Liquidity Creation as Volatility Risk

Itamar Drechsler, Alan Moreira, Alexi Savov*

January 2018

Abstract

We show, both theoretically and empirically, that liquidity creation induces negative exposure to volatility risk. Intuitively, liquidity creation involves taking positions that can be exploited by privately informed investors. These investors' ability to predict future price changes makes their payoff resemble a straddle (a combination of a call and a put). By taking the other side, liquidity providers are implicitly short a straddle, suffering losses when volatility spikes. Empirically, we show that short-term reversal strategies, which mimic liquidity creation by buying stocks that go down and selling stocks that go up, have a large negative exposure to volatility shocks. This exposure, together with the large premium investors demand for bearing volatility risk, explains why liquidity creation earns a premium, why this premium is strongly increasing in volatility, and why times of high volatility like the 2008 financial crisis trigger a contraction in liquidity. Taken together, these results provide a new, asset-pricing view of the risks and rewards to financial intermediation.

JEL Codes:

Keywords: Liquidity, volatility, reversals, VIX, variance premium

*Drechsler and Savov are at New York University and NBER, idrechsl@stern.nyu.edu and asavov@stern.nyu.edu. Moreira is at University of Rochester, alan.moreira@simon.rochester.edu. We thank Liz Casano for excellent research assistance. Suggestions and comments are welcome.

We show, both theoretically and empirically, that liquidity creation—making assets cheaper to trade than they otherwise would be—induces exposure to volatility risk. Given the very large premium investors pay to avoid volatility risk (Carr and Wu, 2008), this explains why liquidity creation earns a premium (Krishnamurthy and Vissing-Jorgensen, 2015), why this premium is strongly increasing in volatility (Nagel, 2012), and why times of high volatility like the 2008 financial crisis trigger a contraction in liquidity (Brunnermeier, 2009). Taken together, these results provide a new, asset-pricing view of the risks and rewards of financial intermediation.

Why does liquidity creation induce exposure to volatility risk? To create liquidity for some investors in an asset, a liquidity provider takes positions that can be exploited by other, privately informed investors (e.g. Kyle, 1985). These investors buy the asset if they think it will rise in value and sell it if they think it will fall. Their ex post payoff therefore resembles a straddle (a combination of a call and a put option). Like any straddle, this payoff is high if volatility rises and low if it falls. By taking the other side, the liquidity provider is implicitly short the straddle, earning a low payoff if volatility rises and a high one if it falls. In other words, the liquidity provider is exposed to volatility risk.

The relation between liquidity creation and volatility risk is fundamental; it arises directly from the presence of asymmetric information (a staple of the literature since Akerlof, 1970). As a result, it applies widely across a variety of market structures. For instance, one way financial institutions create liquidity is by issuing relatively safe securities against risky assets (Gorton and Pennacchi, 1990). In doing so, they are betting against private information possessed by those who originate the assets or in some other way take a position against them (e.g. through derivatives). Consequently, when volatility spikes and this private information becomes more valuable, financial institutions suffer losses, as they did during the 2008 financial crisis.

Financial institutions and other investors also create liquidity by trading in secondary markets such as those for stocks and bonds. We present a model to formalize how this type of liquidity creation induces volatility risk and how this risk drives the liquidity premium. We also use the model to motivate our empirical analysis.

The model builds on the framework of Kyle (1985), but with the added ingredient of

stochastic volatility. A liquidity provider observes order flow, which captures the liquidity demand of informed and uninformed investors alike. She sets prices to break even on average, earning profits against the uninformed to offset losses against the informed. To do so, she has to predict the value of the informed investors' private signal. This value depends on the remaining volatility in the asset's payoff. As a result, prices conditional on order flow depend on expected volatility going forward. When expected volatility is high, assets with negative order flow have lower prices and assets with positive order flow have higher prices. When expected volatility is low, prices depend less on order flow. If, however, volatility rises in the following period, the value of the informed investors' information also rises. This earns them a higher expected profit than was priced in by the liquidity provider who therefore incurs an expected loss. Thus, the liquidity provider has a negative exposure to volatility shocks. The magnitude of this exposure depends on the expected size of these shocks, i.e. on the volatility of volatility.

The model has many assets and liquidity providers are free to trade all of them. Yet they cannot diversify away their volatility risk because volatility follows a factor structure across assets. As Herskovic et al. (2016) show, this assumption is strongly supported in the data. Also strongly supported is that the factors that drive asset-level volatility are highly correlated with aggregate volatility (e.g. Campbell et al., 2001; Bekaert, Hodrick and Zhang, 2012). Hence, our model predicts that liquidity providers have negative exposure to aggregate as well as asset-level volatility. The asset pricing literature shows that aggregate volatility in particular carries a very large negative risk premium (Carr and Wu, 2008; Bollerslev, Tauchen and Zhou, 2009).[1] Therefore, our model predicts that liquidity providers should earn a positive premium for their negative exposure to volatility. This liquidity premium arises purely as compensation for bearing volatility risk.

We test the predictions of our model using U.S. stock return data from 2001 to 2016 (covering the period after "decimalization," when liquidity provision became competitive; see Bessembinder, 2003). Each day, we sort stocks into deciles based on their return (normalized by its rolling standard deviation) and quintiles based on their size (small

[1]Given the high correlation between asset-level and aggregate volatility, it is possible that the priced risk factor is actually asset-level volatility and not aggregate volatility, or that the premium is for both, but this possibility is rarely explored in the literature. Herskovic et al. (2016) is one exception.

stocks are known to be much less liquid). Within each size quintile, we construct long-short portfolios that buy stocks in the low return deciles and sell stocks in the high return deciles. These are known as short-term reversal portfolios in the literature (e.g. Lehmann, 1990). In our model, a large return reflects high order flow and hence high liquidity demand.[2] The reversal portfolio therefore mimics the position of the liquidity provider, hence, as in Nagel (2012), we can use it to analyze the returns to liquidity creation.

Consistent with the model, and with the prior literature, our reversal portfolios earn substantial returns that cannot be explained by exposure to market risk. Among large stocks, which account for the bulk of the market by value, the reversal strategy across the lowest and highest return deciles has an average return of 27 bps over a five day holding period, or about 13.5% per year. The annual Sharpe ratio is 0.6.

Figure 1 plots the return of the large-stock reversal strategy averaged over a 60-day forward-looking window against the level of VIX, a risk-neutral measure of the expected volatility of the S&P 500 over the next 30 days. The figure shows that the reversal return is strongly positively related to VIX (the raw correlation is 46%). In a regression, we find that a one-point higher VIX leads to a 5.37 bps higher reversal return over the next five days, which is large relative to the average return of the strategy. The R^2 of this regression is 2.18%, which is very high for daily data. These findings confirm the main result of Nagel (2012) that VIX predicts reversal returns.[3] They are also a prediction of our model. A high level of VIX is associated not only with high expected volatility but also with high volatility of volatility (and high aversion to volatility risk). In our model this makes liquidity creation riskier and raises the price of liquidity.

The bottom panel of Figure 1 tests this mechanism by plotting a measure of the volatility risk of the reversal strategy. We compute it by running 60-day rolling window regressions of the five-day large-stock reversal return on the daily VIX changes during the holding period. The figure plots the annualized standard deviation of the fitted value from this regression, which captures the systematic volatility of the reversal strategy due to

[2]In practice, returns also reflect the release of public news, which makes them a noisy proxy. To partly reduce this noise, we follow Collin-Dufresne and Daniel (2014) and remove earnings announcement days.

[3]Our methodology differs somewhat from Nagel (2012) in that we focus on large stocks and opt for decile sorting instead of weighting by returns.

VIX changes, i.e. its volatility risk. The figure shows that the volatility risk of the reversal strategy is substantial and that it covaries strongly with the level of VIX (the raw correlation is 58%). This confirms the prediction that when VIX is high the reversal strategy is exposed to more volatility risk, which is consistent with its higher premium.

To formally test if the reversal strategy is exposed to volatility risk, we regress its return on changes in VIX. We find a strong negative beta coefficient: the five-day large-stock reversal return drops by 64 bps if VIX rises by an average of one point per day over the holding period, which is equivalent to 1.3 standard deviations. The beta is large relative to the average return of the strategy (27 bps). It rises to 71 bps if we control for the market return, and the results are similar for small stocks. Thus, the reversal strategies have substantial negative exposure to volatility risk, a central prediction of the model.

The next step in our analysis is to test whether this volatility risk exposure can account for the average returns of the reversal strategies. This prediction differentiates our theory from others, which are based on the inventory risk of liquidity providers (Stoll, 1978; Grossman and Miller, 1988; Gromb and Vayanos, 2002; Duffie, 2010). We note that inventory risk is in principle fully diversifiable, hence it should not command any premium, let alone one as large as the one we find in this recent sample. This tension is present in all models of the liquidity premium based on inventory risk. In contrast, our model shows that there is a tight relationship between liquidity creation and volatility risk, which is nondiversifiable, and thus ties together their premiums.

Our first pricing test runs Fama-MacBeth regressions of the five-day reversal portfolio returns on the daily market returns and VIX changes during the holding period. We find that exposure to VIX changes is strongly priced and substantially reduces the pricing errors of the portfolios. The pricing error of the largest quintile reversal strategy drops from 25 bps to -7 bps. The pricing errors of the next two quintiles are similarly driven down (20 bps to 4 bps and 33 bps to 6 bps, respectively). Only the pricing errors of the two smallest quintiles remain significant (these stocks account for just 0.4% of market value). Hence, outside these very small stocks, our results show that the returns to liquidity provision are well priced by volatility risk exposures.

Our final test asks if volatility risk can price the reversal strategies with the same price

of risk that prevails in other markets. If so, it would suggest that the price of liquidity creation reflects broad economic risks rather than institutional frictions.

The natural place to measure the price of volatility risk is in option markets. The VIX index itself is based on the price of a basket of options whose payoff replicates the realized variance of the S&P 500 over a 30-day period. Since this basket changes from day to day as options near expiration, the change in VIX is not a valid return from which we can estimate a price of risk. However, it is straight-forward to compute such a return by simply holding the basket fixed from one day to the next. We do so using data on S&P 500 options and following the methodology of the Chicago Board Options Exchange (CBOE). The resulting VIX return has a sample average of -1.54% per day.

Whereas VIX has a 30-day horizon, liquidity provision in equity markets typically takes place on a shorter time scale. This is why we follow the literature and focus on reversal returns out to five trading days. While it is not feasible to target such a short horizon with index option data,[4] we are able to move in that direction by calculating the return associated with VIXN, the "near-term VIX", which is one of the two cross sections of options used by CBOE in the computation of VIX (the average horizon of VIXN is 22 days). We find that the VIXN return is -2.01% per day, which is larger even than the VIX return.

We calculate a price of risk associated with exposure to VIX changes by dividing the average VIX return by its loading on VIX changes. This gives a price of -22 bps, i.e. an asset that rises by one percent when VIX rises by one point is predicted to have an average return of -22 bps per day. We do the same for the VIXN return with respect to VIXN changes and obtain a price of risk of -33 bps.

Using these prices of risk and the betas of the reversal portfolios, we calculate predicted returns and pricing errors.[5] We find that these restricted pricing tests perform similarly to the unrestricted Fama-MacBeth regressions, especially when we use the near-term VIXN. Specifically, when we use VIX changes as the priced factor and the VIX re-

[4]Options often behave erratically near expiration, hence the CBOE discards options with less than a week to expiration from the computation of VIX.

[5]This is analogous to including the VIX and VIXN returns as additional test assets and imposing the restriction that they be priced perfectly. This approach is also followed by Constantinides, Jackwerth and Savov (2013).

turn to obtain its price, the pricing error of the large-stock five-day reversal strategy drops from 25 bps to 11 bps and remains marginally significant. On the other hand, when we use VIXN changes as the priced factor and the VIXN return to obtain its price, this pricing error drops to just 1 bp. Thus, near-term volatility risk prices the large-stock reversal strategy almost perfectly. The pricing errors of the second and third largest quintiles also become insignificant. Only the two smallest quintiles retain significant pricing errors, similar to the Fama-MacBeth regressions. Overall, these results indicate that the returns to liquidity creation reflect compensation for the risks of betting against volatility.

The rest of this paper is organized as follows. Section 1 reviews the literature, Section 2 presents the model, Section 3 discusses the data, Section 4 analyzes the empirical results, and Section 5 concludes.

1 Related literature

A large literature in asset pricing studies how liquidity is priced in financial markets. Amihud and Mendelson (1986) and Brennan and Subrahmanyam (1996) show that illiquid stocks have higher expected returns. Chordia, Roll and Subrahmanyam (2000) and Hasbrouck and Seppi (2001) show that different measures of illiquidity comove strongly. Amihud (2002) and Pástor and Stambaugh (2003) show that aggregate illiquidity is priced both in the time series and cross section of stocks. Acharya and Pedersen (2005) provide an equilibrium model that captures these facts. This literature has also shown that liquidity is decreasing in volatility. Chordia, Sarkar and Subrahmanyam (2004) show that increases in volatility are associated with a reduction in liquidity. Hameed, Kang and Viswanathan (2010) show that liquidity provision strategies earn high returns following stock market downturns. Nagel (2012) shows that expected returns to liquidity provision in stocks are increasing in the level of VIX.

A second large literature in asset pricing studies how volatility is priced in financial markets. Carr and Wu (2008) show that investors pay a large premium (the variance risk premium) to hedge aggregate volatility risk. Todorov (2009) and Bollerslev and Todorov (2011) extend these results to jump risk. Bollerslev, Tauchen and Zhou (2009) and Drech-

sler and Yaron (2010) show that the variance risk premium strongly predicts aggregate stock returns. Manela and Moreira (2017) extend these results by over a century. Bao, Pan and Wang (2011) and Longstaff et al. (2011) find similar results for bonds. Drechsler and Yaron (2010), Drechsler (2013), and Dew-Becker et al. (2017) provide theories that explain the behavior of the variance risk premium as driven by macroeconomic risk.

The contribution of our paper is to integrate the literatures on the pricing of liquidity and volatility. First, we show theoretically that liquidity provision induces exposure to volatility risk. Second, we confirm this prediction empirically and show that volatility risk can explain the observed returns to liquidity provision strategies.

Our model is rooted in the theoretical literature on how asymmetric information impacts asset prices and liquidity (e.g. Hellwig, 1980; Grossman and Stiglitz, 1980; Diamond and Verrecchia, 1981; Kyle, 1985; Glosten and Milgrom, 1985). Specifically, we introduce stochastic volatility in the framework of Kyle (1985).[6] We show that informed traders are effectively long a straddle because they can trade in the direction of subsequent price changes. Liquidity providers, who take the other side, are therefore short a straddle. As a result, they suffer losses when volatility comes in higher than expected. This exposure cannot be diversified away because volatility is highly correlated across assets and with market volatility. Our model shows how this leads to the emergence of a liquidity premium as compensation for volatility risk.

A more recent literature focuses on the role of financial frictions in generating fluctuations in liquidity (e.g. Eisfeldt, 2004; He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014). Gromb and Vayanos (2002), Brunnermeier and Pedersen (2008), and Adrian and Shin (2010) show how, when financial institutions are risk averse or subject to a Value-at-Risk (VaR) constraint, a rise in volatility leads to a contraction in the supply of liquidity. Moreira and Savov (2017) show how such a liquidity contraction impacts asset prices and the macroeconomy. Our paper shows how these dynamics can arise as a result of fluctuations in the variance risk premium. Of course, it is also possible that the

[6]Collin-Dufresne and Fos (2016) solve a dynamic version of this framework with stochastic noise trader volatility. While this feature does not create volatility risk exposure for the liquidity provider, our frameworks share the presence of long-lived private information, which contrasts with earlier papers like Admati and Pfleiderer (1988) and Foster and Viswanathan (1990).

variance risk premium itself is a reflection of the scarcity of liquidity, particularly at the short end. We return to this intriguing possibility at the end of the paper.

2 Model

We present a model in the tradition of Kyle (1985) that shows why liquidity provision is fundamentally exposed to aggregate volatility risk and thus inherits its substantial, negative price of risk. The model makes several further predictions and guides the construction and implementation of our empirical tests.

There are three time periods, 0, 1, and an in-between period $t \in (0,1)$, and N assets that are traded at time 0. There are three types of market participants: insiders, liquidity demanders, and liquidity providers. For each asset there is a single insider, a unit mass of liquidity demanders, and a competitive fringe of liquidity providers.

Each asset i has a single payoff $p_{i,1}$ that is realized at time 1 and is given by:

$$p_{i,1} \;=\; \overline{v}_i + \sigma_{i,1}\, v_i, \tag{1}$$

where $\bar{v}_i$ is a constant and $v_i \sim N(0,1)$ is an idiosyncratic shock that is uncorrelated across assets and is scaled by the volatility $\sigma_{i,1}$. Thus, $\sigma_{i,1}$ represents asset i's idiosyncratic volatility. The value of $\overline{v}_i$ is known by all investors prior to time 0 and is the price of the asset before any trading takes place. At time 0 insiders learn the value of v_i, whereas its value remains unknown to the other market participants. Trading then takes place.

All market participants share the same information and expectations about the random variable $\sigma_{i,1}$, which is uncertain until time 1. The realization of $\sigma_{i,1}$ is driven by both a market (i.e., systematic) volatility factor σ_m that is common to all N assets' idiosyncratic volatilities, and an asset-specific component ε_{σ_i}:

$$\sigma_{i,1} \;=\; k_i \sigma_{m,1} + \varepsilon_{\sigma_i}. \tag{2}$$

The parameter k_i gives the loading of asset i's idiosyncratic volatility on the market volatility factor $\sigma_{m,1}$. We assume $k_i > 0$ to capture the observed strong comovement

of idiosyncratic and market volatilities.

At the intermediate time $t \in (0,1)$, news arrives about volatilities that changes market participants' expectations of the value of $\sigma_{i,1}$. To focus on the impact of this volatility news, we assume that only liquidity providers trade at time t, so that changes in prices at this time are due only to updates in volatility expectations.

The single insider in each asset maximizes the value of profits under the economy's pricing kernel:[7]

$$y_i^* \;=\; \arg\max_{y_i} \; E_0^Q \left[y_i \left(p_{i,1} - p_{i,0} \right) \mid v_i \right], \tag{3}$$

where E_0^Q denotes expectation at time-0 under the risk-neutral probabilities implied by the pricing kernel, y_i is the insider's share position in asset i, and $p_{i,0}$ is the time-0 price of the asset. Liquidity traders demand for shares of asset i is given by $z_i \sim N\left(0, \sigma_{z_i}^2\right)$, which we assume is uncorrelated across assets. Finally, there is the unit mass of liquidity providers. We assume liquidity providers are integrated with the rest of the market (no segmentation), and thus also price payouts according to the economy's pricing kernel.

As in Kyle (1985), liquidity providers cannot distinguish between the trades of insiders and liquidity traders and only observe their net order flow, $X_i = z_i + y_i^*$. Since liquidity providers are competitive, they set the time-t price $p_{i,t}$ so that they break even in expectation:

$$p_{i,t} \;=\; E_t^Q \left[p_{i,1} \mid X_i \right]. \tag{4}$$

As in Kyle (1985), we conjecture (and later verify) that the insider's demand is given by

$$y_i^* \;=\; \phi \left(E_0^Q \left[p_{i,1} \mid v_i \right] - \bar{p}_i \right) \;=\; \phi \left(\bar{v}_i + E_0^Q \left[\sigma_{i,1} \right] v_i - \bar{p}_i \right), \tag{5}$$

where ϕ and $\overline{p}_i$ are determined in equilibrium. By Bayes' rule, the price at time $t \in [0,1)$

[7] The assumption of a single insider is inessential. We could instead model a mass of competitive insiders who are risk-averse. The essential requirement is that some consideration limits the intensity of insiders' trading, so that their trading does not fully reveal their private information.

is

$$p_{i,t} \;=\; \overline{v}_i + E_t^Q\left[\sigma_{i,1}\right]\; E_t^Q\left[v_i \mid X_i\right] \;=\; \overline{v}_i + E_t^Q\left[\sigma_{i,1}\right]\lambda X_i, \tag{6}$$

$$\text{where} \quad \lambda \;=\; \frac{\operatorname{cov}_0\left(v_i, X_i\right)}{\operatorname{var}_0(X_i)} \;=\; \frac{\phi E_0^Q\left[\sigma_{i,1}\right]}{\phi^2(E_0^Q\left[\sigma_{i,1}\right])^2 + \sigma_{z_i}^2}. \tag{7}$$

The solution for λ is based only on time-0 information because the insider's demand y_i^* is determined at time 0 and there is no additional information about v_i in the time-t volatility news. Substituting the price schedule (6) into the insider's problem and simplifying gives:

$$\begin{aligned} y_i^* \;&=\; \arg\max_{y_i}\; E_0^Q\left[y_i\left(E_0^Q[\sigma_{i,1}]v_i - \lambda E_0^Q[\sigma_{i,1}](z_i + y_i)\right) \;\middle|\; v_i\right] && (8)\\ &=\; \arg\max_{y_i}\; y_i E_0^Q[\sigma_{i,1}]v_{i,1} - y_i^2\lambda E_0^Q[\sigma_{i,1}] && (9)\\ &=\; \frac{1}{2\lambda}v_i. && (10)\end{aligned}$$

Setting this expression equal to (5) and solving the resulting equation jointly with (7) gives the values of λ, ϕ, and $\overline{p}_i$ and confirms the conjecture (5) for y_i^*. We obtain the following.

Proposition 1. *ϕ, $\overline{p}_i$, and λ are given by*

$$\phi \;=\; \frac{\sigma_{z_i}}{E_0^Q[\sigma_{i,1}]}, \qquad \overline{p}_i \;=\; \overline{v}_i, \qquad \lambda \;=\; \frac{1}{2\sigma_{z_i}}, \tag{11}$$

and asset i's price at time $0 \le t < 1$ is

$$p_{i,t} \;=\; \overline{v}_i + \frac{X_i}{2\sigma_{z_i}}E_t^Q[\sigma_{i,1}]. \tag{12}$$

Thus, the time-0 price is pushed in the direction of the net order flow X_i. This is intuitive since liquidity providers absorb this order flow by taking the opposite position, $-X_i$. The sensitivity of the price to the order flow imbalance is inversely related to the expected amount of uninformed liquidity trading, σ_{z_i}, and proportional to expected volatility. A higher expected volatility makes insiders' information more powerful and thus increases adverse selection, making the price more sensitive to imbalances in order flow.

Unfortunately, neither the sign or magnitude of X_i are directly observable in the data.

While volume is observable, it is unsigned and represents only the gross quantity of trading, whereas X_i is a net quantity. The relationship between net and gross can be complicated and unstable across different assets and markets, depending on the number of individual liquidity traders and the amount of netting that occurs within this group. Fortunately, equation (12) shows that the change in the asset's price at time 0 contains information about the sign and magnitude of liquidity providers' exposure to the asset's volatility. Let $\Delta p_{i,0} = (p_{i,0} - \overline{v}_i)$ denote asset i's time-0 price change. We have the following.

Lemma 1. *The exposure of liquidity providers to asset i's volatility, $-X_i/\sigma_{z_i}$, is proportional to, and has the opposite sign of, its time-0 price change:*

$$-\frac{X_i}{\sigma_{z_i}} \;=\; -\Delta p_{i,0}\frac{2}{E_0^Q[\sigma_{i,1}]} \tag{13}$$

Thus, liquidity providers hold a portfolio of reversals: they take larger long (short) exposures in assets that have larger time-0 price decreases (increases).

Together, equations (12) and (13) show that assets' time-0 price changes give their exposures to volatility news. As the following proposition shows, liquidity providers are *negatively* exposed to increases in volatility on *all* of their positions, both long and short. When volatilities increase, both legs of the reversals go against the liquidity providers– long positions decline while short positions advance. Let $\Delta p_{i,t} = p_{i,t} - p_{i,0}$ denote the time-t price change of asset i. We have the following:

Proposition 2. *The change in the price of asset i at time t is given by*

$$\Delta p_{i,t} \;=\; \frac{X_i}{2\sigma_{z_i}}(E_t^Q[\sigma_{i,1}] - E_0^Q[\sigma_{i,1}]). \tag{14}$$

and hence the beta of liquidity providers' holdings of asset i ($-X_i\Delta p_i$) to shocks to market volatility ($E_t^Q[\sigma_{m,1}] - E_0^Q[\sigma_{m,1}]$) is

$$\beta_{i,\sigma_m} \;=\; \frac{-X_i\Delta p_{i,t}}{E_t^Q[\sigma_{i,1}] - E_0^Q[\sigma_{i,1}]}k_{i,m} \;=\; -\frac{X_i^2}{2\sigma_{z_i}}k_{i,m} \;<\; 0. \tag{15}$$

Thus, all of the liquidity provider's positions, long and short, have negative betas to changes in market volatility. When market volatility increases, the reversal portfolio loses on both sides–long positions decline while short positions advance.

Therefore, the liquidity provider is negatively exposed to systematic, undiversifiable market volatility risk. This is surprising because assets' payoffs in the model are purely *idiosyncratic* in nature, and have no correlation with the market return. Moreover, even if there were market risks, liquidity providers hold zero-investment portfolios that are neither net long nor short. Nevertheless, they have inescapably negative market volatility betas. The explanation is that, even though assets' cash flows are purely idiosyncratic, there is commonality in their expected volatilities. It is this commonality in their second moments that exposes the liquidity provider to an undiversifiable risk in the portfolio's first moment. This is important in practice, since empirically there is a very high degree of commonality in volatilities, both idiosyncratic and systematic, across all assets.

A large literature in volatility and option pricing documents that market volatility risk commands a very substantial and negative price of risk. Formally, this means that $E_0^P[\sigma_{m,1}] < E_0^Q[\sigma_{m,1}]$, i.e., risk-adjusted (Q-measure) volatility expectations are substantially higher than objective expectations, so that selling insurance against market volatility risk earns a large price premium. Since the reversal portfolio has a large negative beta to this risk, the liquidity provider demands a large risk premium to bear it.

Proposition 3. *The expected payoff on the liquidity provider's portfolio from time* 0 *to time* 1 *is:*

$$E_0^P\left[\sum_{i=1}^{N} -X_i \Delta p_{i,1}\right] = \left(\sum_{i=1}^{N} \beta_{i,m}\right)\left(E_0^P[\sigma_{m,1}] - E_0^Q[\sigma_{m,1}]\right) > 0. \tag{16}$$

Hence, the return to liquidity provision reflects the large volatility risk exposure of the reversal portfolios, as captured by their market volatility betas.

For testing the model, it will be useful to create a cross-section of reversal portfolios with differing exposures to market volatility risk. Equations (13) and (14) show this can be done by forming a cross-section of reversal portfolios based on the magnitude of assets' time-0 price changes (returns). Within each reversal portfolio, one can approximate the liquidity providers' position (X_i) in each asset by weighting it by a proxy for σ_{z_i}. Al-

though σ_{z_i} is not directly observable, we can proxy for it using the asset's volume, which is proportional to σ_{z_i}. Thus, the model predicts that we can obtain a cross-section of exposures to market volatility risk by sorting assets into volume-weighted reversal portfolios according to their time-0 returns.

3 Data and summary statistics

Stock screens: We construct reversal portfolios using data from CRSP. We exclude stocks with missing market capitalization or trading volume, and stocks with prices below $1 (penny stocks). Following Collin-Dufresne and Daniel (2014), we also exclude stocks that are within one day of an earnings announcement as recorded in COMPUSTAT. Earnings announcements are overwhelmingly public information-driven (as opposed to liquidity-driven) events, which confounds measurement of the returns to liquidity provision.[8]

Sample selection: The sample is daily from April 9, 2001 to December 31, 2016, which covers 3,958 trading days. The starting date corresponds to "decimalization," the transition from fractional to decimal pricing on the New York Stock Exchange and NASDAQ. As shown by Bessembinder (2003), decimalization saw a large decrease in effective trading costs, consistent with increased competition among liquidity providers. This implies that the returns to liquidity provision prior to decimalization reflect monopolistic rents.

Portfolio formation: Each day, we double-sort stocks independently into quintiles by market capitalization and deciles by normalized return. The sort by market capitalization is motivated by the fact that return reversals decrease with size (Avramov, Chordia and Goyal, 2006). The normalized return is computed as the z-score of each stock's log return relative to its distribution over the past 60 trading days (i.e. by subtracting the mean and dividing by the standard deviation). The normalization ensures that the outer decile portfolios are not dominated by stocks with the highest volatility.

The sort by normalized return captures variation in liquidity demand. In the model,

[8]Consistent with this view, the literature on post-earnings announcement drift shows that earnings announcements are associated with return continuation rather than reversal (e.g. Bernard and Thomas, 1989). Interestingly, Sadka (2006) finds that this continuation is partly explained by liquidity risk, but does not explore a connection to volatility risk.

liquidity demand is given by the order flow variable X_i. Since prices are linear in order flow, price changes (returns) give us a proxy for liquidity demand. Prices can change for non-liquidity reasons such as the release of public news. This makes them a noisy proxy for liquidity demand (this is why we remove earnings announcement days). The identification assumption we are making is that this noise is uncorrelated with exposure to volatility risk.

We weight stocks within each double-sorted portfolio by their dollar volume on the sorting day. This mimics the economic exposure of liquidity providers, helping to capture the risks that they face. Dollar volume is computed as trading volume times lagged price. Lagging the price simply ensures that we are not weighting by the day's return a second time after sorting on it.

Following Nagel (2012), we hold each portfolio for one to five trading days. This captures the returns to low-frequency liquidity provision. Whereas liquidity provision in recent years has been dominated by high-frequency trading, low-frequency liquidity provision remains important as long as imbalances between ultimate buyers and sellers persist for more than one day. The presence of a reversal premium suggests that they do. Our results show that the reversal premium has not declined with the rise of high-frequency trading, suggesting that low-frequency liquidity provision remains scarce.

Aggregate factors: We use the excess CRSP value-weighted market return as the market risk factor. We compute excess returns by subtracting the risk-free rate (the return on the 1-month T-Bill from CRSP).

We obtain the VIX index from the CBOE.

We measure aggregate idiosyncratic volatility, IVOL, as the value-weighted average squared abnormal return across stocks. Abnormal returns are calculated using a market model. IVOL is annualized to give it the same units as VIX.

The VIX return: We use S&P 500 options data from OptionMetrics to calculate the VIX return (the data ends on April 29, 2016). The VIX index is a model-free measure of the implied volatility of the S&P 500 index (as proposed by Britten-Jones and Neuberger, 2000). Specifically, the squared VIX is the price of a basket of options whose payoff is the realized variance of the S&P 500 over the next 30 days. Because this basket changes

from day to day, the change in the squared VIX itself is not a valid return. However, the percentage change in the price of a basket of options used to construct VIX on a given day is a valid return. We refer to it as the VIX return.

To construct the VIX return, we first replicate the VIX itself by following the methodology provided by the CBOE.[9] The replication is very accurate: our replicated VIX has a 99.83% correlation with the official VIX. The VIX return is the percentage change in the price of the basket of options used to construct our replicated VIX.

To maintain its constant target maturity of 30 days, the VIX is computed from a weighted average of the prices of two baskets of options, one with maturity less than 30 days (known as the near-term VIX) and one with maturity greater than 30 days (the far-term VIX). The CBOE publishes these under the tickers VIXN and VIXF, respectively. We construct their returns from the prices of the baskets that replicate them. The VIX return is simply the weighted average of the VIXN and VIXF returns, with the weights set to achieve an average maturity of 30 days.

3.1 Summary statistics

Table 1 presents summary statistics for the reversal portfolios. We focus on long-short strategies across return deciles, e.g. 1–10 (labeled "Lo–Hi"), 2–9, 3–8, 4–7, and 5–6. These are the return reversal strategies. The Lo–Hi strategy is based on the outermost return deciles and hence carries the strongest reversal signal.

From the top panel of Table 1, there are about one hundred stocks on each side (long and short) of each portfolio. The counts are slightly higher in the outer deciles for small stocks and slightly lower for large stocks. The reason is that small stocks have return distributions with fatter tails, which makes them more likely to have extreme returns even after normalizing by their volatility.

The second panel of Table 1 reports market capitalizations. The typical stock in the largest quintile is worth $50 billion, three orders of magnitude larger than the smallest

[9] See the CBOE white paper at https://www.cboe.com/micro/vix/vixwhite.pdf. The current methodology uses SPX weekly options as well as the traditional monthly ones. It was first implemented on October 6, 2014. Prior to this date, only the traditional monthly options were used.

quintile and two orders larger than the middle quintile. The largest stocks account for 96.4% of the total value of all the portfolios, making them the by far the most important quintile in economic terms (the smallest stocks account for less than 0.1%).

The third panel of Table 1 shows the sorting-day returns (without normalizing). Since small stocks are more volatile than large stocks, their sorting-day returns are larger in magnitude. For instance, the sorting-day return of the Lo–Hi strategy for the smallest stocks is -24.36%, whereas for the largest stocks it is -7.45%.

The fourth panel of Table 1 looks at share turnover. It shows that the Lo–Hi strategy has double the turnover of the middle decile strategies. This higher turnover, together with the large sorting-day return and subsequent reversal, indicates that these stocks experience an outward shift in investors' demand for liquidity. Liquidity providers absorb this demand shift, allowing for higher turnover, but charge for it by allowing prices to overshoot so they can revert later.

Finally, the bottom panel of Table 1 shows the illiquidity measure from Amihud (2002), i.e. the absolute value of the return divided by dollar volume (multiplied by 10^6 for readability).[10] As expected, illiquidity is strongly decreasing in size: the largest stocks have illiquidity measures that are three orders of magnitude smaller than for the smallest stocks. Liquidity is thus relatively cheap among the largest stocks (per dollar traded). Yet even among these stocks illiquidity is three times higher for the Lo–Hi strategy than for the middle deciles, indicating that liquidity provision is costly.

Panel A of Table 2 presents summary statistics for the daily changes in VIX, near-term VIX (VIXN), and far-term VIX (VIXF). While the three series look similar, VIXN changes are slightly more volatile with a standard deviation of 2.19 points versus 1.71 points for VIX. The distribution of VIXN changes also has slightly fatter tails with a 90% confidence interval ranging from -2.61 to 3.07 versus -2.22 to 2.39 for VIX.

Panel A of Table 2 also presents summary statistics for the return series associated with VIX, VIXN, and VIXF. The average VIX return is -1.54% per day. This number is similar to estimates of the variance premium (e.g. Carr and Wu, 2008; Bollerslev, Tauchen

[10]Since we want to measure illiquidity on the sorting day, we do not average the measure for each stock over time as is customary. Instead, we average it across stocks within each portfolio and then over time for the portfolio itself.

and Zhou, 2009; Drechsler and Yaron, 2010).[11] As expected, the VIX return is highly right-skewed, with a 99^{th} percentile of 66.61%.

The average VIXN return is -2.01%, which is larger than the average VIX return (the average VIXF return is smaller). The difference is partly explained by the slightly higher volatility of VIXN changes in Panel A. Yet, as Panel B shows, the VIXN return is larger than the VIX return even per unit of beta with respect to its underlying index. In particular, the VIXN return is -35 bps per unit of beta (-2.01% divided by 5.696), while the VIX return is only -22 bps. This shows that the price of short-term volatility risk (as captured by VIXN) is higher than the price of relatively longer-term volatility risk (as captured by VIX).

Finally, the VIXN return has a somewhat lower correlation with VIXN changes than the VIX return has with VIX changes (61% versus 73%, based on the square root of the R^2). This is explained by the fact that the VIXN return depends more strongly on the current day's realized variance (the "dividend" of the strategy), while the VIX return, which is more long-term, depends more on expected future realized variance (the "capital gain"). The changes in both VIX and VIXN omit the realized variance component, hence the lower correlation of the VIXN return with VIXN changes than the VIX return with VIX changes. Nevertheless, both correlations are very high, which suggests that these returns capture well the premium paid for hedging volatility risk.

4 Empirical results

4.1 Reversal strategy returns

Table 3 shows the post-formation returns of the reversal strategies, focusing on the five-day horizon. From Panel A, the Lo–Hi strategy delivers a five-day return of 27 bps for the largest stocks. This corresponds to an annual return of 13.5%, which is economically large. Using the standard deviation in Panel B, the associated annual Sharpe ratio is 0.6.

[11] The variance premium is the realized variance over a period of 30 days divided by the squared VIX at the start of the period. Since the variance premium can only be computed at 30-day horizons, comparing it to the VIX return requires dividing by 21 (the number of trading days in a typical 30-day period). This implicitly assumes a flat term structure of volatility risk premia within the 30-day period.

As expected, the reversal returns decline as we move from the Lo–Hi strategy toward the middle deciles, reaching near zero for the 5–6 strategy. Interestingly, standard deviations also decline, suggesting that there is greater risk in the extreme portfolios, and that this risk is not diversified at the portfolio level.

The reversal returns increase as we move from large stocks toward small stocks. This is consistent with Avramov, Chordia and Goyal (2006), who show that reversal returns are especially large among small stocks. For the smallest stocks, the Lo–Hi strategy delivers a massive 116 bps five-day return (the Sharpe ratio is 0.8). About half of this return (56 bps) is earned on the very first day, which is likely due to the well-known bid-ask bounce effect.[12] This effect does not impact large stocks, which have much narrower bid-ask spreads. Accordingly, the first-day return of the Lo–Hi strategy for the largest stocks is only 3 bps (see Figure 3 below).

Panel C of Table 3 reports the CAPM alphas of the reversal strategies, which are obtained from the time-series regressions:

$$R^p_{t,t+5} \;=\; \alpha^p + \sum_{s=1}^{5} \beta^p_s R^M_{t+s} + \epsilon^p_{t,t+5}, \tag{17}$$

where $R^p_{t,t+5}$ is the cumulative excess return of portfolio p from t to $t+5$ and R^M_{t+s} is the excess market return on day $t+s$, $s = 1 \ldots 5$. Allowing for separate coefficients β^p_s for each day s of the holding period captures the changes in exposure that occur as stocks revert back from the initial spike in liquidity demand on the portfolio formation date.[13] To account for the overlapping nature of the holding-period returns, we compute Newey-West standard errors with five lags and report the associated t-statistics in Panel D.

Panel C shows that the CAPM cannot account for the returns of the reversal strategies.

[12] As shown by Roll (1984), stocks with a low (high) return on the portfolio formation date are more likely to have finished the day at the bid (ask). Since they are equally likely to finish the following day at the bid and ask, on average they experience a reversal of about half the bid-ask spread. This effect does not extend beyond the first day. For the purposes of this paper, it can be argued that the bid-ask bounce is part of the return to liquidity provision. The counter-argument is that it only applies to the first share traded and that it is in part due to the discreteness of the price grid.

[13] We note that since daily market returns have a small negative autocorrelation (-6%), the sum of the daily coefficients does not exactly equal the coefficient that would obtain in a regression on the cumulative market return over the holding period, though the two are very close. The same is true of VIX changes, which we use in Section 4.3 below.

In all cases, the CAPM alphas are about the same as the raw average returns. For instance, the Lo–Hi strategy for the largest stocks has an alpha of 25 bps, only marginally lower than its 27-bps average return. The associated t-statistic is 4.51, strongly rejecting the null hypothesis that the CAPM prices this strategy. The same is true across all size quintiles in the Lo–Hi and 2–9 strategies. Overall, the table shows robust evidence of a substantial reversal premium.

4.2 Predicting reversals using VIX

Table 4 shows results from predictive regressions of reversal strategy returns on the level of VIX on the portfolio formation date. Panel A reports predictive loadings at the five-day horizon (times a hundred). Focusing on the Lo–Hi strategy, the loadings are similar across size quintiles, ranging from 2.94 to 7.01. From Panel B, they are all highly statistically significant except for the smallest stocks where significance is marginal. For the largest stocks, the coefficient is 5.37, implying that a one-point increase in VIX is associated with a 5.37 bps higher return over the subsequent five days. This number is large relative to the 27-bps average return of the strategy.

Panel C of Table 4 shows the R^2 coefficients of the predictive regressions. Focusing again on the Lo–Hi strategy, the R^2 is lowest for the smallest quintile (0.09%) and highest for the largest quintile (2.18%). The difference is explained by the fact that small stocks have much higher idiosyncratic volatility than large ones. Not all of this volatility is diversified away in the reversal portfolios because they contain only about a hundred stocks and because their returns are weighted by dollar volume (equal weighting would lead to more diversification). In terms of economic magnitude, the 2.18% R^2 for the largest quintile Lo–Hi strategy is extremely high for five-day returns.[14]

Overall, Table 4 corroborates and extends the main finding of Nagel (2012) that expected reversal strategy returns are increasing in the level of VIX. Nagel's (2012) results are largely representative of small stocks because they are not value-weighted and rely

[14]In particular, following Campbell and Thompson (2007), it is about three times the strategy's squared five-day Sharpe ratio. Thus, an investor using VIX to time the large-stock reversal strategy would see a three-fold increase in expected return relative to a buy-and-hold strategy.

on raw rather than normalized returns. The evidence presented here shows that the predictability is present, and in terms of R^2 is even stronger, among large stocks.

Nagel (2012) interprets the predictive power of VIX as evidence that liquidity providers face a value-at-risk (VAR) constraint.[15] A rise in VIX causes the VAR constraint to tighten, which leads to less liquidity provision and a higher price of liquidity. Thus, a rise in VIX shifts the liquidity supply curve inward.

We can test whether VIX shifts the liquidity supply curve by looking at the *quantity* of liquidity, as well as the price. The top panel of Figure 2 plots the turnover of the large-stock reversal strategy against VIX. We compute this turnover by first taking the weighted average turnover within the long and short sides of the portfolio, then averaging across the two sides and over a 60-day forward-looking window.

Figure 2 shows that the turnover of the large-stock reversal strategy is strongly increasing in the level of VIX. The raw correlation between the two series is very high at 39%. This shows that VIX positively predicts both the price (the expected return) and the quantity (turnover) of liquidity provision among large stocks. Thus, the dominant feature of the data is that VIX shifts the liquidity *demand* curve.

The bottom panel of Figure 2 shows that when VIX is high the reversal strategy becomes riskier. It plots the annualized volatility of the large-stock reversal strategy, calculated over the same 60-day rolling window as turnover. The figure shows that volatility is strongly increasing in VIX. The raw correlation is 64% and the economic magnitude is large: a one-point increase in VIX is associated with a 0.386-point increase in the reversal strategy's volatility. Therefore, when VIX is high, the reversal strategy becomes much riskier, which could explain why its expected return rises.

4.3 Reversal strategy volatility risk

The next question we tackle is whether the risk in the reversal strategy is diversifiable or if it has a systematic component. It is certainly plausible that some idiosyncratic variance

[15] While VAR constraints are arguably important for small stocks, it is less likely that they significantly impact liquidity provision among large stocks, which are heavily traded and hence less dependent on specialized intermediaries.

remains at the portfolio level, even among the largest stocks. At the same time, we know from the the variance risk premium literature (e.g., Drechsler and Yaron, 2010) that a higher VIX is associated with higher systematic risk, both in terms of the volatility of the market return and the volatility of volatility itself. Our model predicts that the latter type of systematic risk, volatility risk, is important for the returns to liquidity provision.

In this section we test whether the reversal strategy is exposed to volatility risk by regressing its return on VIX changes:

$$R^p_{t,t+5} \quad = \quad \alpha^p + \sum_{s=1}^{5} \beta_s^{p,VIX} \Delta VIX_{t+s} + \sum_{s=1}^{5} \beta_s^{p,M} R^M_{t+s} + \epsilon^p_{t,t+5}. \tag{18}$$

These regressions have the same form as (17), but with the daily VIX changes included alongside the market return. As before, we compute t-statistics based on Newey-West standard errors with five lags. We begin by focusing on the cumulative coefficients, $\sum_{s=1}^{5} \beta_s^{p,VIX}$, and we look at their individual components in the next section.

Table 5 presents the results. Panel A reports the betas from a specification with only VIX changes. Focusing first on the Lo–Hi strategy, the VIX betas are consistently negative and highly statistically significant. They are slightly larger for small stocks but overall similar across size quintiles. They are still negative and statistically significant for the 2–9 strategy and decline steadily toward zero for the inner deciles.

The estimated cumulative beta is -0.64 for the large-stock Lo–Hi strategy. This means that the strategy loses 64 bps if VIX rises on average by one point per day over the five trading days of the holding period. Such a rise is not unusual: it is 1.3 times the five-day standard deviation of VIX changes. Its impact, on the other hand, is substantial: it is two and a half times the strategy's abnormal return.

Panel B of Table 5 adds the market return as an additional control. This addresses a potential concern that the negative VIX exposure reflects market risk rather than volatility risk. The table shows that this is not the case. Betas with respect to VIX changes remain similar to those in Panel A. In the case of the large-stock Lo–Hi strategy, the beta actually increases slightly in magnitude to -0.71 and remains highly significant.

The strong negative correlation between the market return and VIX does make the be-

tas in Panel B noisier than those in Panel A, as reflected in the somewhat lower t-statistics. Overall, however, Table 5 clearly shows that the reversal strategy has a substantial negative exposure to volatility risk, as captured by VIX.

4.4 Reversal strategy dynamics

Figure 3 plots the dynamics of the reversal strategies for large stocks. They allow us to see how average returns, predictive loadings, and volatility risk exposures evolve with horizon. This is useful because there is no single horizon at which we expect liquidity providers to unwind their positions.

From Panel A, average returns rise roughly linearly with horizon, leveling off slightly toward the five-day mark. They are similar for the Lo–Hi and 2–9 strategies and uniformly decrease toward the inner-most deciles. The steady pattern in Panel A is reassuring because it shows that the reversal returns for large stocks are not driven by bid-ask bounce, which can only affect the first day of the holding period. Rather, the pattern is consistent with a stable premium paid over time.

Panel B of Figure 3 shows that the predictive loadings on the level of VIX follow the same steady linear pattern as the average returns. This bolsters the evidence that VIX predicts reversal returns in a robust way.

Panels C and D show that the same is true of the betas with respect to VIX changes, whether we control for the market return (Panel D) or not (Panel C). The increasing pattern of betas indicates that stocks in the reversal portfolio are about equally exposed to volatility risk throughout the holding period. The fact that these risk loadings line up with the average returns in Panel A suggests that they may be able to explain them. We test this proposition in the next section.

4.5 Fama-MacBeth regressions

In this section we run two-stage Fama-MacBeth regressions to see if volatility risk exposure can account for the average returns of the reversal strategies. The first stage of the Fama-MacBeth regressions estimates betas as in Section 4.3 (see (18)). We again sum the

individual coefficients to obtain portfolio-level betas, e.g. $\beta^{p,VIX} = \sum_{s=1}^{5} \beta_s^{p,VIX}$.[16] These betas enter the second-stage regression:

$$R_{t,t+5}^{p} = \lambda_t^{VIX} \beta^{p,VIX} + \lambda_t^{M} \beta^{p,M} + \epsilon_{t,t+5}^{p}. \tag{19}$$

Note that by omitting an intercept in the second-stage regression, we require it to price the zero-beta rate. This prevents a situation where test assets with small differences in betas are priced by an implausibly high price of risk (Lewellen, Nagel and Shanken, 2010).

The second-stage regression provides us with time series of the factor premia. We report the averages of these premia:

$$\lambda^{VIX} = \sum_{t=1}^{T} \lambda_t^{VIX} \qquad \lambda^{M} = \sum_{t=1}^{T} \lambda_t^{M}. \tag{20}$$

We use the time series variation in the factor premia to calculate t statistics based on Newey-West standard errors to account for overlap in the holding periods.

Using the factor premia (20), we calculate a pricing error for each portfolio. We also calculate the root-mean-squared pricing error across all portfolios and a p-value for the hypothesis that the pricing errors are jointly equal to zero.

The results of the Fama-Macbeth regressions are presented in Table 6. Panel A contains the factor premia. The first row reports a specification with the market return as the sole factor. It earns a significant positive risk premium of 3 bps, similar to its average daily return. The root-mean-squared pricing error is 18 bps and the model is rejected with a p-value close to zero. The second row of Panel A adds the change in VIX as a second factor. The change in VIX earns a large and highly significant premium of -49 bps. Thus, an asset that rises by 1% when VIX rises by 1 point is predicted to have an average return of -49 bps per day. The root-mean-squared pricing error is a smaller 14 bps but the model is also rejected.

Panel B shows the pricing errors of the underlying long-short reversal strategies. The

[16] Summing up the individual coefficients effectively imposes the same price of risk on each day of the holding period. This is natural in a frictionless setting but is potentially restrictive under segmented markets (e.g., "slow-moving capital," see Duffie, 2010).

top set of pricing errors are for the one-factor market model. As expected, the results are very close to the CAPM alphas in Table 3. In particular, the Lo–Hi strategy has a large positive pricing error of 113 bps for small stocks and 25 bps for large stocks. These pricing errors are also highly statistically significant.

The pricing errors of reversal strategies decline substantially when we include the change in VIX as a second factor. For the largest stocks, the pricing error of the Lo–Hi strategy drops from 25 bps to -7 bps and becomes insignificant. The pricing errors of the third and fourth largest quintiles are also eliminated. Only the smallest and second smallest quintiles retain significant pricing errors, although they are between a third and a half smaller. This pattern, where small stocks retain substantial pricing errors while large stocks see their pricing errors eliminated, explains why the root-mean-squared pricing error in Panel A drops by a relatively small amount when we add VIX changes as a second factor. By squaring the pricing errors, the root-mean-squared pricing error overweights outliers, which in this case are the smallest stocks.

Figure 4 plots the average returns of the reversal strategies against their predicted values based on the Fama-MacBeth regression estimates. Each size quintile is represented by a different shape and color and contains five data points corresponding to the five long-short strategies across deciles.

The left plot shows clearly that the market model cannot explain the returns of the reversal strategies. The average returns display wide variation along the vertical axis while the predicted returns are confined to a very narrow range along the horizontal axis. Moreover, the predicted returns are all close to zero. Thus, predicted and average returns differ by a wide margin.

By contrast, the right panel of Figure 4 shows that adding VIX changes largely explains the average returns of the reversal strategies. Predicted returns cover a wide range along the horizontal axis, lining up well with the average returns along the 45-degree line. Only the extreme small-stock strategies (the Lo–Hi and 2–9 strategies for the smallest quintile and the Lo–Hi strategy for the second smallest quintile) lie away from the 45-degree line. Thus, the reversal strategies among small stocks earn abnormally high returns even after accounting for volatility risk. This is consistent with some degree of

segmentation in liquidity provision for these stocks. Yet, since they represent only about 0.4% of total market value, this segmentation has low economic significance. For the larger stocks, which are the economically important ones, Table 6 and Figure 4 show that volatility risk can account for the returns to liquidity provision.

4.6 Restricting the price of volatility risk

The next question we address is whether the price of volatility risk needed to price the reversal strategies is consistent with the price of volatility risk that prevails in other asset markets. Answering this question sheds further light on the broader question of whether the returns to liquidity provision reflect intermediation frictions or more widespread economic risks.

The natural place to obtain the price of volatility risk is from options markets. The VIX index itself is computed from the price of a basket of options whose payoff replicates the realized variance of the S&P 500 over a 30-day window. Yet while we can think of VIX as a price, the change in VIX is not a valid return because the VIX basket changes from day to day, hence it does not reflect the premium paid for bearing exposure to VIX. This problem is straight-forward to solve, however, as it simply requires tracking the price of a given VIX basket to the day after it was used to construct VIX. The percentage change in this price over that day is the VIX return, R^{VIX}.

A remaining issue is that liquidity providers are likely to have a shorter horizon than the 30-day window captured by VIX because their inventory turns over faster. This is why throughout the paper we focus on reversal returns over one to five trading days, as does the rest of the literature. Yet constructing a five-day volatility index is infeasible because options are known to behave erratically as they approach expiration (this is why no option with fewer than seven days to expiration is used in the construction of VIX). To get a measure of shorter-term volatility risk, we use VIXN, which is the near-term component used in the construction of VIX (see Section 3). VIXN has an average maturity of 22 days, which is significantly shorter than VIX. Its associated return is the VIXN return, R^{VIXN}.

Table 2 shows the moments of the distribution of R^{VIX} and R^{VIXN}. Their means are

-1.54% and -2.01%, respectively, and their loadings on ΔVIX and $\Delta VIXN$ are 6.938 and 5.696. Thus, the price of ΔVIX exposure is -22 bps per unit, while the price of $\Delta VIXN$ exposure is -35 bps per unit. This shows that the price of shorter-term volatility risk is higher than relatively longer-term volatility risk.

Table 7 shows the pricing errors of the reversal strategies, where we restrict the price of volatility risk to -22 bps when we use VIX as the pricing factor (Panel A) and -35 bps when we use VIXN as the pricing factor (Panel B). We also restrict the price of market risk to 3 bps, equal to the average return of the market factor. To obtain the pricing errors reported in the table, we first calculate the betas of the reversal portfolios in the same way as in (18), but we replace ΔVIX with $\Delta VIXN$ in the case of Panel B. We then multiply these betas by the restricted prices of risk and subtract the resulting predicted returns from the average returns to calculate the pricing errors.

Panel A shows that the restricted price of VIX exposure is able to explain a substantial fraction of the reversal strategy returns, although it does not explain them fully. In particular, the pricing error of the Lo–Hi strategy for large stocks declines by about three fifths, from 25 bps to 11 bps, though it remains marginally statistically significant. The reduction is similar in levels but smaller in percentage terms for the small-stock strategies whose initial pricing errors are quite a bit larger.

Panel B shows that the restricted price of VIXN exposure explains a significantly larger fraction of the reversal strategy returns. Most prominent, the pricing error of the large-stock Lo–Hi strategy drops from 25 bps to just 1 bp and becomes statistically insignificant. Short-term volatility risk is thus able to fully explain the returns to liquidity provision among large stocks. The pricing errors of the third and fourth largest quintiles are similarly driven down and become insignificant. Only the smallest two quintiles retain substantial pricing errors, similar to those of the unrestricted Fama-MacBeth regressions in Section 4.5. Thus, as before, the ability of volatility risk to explain the returns to liquidity provision is concentrated among large stocks, which make up over 99.6% of market value. We now know that it is short-term volatility risk in particular that is able to do so.

Figure 5 depicts the average versus predicted returns of the reversal strategies with

the restricted prices of risk. The left plot uses VIX and corresponds to Panel A of Table 7, while the right plot uses VIXN and corresponds to Panel B. Both plots show a reasonable alignment between average and predicted returns, except in the case of the outer decile strategies among the smallest two quintiles. The fit is significantly better, however, when we use the shorter-term VIXN, consistent with the conclusion from Table 7. In particular, the fit is very similar to the unrestricted Fama-MacBeth regressions in Figure 4, indicating that volatility risk is priced similarly in options markets and in reversal returns. This finding favors the interpretation that the returns to liquidity provision reflect broad economic risks rather than narrow intermediation frictions.

5 Conclusion

Just how broad are the economic risks reflected in the returns to liquidity provision? While the literature on the variance risk premium emphasizes risks embedded in the consumption process of a representative agent, these risks could themselves be a reflection of the economy's dependence on the financial sector. We find this possibility particularly intriguing as it promises to further integrate the asset pricing and financial intermediation literatures.

If we were to speculate about how this integration might proceed, it could go as follows. A spike in volatility like the one that hit the U.S. economy in the summer of 2007 increases the flow of private information into asset markets. Liquidity providers respond by raising the sensitivity of asset prices to order flow, which raises the cost of liquidity for investors. As liquidity becomes more scarce, consumption and investment become misallocated throughout the economy, ultimately resulting in an economic contraction. The ex post contraction implies a high variance risk premium, which, as we showed in this paper, will be reflected in a high liquidity premium. To the extent that there are financial frictions, they would be worsened by the volatility risk exposure of liquidity providers, exacerbating the liquidity crunch, further harming the economy, amplifying the variance risk premium, and so on.

References

Acharya, Viral V, and Lasse Heje Pedersen. 2005. "Asset pricing with liquidity risk." *Journal of financial Economics*, 77(2): 375–410.

Admati, Anat R, and Paul Pfleiderer. 1988. "A theory of intraday patterns: Volume and price variability." *The Review of Financial Studies*, 1(1): 3–40.

Adrian, Tobias, and Hyun Song Shin. 2010. "Liquidity and leverage." *Journal of financial intermediation*, 19(3): 418–437.

Akerlof, George A. 1970. "The market for" lemons": Quality uncertainty and the market mechanism." *The quarterly journal of economics*, 488–500.

Amihud, Yakov. 2002. "Illiquidity and stock returns: cross-section and time-series effects." *Journal of financial markets*, 5(1): 31–56.

Amihud, Yakov, and Haim Mendelson. 1986. "Asset pricing and the bid-ask spread." *Journal of financial Economics*, 17(2): 223–249.

Avramov, Doron, Tarun Chordia, and Amit Goyal. 2006. "Liquidity and autocorrelations in individual stock returns." *The Journal of finance*, 61(5): 2365–2394.

Bao, Jack, Jun Pan, and Jiang Wang. 2011. "The illiquidity of corporate bonds." *The Journal of Finance*, 66(3): 911–946.

Bekaert, Geert, Robert J Hodrick, and Xiaoyan Zhang. 2012. "Aggregate idiosyncratic volatility." *Journal of Financial and Quantitative Analysis*, 47(6): 1155–1185.

Bernard, Victor L, and Jacob K Thomas. 1989. "Post-earnings-announcement drift: delayed price response or risk premium?" *Journal of Accounting research*, 1–36.

Bessembinder, Hendrik. 2003. "Trade Execution Costs and Market Quality after Decimalization." *Journal of Financial and Quantitative Analysis*, 38(4): 747âĂŞ777.

Bollerslev, Tim, and Viktor Todorov. 2011. "Tails, fears, and risk premia." *The Journal of Finance*, 66(6): 2165–2211.

Bollerslev, Tim, George Tauchen, and Hao Zhou. 2009. "Expected stock returns and variance risk premia." *The Review of Financial Studies*, 22(11): 4463–4492.

Brennan, Michael J, and Avanidhar Subrahmanyam. 1996. "Market microstructure and asset pricing: On the compensation for illiquidity in stock returns." *Journal of financial economics*, 41(3): 441–464.

Britten-Jones, Mark, and Anthony Neuberger. 2000. "Option prices, implied price processes, and stochastic volatility." *The Journal of Finance*, 55(2): 839–866.

Brunnermeier, Markus K. 2009. "Deciphering the liquidity and credit crunch 2007–2008." *The Journal of economic perspectives*, 23(1): 77–100.

Brunnermeier, Markus K, and Lasse Heje Pedersen. 2008. "Market liquidity and funding liquidity." *The review of financial studies*, 22(6): 2201–2238.

Brunnermeier, Markus K, and Yuliy Sannikov. 2014. "A macroeconomic model with a financial sector." *The American Economic Review*, 104(2): 379–421.

Campbell, John Y, and Samuel B Thompson. 2007. "Predicting excess stock returns out of sample: Can anything beat the historical average?" *The Review of Financial Studies*, 21(4): 1509–1531.

Campbell, John Y, Martin Lettau, Burton G Malkiel, and Yexiao Xu. 2001. "Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk." *The Journal of Finance*, 56(1): 1–43.

Carr, Peter, and Liuren Wu. 2008. "Variance risk premiums." *The Review of Financial Studies*, 22(3): 1311–1341.

Chordia, Tarun, Asani Sarkar, and Avanidhar Subrahmanyam. 2004. "An empirical analysis of stock and bond market liquidity." *The Review of Financial Studies*, 18(1): 85–129.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam. 2000. "Commonality in liquidity." *Journal of financial economics*, 56(1): 3–28.

Collin-Dufresne, Pierre, and Kent Daniel. 2014. "Liquidity and return reversals." *Unpublished working paper. New York: Columbia University*.

Collin-Dufresne, Pierre, and Vyacheslav Fos. 2016. "Insider trading, stochastic liquidity, and equilibrium prices." *Econometrica*, 84(4): 1441–1475.

Constantinides, George M, Jens Carsten Jackwerth, and Alexi Savov. 2013. "The puzzle of index option returns." *Review of Asset Pricing Studies*, 3(2): 229–257.

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez. 2017. "The price of variance risk." *Journal of Financial Economics*, 123(2): 225–250.

Diamond, Douglas W, and Robert E Verrecchia. 1981. "Information aggregation in a noisy rational expectations economy." *Journal of Financial Economics*, 9(3): 221–235.

Drechsler, Itamar. 2013. "Uncertainty, Time-Varying Fear, and Asset Prices." *The Journal of Finance*, 68(5): 1843–1889.

Drechsler, Itamar, and Amir Yaron. 2010. "What's vol got to do with it." *The Review of Financial Studies*, 24(1): 1–45.

Duffie, Darrell. 2010. "Presidential address: asset price dynamics with slow-moving capital." *The Journal of finance*, 65(4): 1237–1267.

Eisfeldt, Andrea L. 2004. "Endogenous liquidity in asset markets." *The Journal of Finance*, 59(1): 1–30.

Foster, F Douglas, and Sean Viswanathan. 1990. "A theory of the interday variations in volume, variance, and trading costs in securities markets." *The Review of Financial Studies*, 3(4): 593–624.

Glosten, Lawrence R, and Paul R Milgrom. 1985. "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders." *Journal of financial economics*, 14(1): 71–100.

Gorton, Gary, and George Pennacchi. 1990. "Financial intermediaries and liquidity creation." *The Journal of Finance*, 45(1): 49–71.

Gromb, Denis, and Dimitri Vayanos. 2002. "Equilibrium and welfare in markets with financially constrained arbitrageurs." *Journal of financial Economics*, 66(2): 361–407.

Grossman, Sanford J, and Joseph E Stiglitz. 1980. "On the impossibility of informationally efficient markets." *The American economic review*, 70(3): 393–408.

Grossman, Sanford J, and Merton H Miller. 1988. "Liquidity and market structure." *the Journal of Finance*, 43(3): 617–633.

Hameed, Allaudeen, Wenjin Kang, and Shivesh Viswanathan. 2010. "Stock market declines and liquidity." *The Journal of Finance*, 65(1): 257–293.

Hasbrouck, Joel, and Duane J Seppi. 2001. "Common factors in prices, order flows, and liquidity." *Journal of financial Economics*, 59(3): 383–411.

Hellwig, Martin F. 1980. "On the aggregation of information in competitive markets." *Journal of economic theory*, 22(3): 477–498.

Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh. 2016. "The common factor in idiosyncratic volatility: Quantitative asset pricing implications." *Journal of Financial Economics*, 119(2): 249–283.

He, Zhiguo, and Arvind Krishnamurthy. 2013. "Intermediary asset pricing." *The American Economic Review*, 103(2): 732–770.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2015. "The impact of Treasury supply on financial sector lending and stability." *Journal of Financial Economics*, 118(3): 571–600.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica*, 53(6): 1315–1336.

Lehmann, Bruce N. 1990. "Fads, martingales, and market efficiency." *The Quarterly Journal of Economics*, 105(1): 1–28.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken. 2010. "A skeptical appraisal of asset pricing tests." *Journal of Financial economics*, 96(2): 175–194.

Longstaff, Francis A, Jun Pan, Lasse H Pedersen, and Kenneth J Singleton. 2011. "How sovereign is sovereign credit risk?" *American Economic Journal: Macroeconomics*, 3(2): 75–103.

Manela, Asaf, and Alan Moreira. 2017. "News implied volatility and disaster concerns." *Journal of Financial Economics*, 123(1): 137–162.

Moreira, Alan, and Alexi Savov. 2017. "The Macroeconomics of Shadow Banking." *The Journal of Finance*, 72(6): 2381–2432.

Nagel, Stefan. 2012. "Evaporating Liquidity." *The Review of Financial Studies*, 25(7): 2005–2039.

Pástor, L'uboš, and Robert F Stambaugh. 2003. "Liquidity risk and expected stock returns." *Journal of Political economy*, 111(3): 642–685.

Roll, Richard. 1984. "A simple implicit measure of the effective bid-ask spread in an efficient market." *The Journal of finance*, 39(4): 1127–1139.

Sadka, Ronnie. 2006. "Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk." *Journal of Financial Economics*, 80(2): 309 – 349.

Stoll, Hans R. 1978. "The supply of dealer services in securities markets." *The Journal of Finance*, 33(4): 1133–1151.

Todorov, Viktor. 2009. "Variance risk-premium dynamics: The role of jumps." *The Review of Financial Studies*, 23(1): 345–383.

Table 1: Reversal strategy summary statistics

This table shows summary statistics for the reversal portfolio strategies. Each day, stocks are sorted into deciles by normalized return and quintiles by market capitalization. The normalized return is calculated based on the return distribution over the previous 60 trading days. Stocks with share price less than $1 and stocks with an earnings announcement within one trading day of the sorting date are excluded. The portfolios are constructed using dollar volume as weights. The sample is from April 9, 2001 to December 31, 2016.

	Number of stocks				
	Lo-Hi	2-9	3-8	4-7	5-6
Small	137	109	102	100	100
2	119	109	107	107	107
3	105	109	110	111	112
4	94	109	114	115	116
Big	92	112	115	114	113

	Market cap (billions)				
	Lo-Hi	2-9	3-8	4-7	5-6
Small	0.05	0.05	0.05	0.05	0.05
2	0.16	0.16	0.17	0.17	0.17
3	0.43	0.44	0.44	0.44	0.44
4	1.35	1.37	1.37	1.37	1.37
Big	49.57	54.15	56.02	56.02	55.40

	Sorting-day returns (%)				
	Lo-Hi	2-9	3-8	4-7	5-6
Small	−24.36	−6.92	−4.21	−2.34	−0.74
2	−17.54	−6.05	−3.73	−2.07	−0.64
3	−14.77	−5.43	−3.34	−1.87	−0.60
4	−11.97	−4.70	−2.92	−1.64	−0.52
Big	−7.45	−3.43	−2.13	−1.18	−0.38

	Share turnover (%)				
	Lo-Hi	2-9	3-8	4-7	5-6
Small	10.28	7.37	6.71	6.28	6.07
2	7.84	4.45	3.85	3.60	3.42
3	6.41	3.11	2.63	2.46	2.36
4	5.59	2.76	2.44	2.25	2.21
Big	3.28	2.13	1.99	1.89	1.83

	Amihud illiquidity ($\times 10^6$)				
	Lo-Hi	2-9	3-8	4-7	5-6
Small	33.60	21.08	14.42	10.34	8.58
2	5.73	3.98	2.79	2.07	1.70
3	1.36	1.00	0.71	0.52	0.43
4	0.30	0.22	0.16	0.11	0.09
Big	0.03	0.02	0.01	0.01	0.01

Table 2: VIX return summary statistics

This table shows summary statistics for the VIX return (Panel A) and regressions of VIX returns on VIX changes (Panel B). The VIX index is constructed from the prices of two baskets of options, one with expiration less than 30 days and one with expiration greater than 30 days. These are called VIXN and VIXF, respectively, where the additional letters stand for "near" and "far". The VIX return, R^{VIX}, is the percentage change in the price of the portfolio of options used to construct VIX. The VIXN and VIXF returns, R^{VIXN} and R^{VIXF}, are the percentage changes in the prices of the portfolios of options used to construct VIXN and VIXF, respectively. The sample is from April 9, 2001 to April 29, 2016.

Panel A: Summary statistics

	Mean	St.Dev.	P1	P5	P10	P90	P95	P99
ΔVIX	-0.00	1.71	-4.39	-2.22	-1.51	1.55	2.39	5.21
ΔVIXN	-0.00	2.19	-5.74	-2.61	-1.81	1.96	3.07	6.96
ΔVIXF	-0.00	1.53	-3.69	-2.02	-1.45	1.48	2.23	4.95
R^{VIX}	-1.54	16.49	-24.37	-16.77	-13.92	14.25	25.79	66.61
R^{VIXN}	-2.01	20.47	-31.22	-21.63	-17.29	16.57	30.72	75.56
R^{VIXF}	-0.87	12.17	-20.60	-13.57	-10.70	11.02	19.45	47.19

Panel B: Regressions of VIX returns on ΔVIX

	R^{VIX}	R^{VIXN}	R^{VIXF}
	(1)	(2)	(3)
ΔVIX	6.938***		
	(0.106)		
ΔVIXN		5.696***	
		(0.120)	
ΔVIXF			6.244***
			(0.080)
Constant	-1.511***	-1.986***	-0.849***
	(0.184)	(0.264)	(0.123)
Obs.	3,788	3,787	3,787
R^2	0.529	0.372	0.615

Table 3: Reversal strategy returns

Average returns, standard deviations, CAPM alphas, and associated t-statistics for the reversal strategies over a five-day holding period. Stocks are sorted into deciles by normalized return and quintiles by market capitalization. Penny stocks and earnings announcements are excluded. The portfolios are weighted by dollar volume. The CAPM alphas are the intercepts from the regression:

$$R^p_{t,t+5} \;=\; \alpha^p + \sum_{s=1}^{5} \beta^p_s R^M_{t+s} + \epsilon^p_{t,t+5},$$

where $R^p_{t,t+5}$ is the cumulative excess return on portfolio p from t to $t+5$ and R^M_{t+s} is the excess market return on day $t+s$, $s = 1 \ldots 5$. The t-statistics are based on Newey-West standard errors with five lags to account for the overlap in returns. The sample is from April 9, 2001 to December 31, 2016.

Panel A: Average returns

	5-day average return (%)				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	1.16	0.56	0.21	0.05	0.04
2	0.65	0.30	0.17	0.03	−0.03
3	0.35	0.24	0.01	0.11	−0.01
4	0.22	0.23	0.13	0.06	0.01
Big	0.27	0.25	0.18	0.11	0.05

Panel B: Standard deviations

	5-day standard deviation (%)				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	10.54	7.11	6.19	5.99	5.89
2	6.44	4.59	4.03	3.82	3.74
3	5.11	3.18	2.71	3.02	2.28
4	3.88	2.47	2.10	1.92	1.77
Big	3.25	2.28	1.78	1.55	1.31

Panel C: CAPM alphas

	5-day CAPM alpha (%)				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	1.14	0.55	0.20	0.04	0.04
2	0.62	0.30	0.16	0.02	−0.03
3	0.34	0.23	−0.00	0.10	−0.01
4	0.20	0.23	0.13	0.06	0.01
Big	0.25	0.24	0.18	0.11	0.05

Panel D: CAPM alpha t-statistics

	5-day CAPM alpha t-statistic				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	6.57	4.88	1.97	0.44	0.39
2	5.83	3.96	2.43	0.33	−0.50
3	3.85	4.31	−0.03	2.25	−0.14
4	3.13	5.54	3.82	2.00	0.31
Big	4.51	6.48	6.24	4.35	2.50

Table 4: Predicting reversal strategy returns using VIX

Coefficients, t-statistics, and R^2 from predictive regressions of reversal strategy returns over a five-day holding period on the level of VIX on the formation date. Stocks are sorted into deciles by normalized return and quintiles by market capitalization. Penny stocks and earnings announcements are excluded. The portfolios are weighted by dollar volume. The predictive regressions are

$$R^p_{t,t+5} \;=\; a_p + b_p VIX_t + \epsilon^p_{t,t+5},$$

where $R^p_{t,t+5}$ is the cumulative excess return on portfolio p from day t to day $t+5$ and VIX_t is the VIX index on day t, the portfolio formation date. The t-statistics are based on Newey-West standard errors with five lags to account for the overlap in returns. The sample is from April 9, 2001 to December 31, 2016.

Panel A: Predictive coefficients

	5-day return VIX loading ($\times 10^2$)				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	3.53	3.48	3.51	2.72	0.16
2	7.01	3.14	2.68	1.40	−0.27
3	4.84	2.98	1.16	0.93	−0.10
4	2.94	2.33	1.52	−0.04	0.44
Big	5.37	3.69	1.74	0.67	0.08

Panel B: t-statistics

	5-day return VIX t-statistic				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	1.95	2.62	2.55	2.16	0.13
2	3.70	2.77	3.66	2.05	−0.34
3	4.03	3.42	1.60	1.59	−0.21
4	2.80	3.70	2.83	−0.11	1.24
Big	3.98	4.64	3.26	1.70	0.22

Panel C: R^2

	5-day return VIX R^2 (%)				
	Lo–Hi	2–9	3–8	4–7	5–6
Small	0.09	0.19	0.26	0.16	0.00
2	0.95	0.37	0.35	0.11	0.00
3	0.72	0.70	0.15	0.08	0.00
4	0.46	0.71	0.42	0.00	0.05
Big	2.18	2.11	0.77	0.15	0.00

Table 5: Volatility risk of the reversal strategy

This table shows the betas of the reversal portfolios on VIX changes at the five-day horizon. The betas are estimated from the regression:

$$R^p_{t,t+5} = \alpha_p + \sum_{s=1}^{5} \beta_s^{p,VIX} \Delta VIX_{t+s} + \sum_{s=1}^{5} \beta_s^{p,M} R^M_{t+s} + \epsilon^p_{t,t+5},$$

where $R^p_{t,t+5}$ is the cumulative excess return on portfolio p from day t to day $t+5$, and ΔVIX_{t+s} and R^M_{t+s} are the change in VIX and the excess market return on day $t+s$, respectively. Panel A omits the market return while Panel B includes it. The panels report the summed coefficients $\sum_{s=1}^{5} \beta_p^{VIX,s}$ and $\sum_{s=1}^{5} \beta_p^{M,s}$ (alternative horizons are plotted in Figure 3.) The t-statistics are based on Newey-West standard errors with five lags to account for the overlap in returns. The sample is from April 9, 2001 to December 31, 2016.

Panel A: Δ VIX betas

	5-day ΔVIX beta					5-day ΔVIX beta t-statistic					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	−0.81	−0.49	−0.57	−0.26	−0.31	Small	−2.66	−3.29	−3.02	−1.31	−1.78
2	−0.82	−0.34	−0.24	−0.31	0.03	2	−3.47	−2.21	−1.82	−2.39	0.25
3	−0.57	−0.26	−0.36	−0.32	−0.01	3	−3.63	−2.52	−3.82	−4.16	−0.09
4	−0.54	−0.26	−0.18	0.01	−0.04	4	−4.30	−2.78	−2.70	0.13	−0.93
Big	−0.64	−0.34	−0.09	−0.01	−0.01	Big	−4.28	−3.09	−1.34	−0.25	−0.31

Panel B: Δ VIX betas, controlling for the market return

	5-day ΔVIX beta					5-day ΔVIX beta t-statistic					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	−0.75	−0.60	−0.80	−0.13	−0.46	Small	−1.53	−2.40	−2.43	−0.43	−1.75
2	−0.65	−0.44	−0.35	−0.04	−0.08	2	−1.87	−1.75	−1.78	−0.20	−0.47
3	−0.62	−0.12	−0.25	−0.14	−0.14	3	−2.23	−0.72	−1.71	−0.85	−1.26
4	−0.33	−0.30	−0.34	0.04	−0.01	4	−1.41	−2.28	−3.00	0.48	−0.13
Big	−0.71	−0.37	−0.18	0.05	−0.13	Big	−3.33	−2.42	−1.46	0.62	−1.62

Table 6: Fama-Macbeth regressions

The table shows factor premia and pricing errors from Fama-Macbeth regressions of the five-day reversal portfolios. In the first stage, we regress the portfolio returns on the market return and the change in the VIX index on each day of the holding period. We then sum the coefficients to obtain portfolio-level betas. In the second stage, we regress the portfolio returns on these betas (without an intercept) to obtain factor premia and pricing errors. Next to each premium and pricing error are the associated t-statistics, which are based on a Newey-West standard error. Also reported is a root-mean-squared pricing error (r.m.s.) and associated p-value. The sample is from April 9, 2001 to December 31, 2016.

Panel A: Factor premia

	Market	t-stat.	Δ VIX	t-stat.	R.m.s.	p-value
(1)	0.03	2.06	.	.	0.18	0.00
(2)	0.05	3.04	−0.49	−8.57	0.14	0.00

Panel B: Pricing errors

(1) Market only

	Pricing error					t statistic					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	1.13	0.55	0.20	0.04	0.03	Small	6.52	4.86	1.97	0.42	0.37
2	0.61	0.29	0.16	0.02	−0.03	2	5.78	3.94	2.43	0.30	−0.50
3	0.33	0.23	−0.00	0.10	−0.00	3	3.82	4.33	−0.08	2.16	−0.13
4	0.20	0.22	0.13	0.06	0.01	4	3.08	5.58	3.81	2.01	0.30
Big	0.25	0.23	0.18	0.11	0.05	Big	4.47	6.47	6.21	4.38	2.58

(2) Market and Δ VIX

	Pricing error					t statistic					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	0.79	0.28	−0.17	−0.02	−0.17	Small	5.08	2.59	−1.88	−0.20	−1.88
2	0.32	0.10	0.00	−0.01	−0.06	2	3.29	1.39	0.04	−0.11	−1.08
3	0.06	0.17	−0.12	0.03	−0.07	3	0.71	3.30	−2.74	0.65	−1.83
4	0.04	0.09	−0.02	0.08	0.00	4	0.70	2.35	−0.68	2.57	0.09
Big	−0.07	0.07	0.10	0.13	−0.00	Big	−1.50	1.99	3.63	5.22	−0.10

Table 7: Fama-Macbeth regressions with a restricted price of risk

The table shows pricing errors from Fama-Macbeth regressions with a restricted price of risk based on the returns of the VIX and near-term VIX (VIXN) replicating portfolios. Panel A is based on the VIX return and Panel B is based on the VIXN return). The restricted price of risk is obtained by dividing the average VIX and VIXN returns by their betas with respect to changes in VIX and VIXN, respectively (see Table 2). The price of market risk is just the average market return over the sample. These restricted prices of risk are multiplied by the betas of the reversal portfolios with respect to the market and VIX or VIXN changes (as reported in Table 5). This gives a predicted return for each reversal portfolio. The reported pricing errors are the differences between the average returns and the predicted returns for long-short strategies between the low- and high-return deciles within each size quintile. The sample is from April 9, 2001 to December 31, 2016.

Panel A: Market and ΔVIX

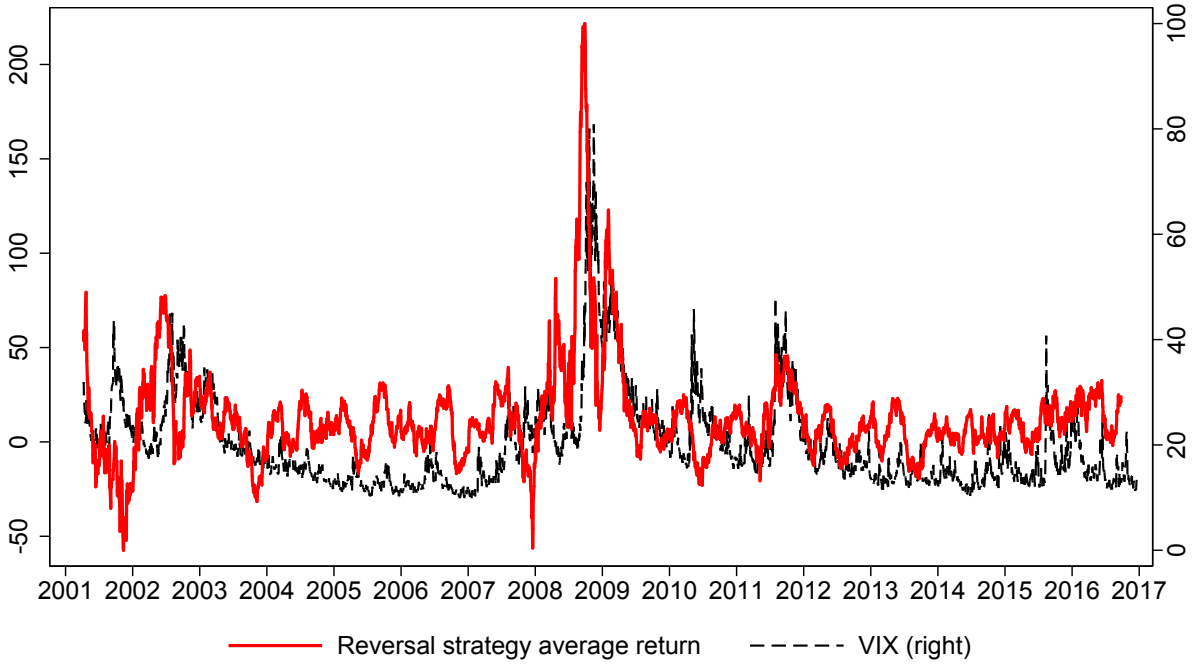
	Pricing error					*t* statistic					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	0.99	0.44	0.04	0.02	−0.05	Small	5.76	3.86	0.43	0.19	−0.58
2	0.50	0.21	0.09	0.01	−0.04	2	4.66	2.83	1.42	0.20	−0.77
3	0.22	0.21	−0.05	0.07	−0.03	3	2.48	3.88	−1.08	1.60	−0.91
4	0.14	0.17	0.06	0.07	0.01	4	2.11	4.14	1.89	2.27	0.23
Big	0.11	0.16	0.14	0.12	0.03	Big	2.04	4.52	5.01	4.80	1.38

Panel B: Market and ΔVIXN

	Pricing error					*t* statistic					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	0.87	0.34	0.04	−0.04	−0.09	Small	5.04	3.02	0.43	−0.37	−0.94
2	0.42	0.16	0.07	−0.00	−0.05	2	3.93	2.16	1.01	−0.07	−0.80
3	0.12	0.12	−0.08	0.04	−0.03	3	1.38	2.19	−1.80	0.83	−0.87
4	0.05	0.10	0.01	0.07	−0.01	4	0.74	2.48	0.18	2.48	−0.46
Big	0.01	0.13	0.11	0.11	0.03	Big	0.21	3.57	3.81	4.51	1.33

Figure 1: Reversal strategy average return and volatility risk

The figure shows the average return and volatility risk of the large-stock reversal strategy against the VIX index. The reversal strategy takes a long position in the lowest normalized return decile portfolio and a short position in the highest normalized return decile portfolio within the largest stock quintile. It holds these positions for five trading days. The average return of the reversal strategy is computed over a 60-day forward-looking window. The systematic volatility of the reversal strategy due to VIX changes is obtained from 60-day rolling regressions of the five-day return on the corresponding five daily changes in the VIX index and taking the annualized standard deviation of the fitted component. The sample is from April 9, 2001 to December 31, 2016.



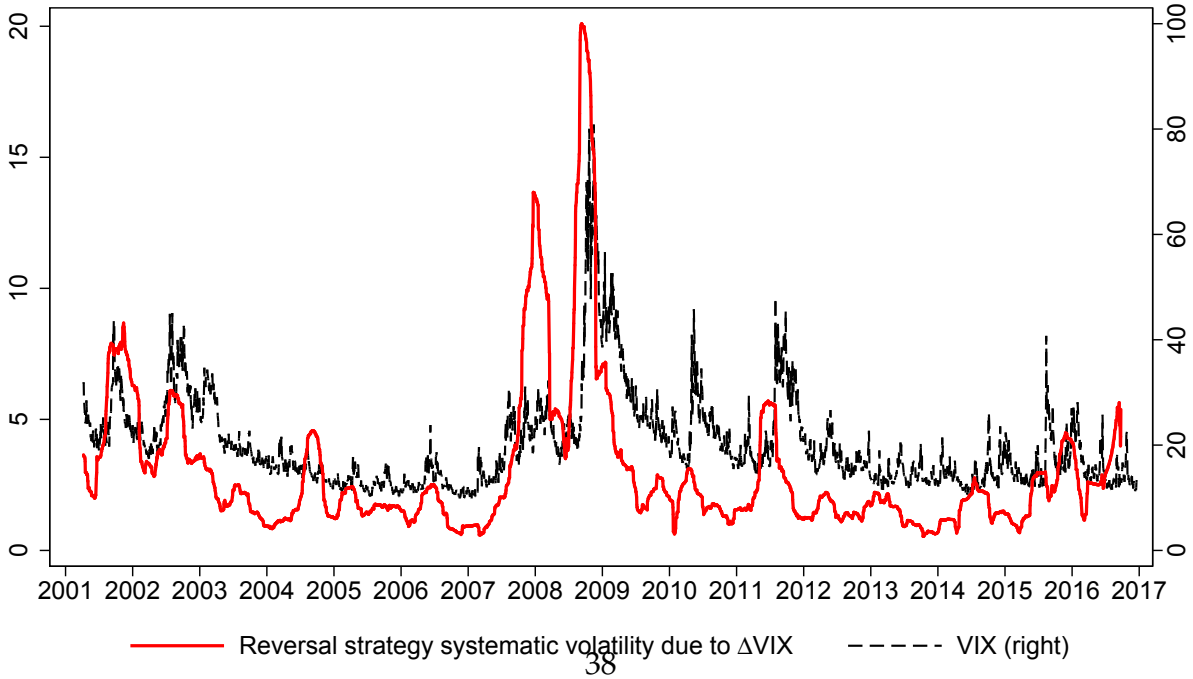
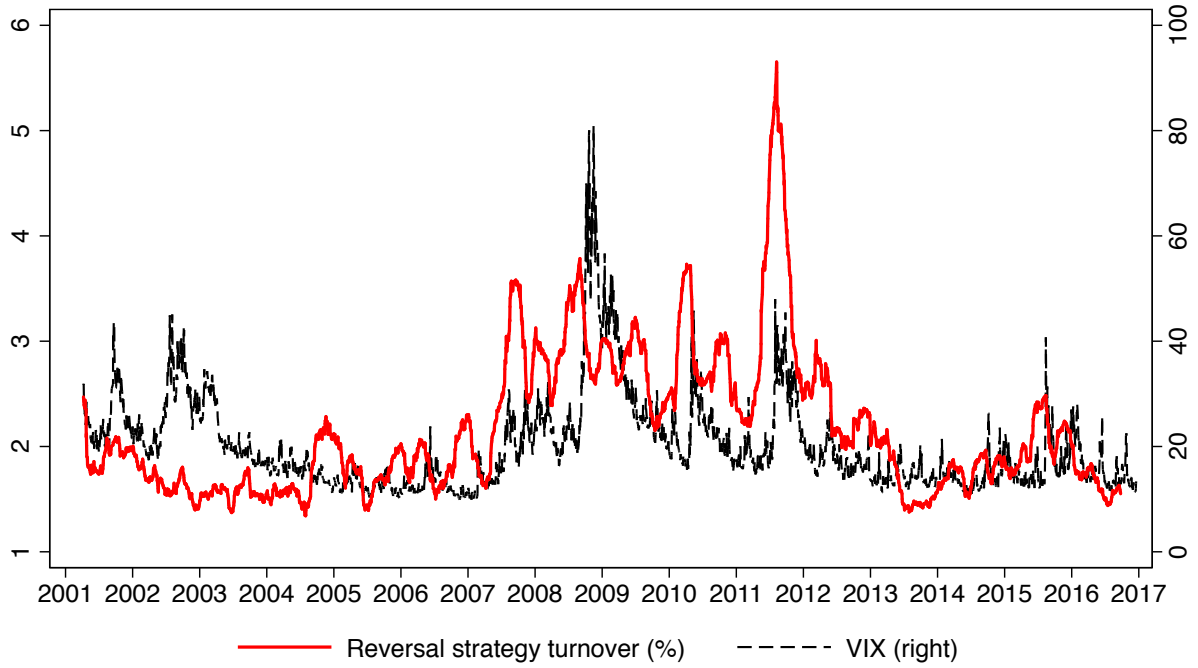


Figure 2: Reversal strategy turnover and volatility

The figure shows the volatility and turnover of the large-stock reversal strategy against the VIX index. The reversal strategy takes a long position in the lowest normalized return decile portfolio and a short position in the highest normalized return decile portfolio within the largest stock quintile. It holds these positions for five trading days. The turnover of the reversal strategy is the average turnover of the long and short side on the portfolio formation date. The figure plots annualized volatility and daily turnover averaged over a 60-day forward-looking window. The sample is from April 9, 2001 to December 31, 2016.



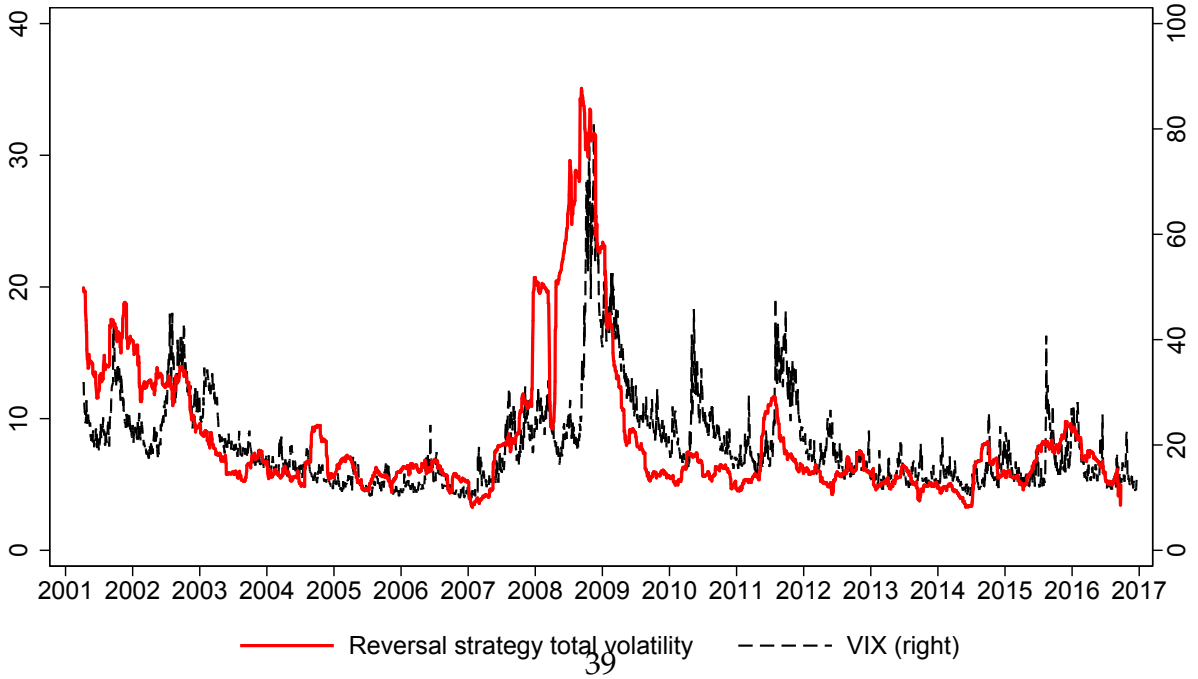
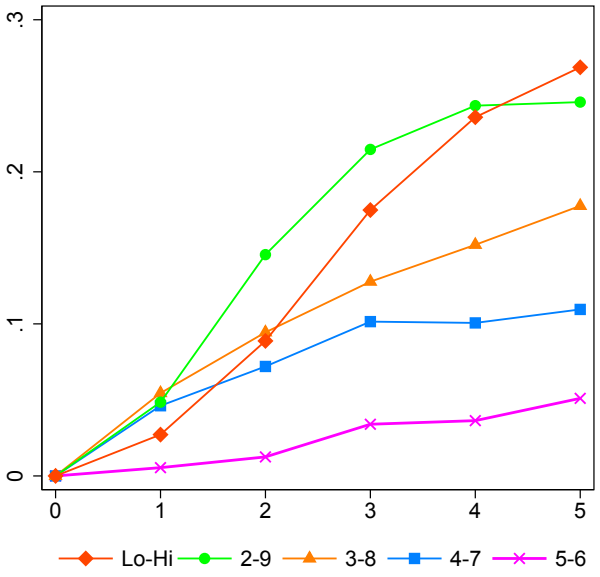


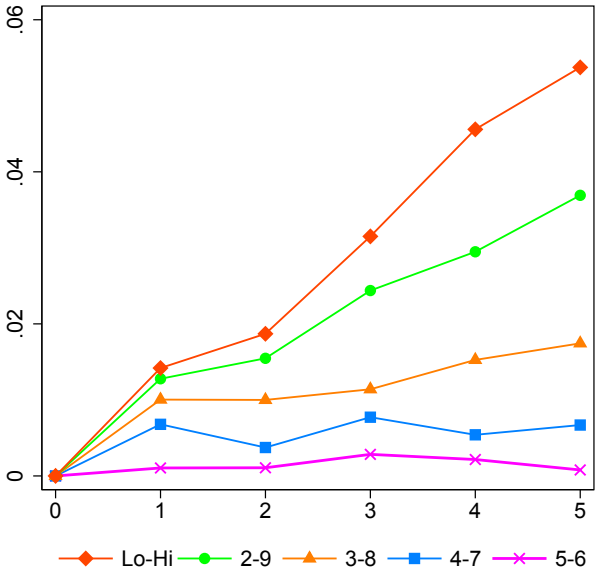
Figure 3: Dynamics of the reversal strategy

The figure shows the average returns, predictive loadings on VIX, and cumulative exposures to VIX changes for reversal strategy portfolios of large-cap stocks. The Lo–Hi reversal strategy takes a long position in the lowest normalized return decile portfolio and a short position in the highest normalized return decile portfolio within the largest stock quintile. The other strategies are defined analogously using the inner deciles. Each strategy is held for one to five trading days. The horizontal axis shows the holding period. On the vertical axis, Panel A plots the average return, Panel B plots the predictive loading on the level of VIX as of the formation date, Panel C plots the cumulative betas with respect to VIX changes, and Panel D plots the cumulative betas while controlling for the market return (see (18)). The sample is from April 9, 2001 to December 31, 2016.

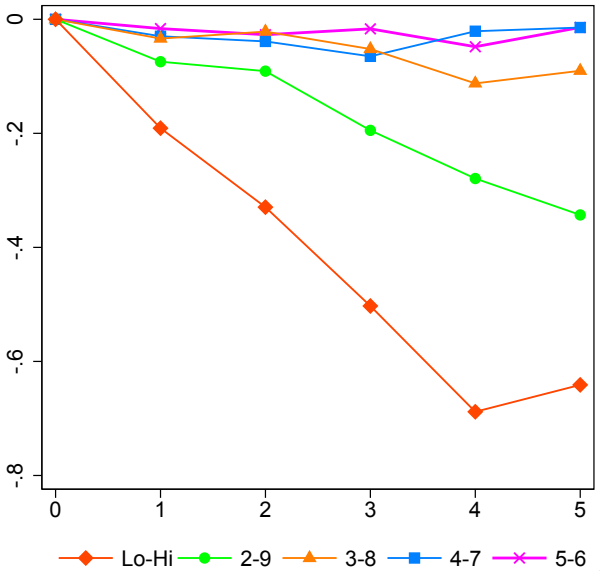
Panel A: Average returns



Panel B: Predictive loadings on VIX



Panel C: Cumulative exposure to ΔVIX



Panel D: Cumulative exposure to ΔVIX, controling for R^M

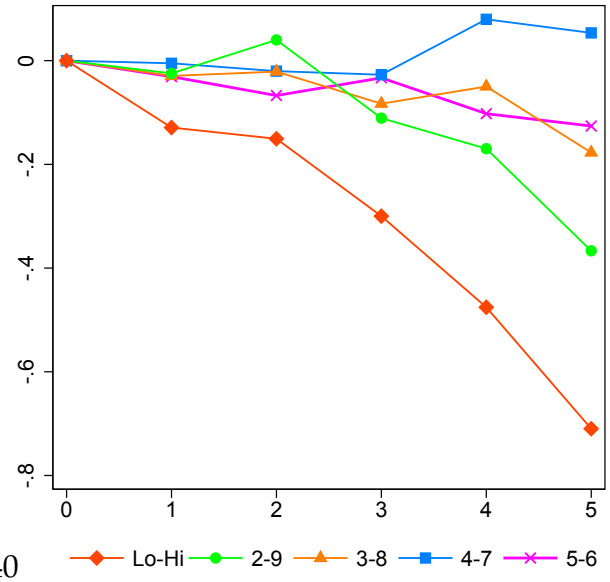


Figure 4: Fama-MacBeth regressions: average versus predicted returns

The figure shows average and predicted returns from Fama-MacBeth regressions of the five-day reversal portfolios. In the first stage, we regress the portfolio returns on the market return and the change in the VIX index on each day of the holding period. We then sum the coefficients to obtain portfolio-level betas. In the second stage, we regress the portfolio returns on these betas (without an intercept) to obtain factor premia. The panels plot the average returns of long-short reversal strategies (there are five of these for each size quintile) against their predicted returns based on the Fama-MacBeth regression estimates. The sample is from April 9, 2001 to December 31, 2016.

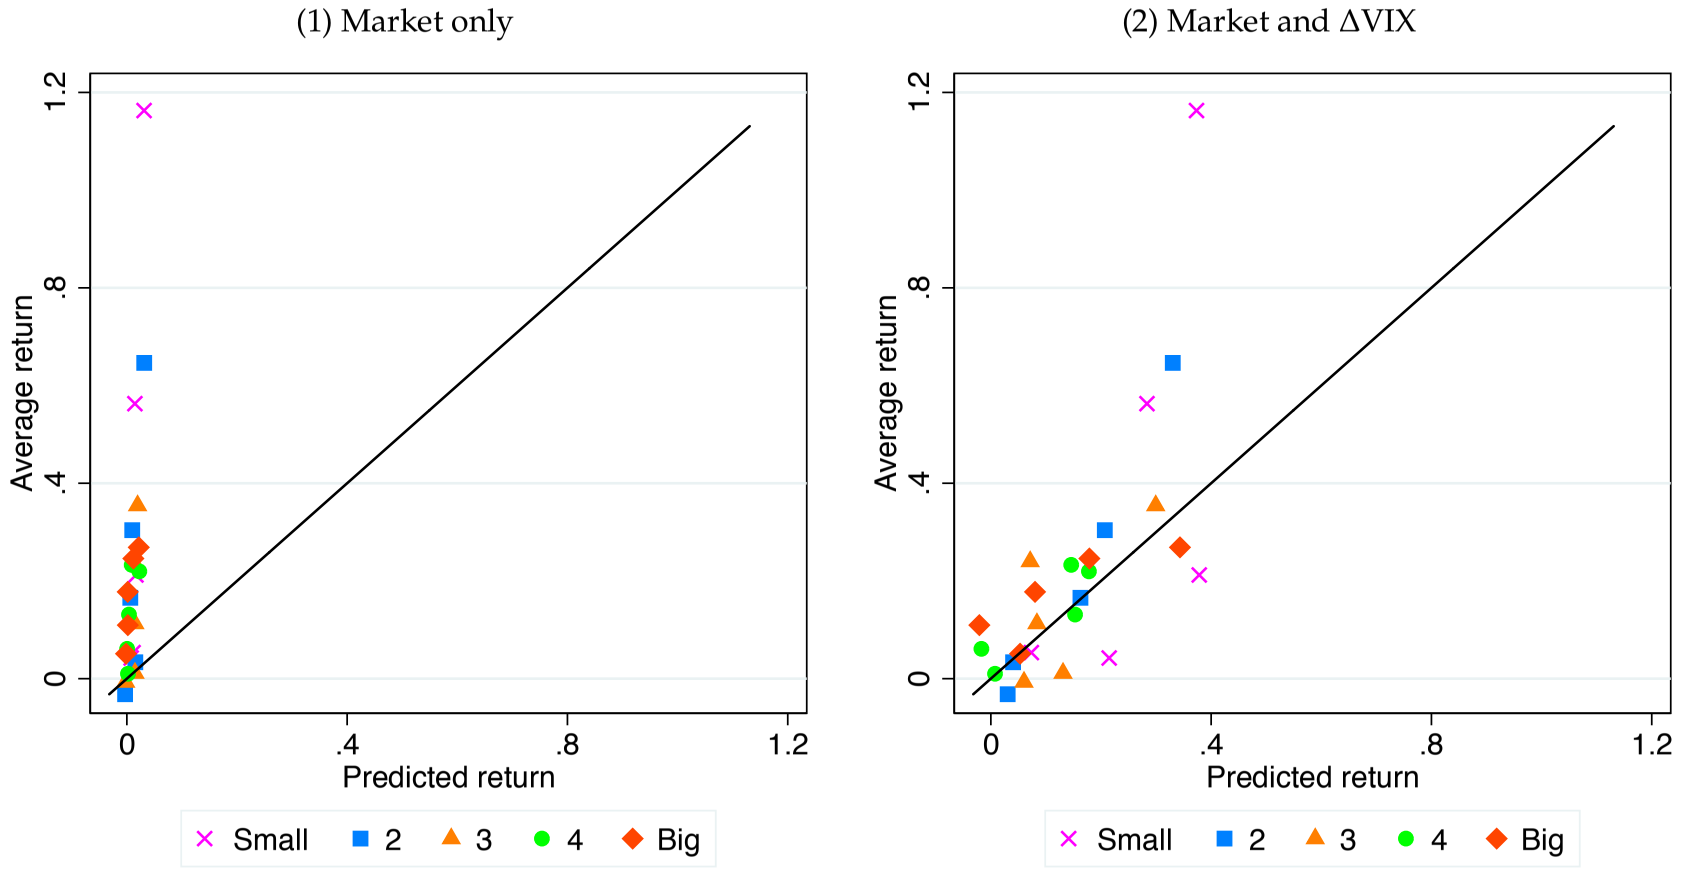


Figure 5: Fama-Macbeth regressions with a restricted price of risk: average versus predicted returns

The figure shows average and predicted returns from Fama-MacBeth regressions with a restricted price of risk based on the returns of the VIX and near-term VIX (VIXN) replicating portfolios. The restricted price of risk is obtained by dividing the average VIX and VIXN returns by their betas with respect to changes in VIX and VIXN, respectively (see Table 2). The price of market risk is just the average market return over the sample. These restricted prices of risk are multiplied by the betas of the reversal portfolios with respect to the market and VIX or VIXN changes (as reported in Table 5). This gives a predicted return for each reversal portfolio. The panels plot the average returns of long-short reversal strategies (there are five of these for each size quintile) against their predicted returns. The sample is from April 9, 2001 to December 31, 2016.

