D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

W. K. Newey and K. D. West. Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies*, 61(4):631–653, 1994.

T. Odean. Are investors reluctant to realize their losses? *Journal of Finance*, 5:1775–1798, 1998.

M. Oehmke. Gradual arbitrage. Working paper, Columbia University, 2009.

T. Pratt. Equity derivatives vanish and explanations multiply. *The Investment Dealers' Digest*, 61: 15, 1995.

Z. Rehman and G. Vilkov. Risk-Neutral Skewness: Return Predictability and Its Sources. *SSRN eLibrary*, 2012.

R. J. Rogalski and J. K. Seward. Corporate issues of foreign currency exchange warrants : A case study of financial innovation and risk management. *Journal of Financial Economics*, 30(2): 347–366, 1991.

M. Rubinstein. Implied binomial trees. *The Journal of Finance*, 49(3):771–818, 1994.

H. Shefrin and M. Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, 40(3):777–790, 1985.

H. Shefrin and M. Statman. Behavioral aspects of the design and marketing of financial products. *Financial Management*, 22(2):123–134, 1993.

A. Shleifer and R. W. Vishny. The limits of arbitrage. *The Journal of Finance*, 52(1):35–55, 1997.

P. A. Stoimenov and S. Wilkens. Are structured products 'fairly' priced? an analysis of the German market for equity-linked instruments. *Journal of Banking & Finance*, 29(12):2971–2993, 2005.

Stuttgart Stock Exchange. *Implementation Rules relating to the Exchange Rules.* Stuttgart Stock Exchange, Retrieved August 1st, 2012, from https://www.boerse-stuttgart.de/media/dokumente/regelwerke/eng/Implementation Rules Derivative Securities.pdf, 2010a.

Stuttgart Stock Exchange. *Exchange Rules of the Stuttgart Stock Exchange.* Stuttgart Stock Exchange, Retrieved August 1, 2012, from https://www.boerse-stuttgart.de/media/dokumente/regelwerke/eng/Exchange Rules.pdf, 2010b.

M. Szymanowska, F. A. De Roon, T. E. Nijman, and R. Van den Goorbergh. An Anatomy of Commodity Futures Risk Premia. *SSRN eLibrary*, 2011.

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. B. Trolle and E. S. Schwartz. Variance risk premia in energy commodities. *Journal of Derivatives*, 17(3):15–32, 2010.

W. Wasserfallen and C. Schenk. Portfolio insurance for the small investor in Switzerland. *The Journal of Derivatives*, 3:37–43, 1996.

S. Wilkens and P. A. Stoimenov. The pricing of leverage products: An empirical investigation of the German market for 'long' and 'short' stock index certificates. *Journal of Banking & Finance*, 31(3):735–750, 2007.

S. Wilkens, C. Erner, and K. Röder. The pricing of structured products in Germany. *Journal of Derivatives*, 11(1):55–69, 2003.

Appendix A

Appendix to Chapter 2

A.1 Datasets

A.1.1 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.

A.1.1.1 Data Cleaning

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.

Further, because of the lack of open interest or volume data for futures options in the sample, some alternative filter for liquidity and information content of the price has to be implemented. To this end, option prices that are less than 8 times the minimum tick size of the option are automatically discarded because they are likely to be rather illiquid and prices tend to become too noisy relative to the information they contain.

A.1.1.2 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 $T \in 1w, 1m, 2m, \ldots, 12m$) 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

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

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:

$$F_{i,i+1} = \frac{100 - ED_i}{100} \tag{A.1}$$

$$DF_{i,i+1} = (1 + F_{i,i+1} \frac{d_{i+1} - d_i}{360})^{-1} \tag{A.2}$$

$$DF_{i+1} = DF_i \cdot DF_{i,i+1} \tag{A.3}$$

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

A.1.1.3 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 from this first step, 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. Because the method is non-parametric, model risk as it would be present when analyzing individual options or extracting distribution parameters is not an issue. Even the Black-Scholes implied volatility is purely used as a tool to transform American-type option prices into European-type ones, which the theory is based on.

In particular, BKM and studies that apply their measures (Dennis and Mayhew, 2002) recommend using at least 3 option observations per side. The high density of available strikes in my sample of commodity futures options makes this requirement rarely binding. Generally, even after all other filters are applied, both skewness measures are constructed on the basis of at least 10, but often up to 50 individual price observations per side. Only very close to expiration can it happen that the number of informative observations (i.e. after filtering for a certain multiple of the minimum tick size) decreases below the required 3 per side, at which point the skewness measure for that date is discarded.

A.1.1.4 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.1.2 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.

A.2 Theory of Option-Implied Measures

A.2.1 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:

$$\begin{aligned} H(S) = [H(S_0) - H_S(S_0)S_0] + H_S(S_0)S + &\int_0^{S_0} H_{SS}(K)\,(K-S)^+\,dK \\ + &\int_{S_0}^{\infty} H_{SS}(K)\,(S-K)^+\,dK \end{aligned} \tag{A.4}$$

The derivation can be found in the appendix of Carr and Madan (2001) 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.

$$V_0[H(S)] = \mathbb{E}_0^{\mathbb{Q}}\left[e^{-r\tau}H(S_T)\right] = \left[H(S_0) - H_S(S_0)S_0\right]e^{-r\tau} + H_S(S_0)S_0 \\ + \int_0^{S_0} H_{SS}(K)P_0(K)dK + \int_{S_0}^{\infty} H_{SS}(K)C_0(K)dK \tag{A.5}$$

Here, we used that $\mathbb{E}_0^{\mathbb{Q}}\left[e^{-r\tau}S_T\right] = S_0$. 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.

A.2.2 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

$$\mathrm{BKMSKEW}_{t,T} = \frac{\mathbb{E}_t^{\mathbb{Q}}\left[\left(r_{t,T} - \mathbb{E}_t^{\mathbb{Q}}[r_{t,T}]\right)^3\right]}{\mathbb{E}_t^{\mathbb{Q}}\left[\left(r_{t,T} - \mathbb{E}_t^{\mathbb{Q}}[r_{t,T}]\right)^2\right]^{3/2}} \tag{A.6}$$

$$= \frac{e^{r\tau}W_{t,T} - 3\mu_{t,T}e^{r\tau}V_{t,T} + 2\mu_{t,T}^3}{[e^{r\tau}V_{t,T} - \mu_{t,T}^2]^{3/2}} \tag{A.7}$$

where

$$\mu_{t,T} = \mathbb{E}_t^{\mathbb{Q}}\left[\log(\frac{S_T}{S_t})\right] \tag{A.8}$$

$$= e^{r\tau} - 1 - \frac{e^{r\tau}}{2}V_{t,T} - \frac{e^{r\tau}}{6}W_{t,T} - \frac{e^{r\tau}}{24}X_{t,T} \tag{A.9}$$

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.

A.2.3 Aggregation Property and Generalized Variance

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) echoes similar results in the literature (Jiang and Tian, 2005; Martin, 2011) which state 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 in the presence of jumps. To this end, Neuberger (2011) 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

$$\mathbb{E}_t^{\mathbb{Q}}\left[g(X_T - X_0)\right] = \mathbb{E}_t^{\mathbb{Q}}\left[\sum^{T} g(\delta X)\right] \tag{A.10}$$

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. The latter is defined as $v^f(s) = \mathbb{E}_t\left[f(S_T - S_0)\right]$ where for f it must hold that $\lim_{x\to 0} f(x)/x^2 = 1$.

A.2.4 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

$$\mathbb{E}_0^{\mathbb{Q}}\left[\left(\log\frac{S_T}{S_0}\right)^2\right] = \mathbb{E}_0^{\mathbb{Q}}\left[\sum^{T}\left(\log\frac{S_t}{S_{t-1}}\right)^2\right] \tag{A.11}$$

is not always exact[48]. As pointed out in Jiang and Tian (2005), this relationship only holds for fully continuous processes without jumps. Neuberger (2011) and Kozhan et al. (2011) propose two alternative functional forms, which have both the property of generalized variance and the aggregation property: $g^L(s) = 2(e^s - 1 - s)$ and $g^E(s) = 2(s \cdot e^s - e^s + 1)$. For instance, for

[48] An alternative approach by Martin (2011) achieves consistency by proposing an alternative definition of both legs of the variance swap.

$g^L(s) = 2(e^s - 1 - s)$, it it is easy to see that, even in the presence of jumps, it holds that

$$\mathbb{E}_0^{\mathbb{Q}}\left[g^L\left(\log\left(\frac{S_T}{S_0}\right)\right)\right] = 2\mathbb{E}_0^{\mathbb{Q}}\left[\frac{S_T}{S_0} - 1 - \log\left(\frac{S_T}{S_0}\right)\right] \tag{A.12}$$

$$= 2\mathbb{E}_0^{\mathbb{Q}}\left[\sum^{T}\left[\frac{S_t}{S_{t-1}} - 1 - \log\left(\frac{S_t}{S_{t-1}}\right)\right]\right] \tag{A.13}$$

Rewrite the right-hand side of Equation A.12 as $v_{0,T}^L := 2\left(\log S_0 - \mathbb{E}_0^{\mathbb{Q}}[\log S_T]\right)$ and call it the implied variance of a security that pays $\log S_T$ at time T, or log variance. Using the spanning approach, the implied log variance 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) and is used in this paper as well. 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. In a similar fashion, following Neuberger (2011), one can write

$$\mathbb{E}_0^{\mathbb{Q}}\left[g^E\left(\log\left(\frac{S_T}{S_0}\right)\right)\right] = 2\mathbb{E}_0^{\mathbb{Q}}\left[\frac{S_T}{S_0}\log\left(\frac{S_T}{S_0}\right) - \frac{S_T}{S_0} + 1\right] \tag{A.14}$$

$$= 2\mathbb{E}_0^{\mathbb{Q}}\left[\frac{S_T}{S_0}\log S_T - \frac{S_T}{S_0}\log S_0\right] \tag{A.15}$$

$$= 2\left[\frac{\mathbb{E}_0^{\mathbb{Q}}[S_T \log S_T]}{S_0} - \log S_0\right] \tag{A.16}$$

and call the last expression $v_{0,T}^E$, the implied variance of the entropy contract paying $S_T \log S_T$. The replicating weights in this case are $H_{SS}^E(K) = \frac{2}{S_0 K}$. The log and entropy variances are used to define the skewness swap below.

A.2.5 Skewness Swap

Similar to the derivation of a consistent variance swap, Neuberger (2011) and Kozhan et al. (2011) also propose 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

$$v_{0,T}^L = -2\mathbb{E}_0^{\mathbb{Q}}\left[\log\left(\frac{S_T}{S_0}\right)\right] \tag{A.17}$$

and similarly, one can show that

$$v_{0,T}^E = 2\mathbb{E}_0^{\mathbb{Q}}\left[\left(\frac{S_T}{S_0}\right)\log\left(\frac{S_T}{S_0}\right)\right] \tag{A.18}$$

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

$$\mathrm{MFIS}_{0,T} = \mathbb{E}_t^{\mathbb{Q}}\left[g^Q\left(s_T - s_0, v^E(s_T - s_0)\right)\right] \tag{A.19}$$

$$= 6\mathbb{E}_t^{\mathbb{Q}}\left[\left(\frac{S_T}{S_0} + 1\right)\log\frac{S_T}{S_0}\right] \tag{A.20}$$

$$= 3(v_{0,T}^E - v_{0,T}^L) \tag{A.21}$$

The option replicating weights are $H_{SS}^Q(K) = \frac{2(K-S_0)}{S_0K^2}$ using the spanning approach. Realized skewness over the period $[0,T]$ can be computed exactly as

$$\sum^T g^Q(\delta s, \delta v^E) = \sum^T\left[3\delta v^E e^{\delta s} + 6\big(\delta s(e^{\delta s} + 1) - 2(e^{\delta s} - 1)\big)\right] \tag{A.22}$$

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

$$6\left(\log(1 + r_{\delta t})(r_{\delta t} + 2) - 2r_{\delta t}\right) = r_{\delta t}^3 + O(r_{\delta t}^4) \tag{A.23}$$

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. In this case, and as it turns out in practice, the first term is far more important. Using slightly different notation, Kozhan et al. (2011) write implied skewness under the assumption of continuous re-balancing as

$$\mathrm{MFIS}_{0,T} = 3\mathbb{E}^{\mathbb{Q}}\left[\int_0^T dv_{t,T}^E\left(\frac{dS_t}{S_t}\right)\right], \tag{A.24}$$

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

$$\mathrm{MFIS}_{0,T} = 3(v_{0,T}^E - v_{0,T}^L) = \mathbb{E}^{\mathbb{Q}}\left[r_{0,T}^3 + O(r_{0,T}^4)\right] \tag{A.25}$$

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.