

Liquidity and Volatility

Itamar Drechsler, Alan Moreira, Alexi Savov*

September 2021

Abstract

Liquidity provision is a bet against private information: if private information turns out to be higher than expected, liquidity providers lose. Since information generates volatility, and volatility co-moves across assets, liquidity providers have a negative exposure to aggregate volatility shocks. Aggregate volatility shocks carry a large premium, hence this negative exposure explains why liquidity provision earns high average returns. We show this by incorporating uncertainty about the amount of private information into an otherwise standard model of liquidity provision. We test the model in the cross section of short-term reversals, which mimic the portfolios of liquidity providers. As predicted by the model, reversals have large negative betas to aggregate volatility shocks. These betas explain reversals' average returns with the same price of volatility risk that prevails in option markets. Volatility risk thus explains the liquidity premium among stocks, and why it increases in volatile times. Our results provide a novel view of the risks and returns to liquidity provision.

JEL Codes:

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

*Drechsler: University of Pennsylvania and NBER, idrechsl@wharton.upenn.edu; Moreira: University of Rochester and NBER, alan.moreira@simon.rochester.edu; Savov: New York University and NBER, asavov@stern.nyu.edu. We thank Markus Brunnermeier, Pete Kyle, Mathieu Fournier, Yunzhi Hu, David Schreindorfer, Daniel Schmidt, Mark Westerfield, and Zhaogang Song for discussions; Kent Daniel, Greg Duffee, Ron Kaniel, Bryan Kelly, Maureen O'Hara, and Christian Opp for their comments; as well as participants at NBER Asset Pricing, Cornell, CBOE Derivatives Conference, European Finance Association, Insead Finance Conference, American Finance Association, University of Texas at Austin, Washington University at St. Louis, Toulouse School of Economics, Duke-UNC Asset Pricing Conference, JHU-Carey Finance Conference, University of Rochester, Temple University, Norges Bank, Federal Reserve Board, Securities and Exchange Commission, UC Davis, and Wharton for their comments. We thank Liz Casano for excellent research assistance. An earlier draft circulated under the title "Liquidity Creation as Volatility Risk".

Liquidity provision—intermediation that lowers the cost of buying and selling assets—is a central function of the financial system. In this paper, we show theoretically and empirically that liquidity providers have a negative exposure to market volatility shocks. Since market volatility shocks carry a large negative risk premium, the variance risk premium, the negative exposure of liquidity providers explains why liquidity provision earns a substantial positive premium, the liquidity premium. It also explains why this premium rises in volatile times (Nagel, 2012). Our results thus show that liquidity provision and volatility insurance share the same risks and returns.

Why are liquidity providers exposed to market volatility shocks? The basic problem liquidity providers face is adverse selection from investors with private information. In standard models such as Kyle (1985), the amount of private information is constant and known to liquidity providers. It therefore poses no risk to them. This is a simple assumption but clearly unrealistic. Information generates volatility, hence the fact that volatility is highly time-varying implies that information is highly time-varying. Some of this variation is due to public news, but some is undoubtedly due to fluctuations in the amount of private information, which creates volatility through trading.

To capture this, we extend the model of Kyle (1985) by making the amount of private information liquidity providers face a latent variable. Liquidity providers therefore trade at prices that depend on their estimate of the amount of private information. When they estimate private information to be low, they perceive little adverse selection and trade at prices that are relatively insensitive to order flow. They buy the assets investors are selling without a big fall in price and sell the assets investors are buying without a big rise in price. Conversely, when they estimate private information to be high, they perceive a lot of adverse selection, and the sensitivity of prices to order flow increases.

The return liquidity providers earn therefore depends on whether the true amount of private information is higher or lower than they estimate. If it is higher, prices are not sensitive enough to order flow. Liquidity providers absorb too much order flow from informed investors, accumulating overvalued long positions and undervalued short positions. Their ex post return will be low. Liquidity providers thus have a negative exposure to the amount of private information in the assets they trade.

It follows that they also have a negative exposure to news about the amount of private information. Suppose news comes out that the order flow liquidity providers absorbed

contained more information than they estimated. This news reveals that their long positions are overvalued and their short positions are undervalued. As prices adjust to reflect this, liquidity providers lose.

We argue that shocks to an asset's volatility are a source of such news about the information contained in its recent order flow. Since information generates volatility, unexpectedly high volatility today reflects an unexpectedly large amount of information coming to light. The key assumption we make is that some of this unexpected information today was private information yesterday. When that is the case, today's positive volatility shock implies that yesterday's order flow contained more information than liquidity providers expected. This in turn implies that yesterday's prices were not sensitive enough to order flow. The result is that, following a positive volatility shock, prices today move further in the same direction as yesterday: assets that fell fall more, and assets that rose rise more.¹ And since liquidity providers bought the assets that fell and sold assets that rose, this continuation results in their loss. Liquidity providers thus have a negative exposure to asset volatility shocks.

If volatility shocks were uncorrelated across assets, liquidity providers could easily diversify this exposure by providing liquidity across many assets. However, as we show, the volatilities of individual assets co-move strongly. The co-movement is present in assets' idiosyncratic volatilities, the component we are most interested in since private information is likely to be idiosyncratic in nature. Moreover, as we also show, idiosyncratic volatilities also co-move strongly with the volatility of the market portfolio. Intuitively, there are quiet times when not much information is coming out, and turbulent times when the flow of information is intense at both the idiosyncratic and aggregate levels.

This co-movement of information flows gives liquidity providers an undiversifiable negative exposure to the overall amount of private information and to market volatility. And since market volatility carries a large negative risk premium, the variance risk premium (Carr and Wu, 2008), no-arbitrage implies that liquidity provision earns a positive premium, the liquidity premium.

To generate this liquidity premium, we further extend the standard Kyle-based model,

¹Collin-Dufresne and Fos (2015) document this in the context of activist trades. They show that standard liquidity measures do not reveal the underlying amount of private information, that private information takes a long time to enter into prices, and that it only enters fully when regulation mandates its public disclosure. Disclosure thus acts as a news shock that reveals that the amount of private information was higher than expected. This causes prices to move further in the direction of activists' previous trades.

which is risk neutral, to include risk pricing. The model shows that liquidity providers have a negative exposure to the amount of private information, and that if this amount co-moves with market volatility, then liquidity providers demand a risk premium. This risk premium is given by the product of their market volatility betas and the variance risk premium. Thus, the premiums earned for providing liquidity and insuring against volatility shocks are actually compensation for the same risk.²

The model thus makes two central testable predictions: (i) liquidity providers' portfolios have negative market volatility betas, and (ii) the average returns of these portfolios are explained by the product of their market volatility betas and the variance risk premium. We test these predictions, and others, in the cross section of short-term stock reversals. These are trading strategies that mimic the portfolios of liquidity providers by buying stocks whose prices have recently fallen and selling stocks whose prices have recently risen (Lehmann, 1990; Lo and MacKinlay, 1990). Reversals thus capture the risks and returns to liquidity provision in stocks.

We build reversal portfolios using daily U.S. stock returns from 2001 to 2020, the period after “decimalization,” when liquidity provision became competitive (Bessembinder, 2003). Each day we sort stocks into quintiles by market capitalization and deciles by the day's return, normalized by its rolling standard deviation as implied by our model. Also as implied by our model, we weight the portfolios by dollar volume to proxy for the size of liquidity providers' holdings. We then construct long-short reversal strategies within each size quintile by buying the low-return deciles and selling the high-return deciles: Lo–Hi, 2–9, 3–8, 4–7, and 5–6. The outermost-decile strategies carry the strongest reversal signal and therefore capture the most intensive liquidity provision. This allows us to test our model in the cross section.

Consistent with the model, and with prior literature, the reversal strategies earn substantial returns that are not explained by their market betas. For instance, among large stocks, the Lo–Hi reversal strategy has an average five-day return of 20 bps, or about 10% per year, and an annual Sharpe ratio of 0.5.

This premium arises in our model because reversal strategies are exposed to market

²This makes our model very different from existing models, which focus on inventory risk (e.g., Stoll, 1978; Grossman and Miller, 1988; Nagel, 2012). In these models, the risk liquidity providers face is diversifiable. To generate a liquidity premium, they assume that markets are deeply segmented, which forces liquidity providers to hold concentrated portfolios. While this may apply in some cases, such as for small and thinly-traded assets, it is unlikely to explain the liquidity premium for large and heavily-traded assets.

volatility risk. We estimate this exposure by regressing their returns on changes in the VIX index, which measures the risk-adjusted expected volatility of the S&P 500 at a 30-day horizon. As the model predicts, the estimated beta is negative: a one-point increase in VIX induces a 20-bps drop in the strategy's return, a large amount relative to its average.

Figure 1 looks at the time variation in this average return by plotting it over rolling 60-day forward-looking windows against VIX. The plot shows that VIX strongly predicts the reversal strategy's return: the two series have a correlation of 46%. A predictive regression shows that a one-point increase in VIX predicts a 9-bps higher reversal return, which is again large relative to its average. This finding confirms that the main result in Nagel (2012) holds for large stocks.

This predictability arises in our model because VIX predicts the volatility risk of the reversal strategy. We test this prediction by re-running our regression of reversal returns on VIX changes on a rolling basis. We then take the volatility of the regression's fitted value (the beta times the volatility of VIX changes), which captures the volatility risk of the reversal strategy (the component of its total volatility due to market volatility shocks). Panel B of Figure 1 plots this volatility risk against VIX itself. The two series track each other very closely, including during the 2008 financial crisis and the 2020 Covid crisis. The relationship is also strong outside of crises, with an overall correlation of 58%. This shows that reversals bear more volatility risk when VIX is high, matching the pattern in average returns in Panel A. Figure 1 thus shows that liquidity provision earns a higher premium when it bears more market volatility risk, as predicted by our model.

In our model, liquidity provision has market volatility risk because the amount of private information co-moves with market volatility. Figure 2 provides evidence for this. Although the amount of private information is not directly observed, we can proxy for it with the volatility it generates. Panel A of Figure 2 plots the value-weighted average of stocks' idiosyncratic volatilities (computed from their beta-adjusted returns) against VIX. The two series are very highly related with a correlation of 92%.³

Panel B of Figure 2 looks directly at the amount of private information and its co-movement with market volatility. Specifically, it plots VIX against the component of idiosyncratic volatility that is explained by order flow.⁴ Unlike public news, private infor-

³In Internet Appendix IA.4 we document the striking fact that VIX predicts average idiosyncratic volatility as well as it does market volatility even though VIX is the risk-adjusted expectation of market volatility.

⁴We compute this component from intraday data by regressing five-minute returns on order flow. We

mation is revealed through the trading process. Order flow-induced volatility therefore provides an estimate of the amount of private information. The figure shows that this estimate co-moves tightly with VIX, with a correlation of 70%. This supports the model's assumption that the amount of private information and market volatility co-move.

Our main empirical analysis tests the model's predictions on the full cross section of reversal portfolios. We first show that the reversal strategies' average returns display the model's predicted pattern: within each size quintile, the outermost Lo-Hi strategy displays the largest return, the innermost 5-6 strategy the smallest, and the others lie in between. We confirm that these returns are not explained by market beta.

Next, we look at the strategies' market volatility betas. We find that the betas are strongly negative. Moreover, they are largest for the Lo-Hi strategies and approach zero for the 5-6 strategies, as predicted. We control for market beta to ensure it does not drive this pattern. The pattern obtains even though VIX was not involved in the construction of the portfolios. Thus, consistent with the model, the reversal strategies have large market volatility betas that align with their average returns.

We use Fama-MacBeth regressions to formally test whether volatility risk explains the returns of the reversal strategies. The estimated premium for the volatility betas is large and significant: -1.08% per unit of beta over five days. Accounting for the volatility betas reduces the pricing errors of the reversal strategies significantly. The pricing error of the large-cap Lo-Hi strategy drops from 17 to -2 bps and becomes statistically insignificant. The pricing error of the second-largest quintile Lo-Hi quintile similarly drops from 13 to -1 bps. Only the very smallest stocks' Lo-Hi pricing error remains large and significant. These stocks, which account for less than 0.2% of total market value, have similar volatility betas as larger stocks, but their reversal returns are too high to be explained by them. Outside these stocks, the volatility risk betas do a good job of explaining reversal returns and hence the liquidity premium among stocks.

Since market variance is traded directly in the index options market, we use index option returns to estimate the premium for market variance risk and test whether it explains the reversal strategies' returns given their betas. Following the literature, we do so using the basket of options that make up VIX. From the returns on this basket, we estimate that

take the daily fitted value as the order flow-induced component of the day's return. We then square these components, value-weight them across stocks, and sum over 30 calendar days (to match the horizon of VIX). Finally, we take the square root to convert the measure to a volatility.

the option-implied risk premium per unit of market volatility beta is -1.03% over five days. This number is consistent with other estimates in the literature.

We restrict the premium for volatility risk to this option-implied value (and the premium for market risk to the market's average excess return) and re-calculate the pricing errors of the reversal strategies.⁵ The restricted pricing test performs very similarly to the unrestricted Fama-MacBeth regressions, reflecting their similar price of risk estimates. The pricing errors of the largest and second-largest quintiles' Lo-Hi portfolios decrease from 17 and 13 bps to -1 bps. The pricing error of an overall liquidity portfolio, which we construct by combining the reversal strategies in a way implied by the model, drops from 17 to 1.5 bps. Thus, volatility risk explains the overall liquidity premium in stocks with the same price of risk as in option markets.

As before, small-stock reversals stand out from the overall picture. While their VIX betas are similar to those of larger stocks, they are too small to explain their very large reversal returns. This suggests that undiversified inventory risk plays a role for these stocks. To analyze this, we extend the model to incorporate inventory costs. As suggested by Nagel (2012), we allow these costs to increase in VIX. This makes reversal returns sensitive to VIX independent of our main mechanism. However, the model shows that this sensitivity is highly transitory: the impact of a VIX shock on reversal returns fully dissipates by the end of liquidity providers' holding period. The reason is that it is due to the higher inventory cost, which liquidity providers pay only so long as they hold the asset in inventory. We find that for small stocks the impact of VIX shocks is indeed transitory, consistent with a role for inventory costs.

In contrast, we find a very different pattern for medium and large stocks. For these stocks' reversals, the impact of VIX shocks is permanent, with no sign of decay well beyond plausible horizons for liquidity providers' holding periods. This is consistent with the main mechanism of our model, where an increase in private information has a permanent impact because it implies a decrease in the fundamental value of liquidity providers' holdings. Our model thus explains both the average returns and dynamic responses of large- and medium-cap stock reversals, consistent with the view that the liquidity premium in stocks is earned as compensation for volatility risk.⁶

⁵This is equivalent to including the VIX portfolio and the market as test assets and requiring that they be priced without error, see, e.g., Constantinides, Jackwerth and Savov (2013).

⁶In Appendix IA.5 we further look at a new set of test assets inspired by the model. We sort stocks

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

1 Related literature

Our paper brings together the large literatures on liquidity and volatility. The theoretical literature on liquidity emphasizes the role of asymmetric information (e.g. Hellwig, 1980; Grossman and Stiglitz, 1980; Diamond and Verrecchia, 1981; Kyle, 1985; Glosten and Milgrom, 1985; Admati and Pfleiderer, 1988). In these models, liquidity providers know exactly how much asymmetric information they face. This allows them to always break even across a sufficiently diversified portfolio. The same is true in models where asymmetric information varies over time but is still known in equilibrium (Foster and Viswanathan, 1990; Eisfeldt, 2004; Collin-Dufresne and Fos, 2016). Our theoretical contribution is to study the more realistic case where the amount of asymmetric information is not perfectly known. We show that this naturally implies that liquidity providers are exposed to volatility shocks. Moreover, this exposure becomes undiversifiable when volatility shocks co-move across assets, as they do in practice.⁷

Our paper builds on the empirical literature that uses short-term reversals to proxy for the portfolios of liquidity providers (e.g. Lehmann, 1990; Lo and MacKinlay, 1990; Hameed, Kang and Viswanathan, 2010; Nagel, 2012; Collin-Dufresne and Daniel, 2014). We contribute to this literature by showing that reversals have large negative exposures to market volatility shocks, as predicted by our model. Our paper is also related to the broader empirical literature that studies the asset pricing effects of illiquidity (Amihud and Mendelson, 1986; Brennan and Subrahmanyam, 1996; Easley and O'Hara, 2004; Chordia, Sarkar and Subrahmanyam, 2004; Amihud, 2002; Pástor and Stambaugh, 2003;

into quintiles based on the co-movement of their idiosyncratic volatility with market volatility. We then form reversals within these quintiles. These reversals are a good test of our model because it predicts that stocks with high volatility co-movement are exposed to more undiversifiable private information risk. As a result, providing liquidity in these stocks should come with a more negative volatility risk exposure and a commensurately higher average return. This is exactly what we find.

⁷Campbell et al. (2001) and Herskovic et al. (2016) find strong co-movement between asset volatilities and market volatility at the annual frequency. We complement their results by showing that the co-movement is also present and even stronger at higher frequencies. Relatedly, Chordia, Roll and Subrahmanyam (2000) and Hasbrouck and Seppi (2001) find strong co-movement in illiquidity across assets.

Acharya and Pedersen, 2005). We contribute to this literature by connecting the liquidity premium to the variance risk premium.

Our model generates a liquidity premium without relying on market segmentation frictions. This contrasts with inventory models, which assume limited diversification at the level of the liquidity provider (e.g., Stoll, 1978; Grossman and Miller, 1988; Duffie, 2010; Nagel, 2012). Among these papers, Nagel (2012) is closest to our work. Nagel (2012) builds a model with inventory frictions that explains the time series behavior of the liquidity premium under the assumption that VIX proxies for the severity of these frictions. We instead build a model without inventory frictions that shows that liquidity providers are fundamentally exposed to volatility risk. In addition to explaining the variation in the liquidity premium, the model makes the quantitative prediction that its level is explained by the level of the variance risk premium. We find that this is indeed the case.⁸

Our explanation for the liquidity premium also contrasts with the literature on intermediary asset pricing, which assumes segmentation between financial institutions and outside investors. This makes liquidity scarce at the aggregate level (Holmström and Tirole, 1998). The segmentation could be due to an equity capital constraint (e.g. He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014; Rampini and Viswanathan, 2019), a Value-at-Risk (VaR) constraint (Gromb and Vayanos, 2002; Brunnermeier and Pedersen, 2008; Adrian and Shin, 2010) or a collateral constraint (Kiyotaki and Moore, 1997; Geanakoplos, 2003; Gertler and Kiyotaki, 2010; Moreira and Savov, 2017). In our model financial institutions price payoffs using the same stochastic discount factor as outside investors, hence there is no segmentation. Yet liquidity remains scarce, even at the asset level, because of systematic fluctuations in the amount of asymmetric information across assets. Also in contrast to this literature, the risk liquidity providers face is due to information, not a tightening of financial constraints.

We draw on the option pricing literature, which shows that investors pay a large premium to insure aggregate volatility risk (Carr and Wu, 2008; Todorov, 2009; Bollerslev and Todorov, 2011). This literature further shows that the variance risk premium is highly time-varying and a powerful predictor of the equity premium (Bollerslev, Tauchen and

⁸Small stocks are an exception. The reversal returns of small stocks are too high to be explained by their volatility risk, hence inventory frictions likely play a role. Consistent with this, the methodology in Nagel (2012) places a large weight on small stocks by focusing on an equally-weighted reversal portfolio. Similarly, Hasbrouck (1988) finds that inventory costs play a role among small stocks but not large ones.

Zhou, 2009; Drechsler and Yaron, 2010; Bao, Pan and Wang, 2011; Longstaff et al., 2011; Manela and Moreira, 2017). A strand of this literature argues that the variance risk premium originates from macroeconomic risk (Drechsler and Yaron, 2010; Drechsler, 2013; Dew-Becker et al., 2017). Our paper does not take a stand on this question but instead takes the variance risk premium as given. Our contribution is to apply the insights of the volatility risk literature to the pricing of liquidity.

2 Model

We present a model based on Kyle (1985) with the key difference that liquidity providers do not know how much private information they face. There are three dates: 0, $\tau \in (0, 1)$, and 1. The risk-free rate is normalized to zero at each date. There are N risky assets $i = 1, \dots, N$ in zero net supply. For each asset there are unit masses of informed traders and liquidity traders. There is also a competitive fringe of liquidity providers that are active across all assets. Trading takes place on dates 0 and τ . Final payoffs are realized on date 1 and are given by

$$p_{i,1} = \bar{v}_i + v_i, \quad (1)$$

where \bar{v}_i is a constant and $v_i \sim N(0, \sigma_{v,i})$ is an idiosyncratic shock that is uncorrelated across assets. One can easily add an aggregate shock, but we leave it out as it is orthogonal to the mechanism we are studying. The value of \bar{v}_i is known by everyone ahead of time. Since v_i is idiosyncratic, \bar{v}_i is the price of asset i before any trading takes place. Thus, $\sigma_{v,i}$ is the volatility of asset i 's price over the whole time period from 0 to 1.

On date 0 informed traders learn v_i , whereas others do not. We follow Nagel (2012) and assume that informed traders demand y_i units of the asset, where

$$y_i = \phi v_i. \quad (2)$$

The parameter ϕ controls how aggressively informed traders trade in the direction of their private signal v_i .⁹ The uninformed liquidity traders demand z_i units of the asset, where

⁹Kyle (1985) is a version of this model with a monopolistic informed trader, which endogenizes ϕ .

$z_i \sim N(0, \sigma_{z,i}^2)$ is uncorrelated across assets.

We assume liquidity providers share risks perfectly with the rest of the economy, i.e. there is no market segmentation. As a result, they price date-1 payoffs using the aggregate stochastic discount factor (SDF) Λ_1 . The price of asset i is:

$$p_{i,t} = E_t[\Lambda_1 p_{i,1}] \equiv E_t^Q[p_{i,1}] \quad \text{for } t \in \{0, \tau\}, \quad (3)$$

where Q is the risk-adjusted (i.e., risk-neutral) probability measure corresponding to Λ_1 . By taking expectations under Q we obtain expressions that look like those in Kyle (1985), but with the difference that they take risk premia into account. This is important because we are interested in deriving the risk premium for liquidity provision.¹⁰

As in Kyle (1985), liquidity providers cannot distinguish between the trades of informed traders and liquidity traders. They only observe their sum, *net* order flow, on date 0:

$$x_i = y_i + z_i = \phi v_i + z_i. \quad (4)$$

We now depart from Kyle (1985) by relaxing the assumption that liquidity providers know the volatility of the informed traders' signal, $\sigma_{v,i}$. This means that they do not know how much private information is in the market. Moreover, we incorporate the fact that the amount of private information co-moves across assets as we saw in Figure 2. We capture this by writing $\sigma_{v,i}^2$ as the sum of a common component σ_v^2 and an idiosyncratic component $\varsigma_{v,i}^2$:

$$\sigma_{v,i}^2 = k\sigma_v^2 + \varsigma_{v,i}^2, \quad (5)$$

where $k > 0$ is the loading of idiosyncratic variance on the common factor.¹¹ The common factor creates variation in the average amount of private information across assets, which makes their idiosyncratic volatilities co-move, as is the case empirically.

It is important that liquidity providers cannot perfectly infer the volatility $\sigma_{v,i}$ of an

¹⁰The literature typically assumes that all agents are risk neutral and is thus silent on risk premia.

¹¹In Internet Appendix IA.5 we further allow stocks to have different loadings k_i on the common factor. We then estimate these loadings empirically and use them to test the model on a new set of test portfolios.

individual asset or the common component σ_v from order flow.¹² Analogously to Kyle (1985), but for second moments, we make this inference imperfect by introducing uncertainty about the amount of liquidity-driven trading. Specifically, we assume that $\sigma_{z,i}$ also has a common component, σ_z , and that neither is directly observable:

$$\sigma_{z,i}^2 = \sigma_z^2 + \varsigma_{z,i}^2. \quad (6)$$

By observing the cross section of order flow, liquidity providers can infer the combination $k\phi^2\sigma_v^2 + \sigma_z^2$ of the common components, but cannot perfectly isolate the amount of private information σ_v^2 from the amount of liquidity demand σ_z^2 . Thus, liquidity providers remain uncertain about the amount of private information in the market.

We keep the the environment conditionally Gaussian—and the learning tractable—by allowing liquidity providers to observe the variance of each asset’s order flow

$$\sigma_{x,i}^2 = \phi^2\sigma_{v,i}^2 + \sigma_{z,i}^2. \quad (7)$$

They could achieve this for instance by observing order flow across a large number of similar assets. This simplifies their learning problem without changing the underlying mechanism and allows us to obtain closed-form expressions.

We assume that news arrives at the intermediate date τ . The news leads liquidity providers to update their expected amounts of private information in each asset, $E_t^Q [\sigma_{v,i}^2]$. To focus on the impact of this shock, we assume there are no other changes in asset demand on date τ , hence prices adjust for informational reasons only.

The following proposition solves for prices as a function of the date-0 order flow x_i :

Proposition 1. *The price of asset i on date $t \in \{0, \tau\}$ is given by*

$$p_{i,t} = \bar{v}_i + \phi \frac{E_t^Q [\sigma_{v,i}^2]}{\sigma_{x,i}^2} x_i. \quad (8)$$

All proofs are in Internet Appendix IA.3. As in Kyle (1985), prices are sensitive to order flow x_i because it contains information about final payoffs due to trading by informed

¹²Empirically, Collin-Dufresne and Fos (2015) find that standard measures of illiquidity do not reveal trading by activists even though this trading has a large effect on prices once it becomes public. This shows that liquidity providers face significant uncertainty about the amount of private information in the market.

investors. However, unlike in Kyle (1985), liquidity providers do not know exactly how much information x_i contains because this depends on the amount of private information $\sigma_{v,i}^2$, which they do not observe. As a result, they optimally set the sensitivity of prices based on their expectation of this quantity, $E_t^Q[\sigma_{v,i}^2]$. Note that the expectation is risk-adjusted because liquidity providers price payoffs using the aggregate SDF.

Liquidity providers absorb the order flow of other investors, so their position is $-x_i$ in each asset. Proposition 1 shows that we can characterize this position using the change in an asset's price on date 0, $\Delta p_{i,0} \equiv p_{i,0} - \bar{p}_i$:

Lemma 1. *The position of liquidity providers in asset i , $-x_i$, is proportional to the date-0 change in the price of the asset:*

$$-x_i = -\frac{1}{\phi} \left(\frac{\sigma_{x,i}^2}{E_0^Q[\sigma_{v,i}^2]} \right) \Delta p_{i,0}. \quad (9)$$

Hence, liquidity providers hold a portfolio of reversals: they take long positions in assets whose price has declined and short positions in assets whose price has increased.

Lemma 1 shows that liquidity providers trade reversals, hence we will use reversals to test the model's predictions.¹³ The next step is to characterize their returns. As Proposition 1 shows, the price of an asset is exposed to shocks to its conditional volatility. The sign and magnitude of this exposure are a function of the asset's date-0 order flow x_i :

Lemma 2. *Let $\Delta p_{i,\tau} = p_{i,\tau} - p_{i,0}$ be the change in asset i 's price between date 0 and τ . Then*

$$\Delta p_{i,\tau} = \phi \frac{x_i}{\sigma_{x,i}^2} \left(E_\tau^Q[\sigma_{v,i}^2] - E_0^Q[\sigma_{v,i}^2] \right). \quad (10)$$

Thus, an asset's exposure to volatility shocks is proportional to its date-0 order flow x_i .

Lemma 2 shows that a positive shock to volatility increases the price of an asset if it had positive date-0 order flow (investors were buying) and decreases it if it had negative date-0 order flow (investors were selling). An increase in volatility reveals that there was more private information about the asset than liquidity providers expected. More private

¹³If there is also public news on date 0, then reversals become a noisy proxy for liquidity providers' portfolios. We allow for this in Internet Appendix IA.2 and show that the model's main predictions carry through.

information implies that prices were not sensitive enough to order flow. As a result, prices adjust further in the direction of that order flow.

Volatility shocks in our model push prices in the direction of past order flow because they reveal how much more information it contained. It is worth noting that not all volatility shocks have this effect. Some reflect information that is never in private hands before it becomes public. They can be modeled by adding a component to asset payoffs in Eq. (1) that is not known to anyone before date 1. The volatility of this component would not be correlated with the amount of information contained in order flow, and would therefore not affect prices in a way that is correlated with liquidity providers' positions. Thus, for volatility shocks to affect liquidity providers, it must be the case that some of the information they reflect was known to someone ahead of time. This is the case for information that is born private before it becomes public.¹⁴

Since order flow is positively related to the initial price change $\Delta p_{i,0}$, Lemma 2 shows that positive volatility shocks induce price *continuation*. And since liquidity providers hold reversals, continuation causes them to incur losses. Liquidity providers thus have a negative exposure to volatility shocks:

Lemma 3. *The change in the value of liquidity providers' position in asset i from date 0 to τ is*

$$-\Delta p_{i,\tau} x_i = -\phi \frac{x_i^2}{\sigma_{x,i}^2} \left(E_\tau^Q \left[\sigma_{v,i}^2 \right] - E_0^Q \left[\sigma_{v,i}^2 \right] \right). \quad (11)$$

Thus, liquidity providers are short a portfolio of variance swaps: they incur losses if there is an increase in $E^Q \left[\sigma_{v,i}^2 \right]$ (asset i 's variance swap rate), and earn a profit if it falls.

Lemmas 1 and 3 show that we can think of liquidity provision in terms of two trading

¹⁴This does not mean that investors trade on insider information. It is enough that some investors are earlier than others in identifying shocks that will drive an asset's price. There are many examples of this from the recent Covid crisis. One is Chris Hansen of Valliant Capital: "Chris Hansen saw it coming. Mr. Hansen... had an early conviction the novel coronavirus would wreak havoc on the global economy.... In February, Valiant started placing bets against, or shorting, levered companies it viewed as likely to be hurt from an economic slowdown caused by the virus. As the market plummeted, Valiant posted gains on those bets, including some on cruise lines, international airlines and travel companies" (*The Wall Street Journal*, April 2, 2020b). Another example is Jim Davis of Woodson Capital Management: "Managers who bet certain stocks would rise and others would fall had their best year in a decade.... For Woodson, which gained 15% in March when the S&P 500 lost 12.5%, existing bets that physical stores would suffer while e-commerce thrived helped. So did a stake in fitness company Peloton Interactive Inc., which surged more than 300% this year through November" (*The Wall Street Journal*, December 25, 2020a). Our model captures investors who, like Hansen and Davis, "saw it coming" and bought and sold assets accordingly.

strategies: reversals and selling variance swaps.¹⁵ These strategies seem unrelated, as reversals are a bet against private information and selling variance swaps is a bet against volatility. Yet they are tightly connected in our model since volatility reflects the flow of information, which includes the private information liquidity providers bet against. Thus, liquidity providers have a built-in negative volatility risk exposure.

Since liquidity providers risk-adjust returns using the aggregate SDF (they are fully diversified), any premium they demand can only be due to undiversifiable risk. Yet asset payoffs v_i are uncorrelated, hence one might think that there is no such risk here. If that were the case, liquidity providers would compete the cost of liquidity provision down to zero. However, it is not the case. Because liquidity providers are exposed to assets' idiosyncratic volatilities, and idiosyncratic volatilities load on the common component σ_v^2 , liquidity providers are exposed to this common component. It thus becomes a source of undiversifiable risk in their portfolios.

As Figure 2 shows, the common component in idiosyncratic volatilities, including the amount due to private information, is highly correlated with the volatility of the market portfolio. We capture this by assuming that σ_v^2 covaries with market volatility σ_m^2 .¹⁶

$$\sigma_v^2 = \sigma_m^2 + \varepsilon_v, \quad (12)$$

where ε_v is orthogonal to σ_m^2 and all other shocks. As with σ_v^2 , market participants do not know σ_m^2 and must form risk-adjusted expectations about it, $E_\tau^Q[\sigma_m^2]$, at each point in time. Combining Lemma 3 and Eq. (12), liquidity providers have a negative exposure to market volatility risk. We characterize it with the following proposition:

Proposition 2. *The beta of liquidity providers' position in asset i with respect to shocks to risk-*

¹⁵A variance swap on asset i pays out the asset's realized variance, $\sigma_{v,i}^2$. The variance swap rate on date t is the risk-adjusted expectation of this payoff, $E_t^Q[\sigma_{v,i}^2]$.

¹⁶We can write the aggregate market portfolio's payoff as: $p_{m,1} = \bar{v}_m + v_m$ where $v_m \sim N(0, \sigma_m)$. It is not important for us to separate public and private information about the market, so v_m includes both.

adjusted expected market variance $E_\tau^Q [\sigma_m^2]$ (i.e., their market volatility beta) is

$$\beta_{i,\sigma_m} = \frac{\text{Cov} \left(-\Delta p_{i,\tau} x_i, E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)}{\text{Var} \left(E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)} = -\frac{\phi k x_i^2}{\sigma_{x,i}^2} \quad (13)$$

$$= -\frac{k}{\phi} \left(\frac{\sigma_{x,i}}{E_0^Q [\sigma_{v,i}^2]} \Delta p_{i,0} \right)^2 < 0. \quad (14)$$

Thus, all of liquidity providers' positions, long ($-x_i > 0$) and short ($-x_i < 0$), have negative market volatility betas. When market volatility rises, liquidity provider lose on both sides: their long positions fall and their short positions rise.

Proposition 2 shows that liquidity providers are negatively exposed to market volatility shocks. A large literature in volatility and option pricing documents that market volatility shocks carry a very large negative price of risk, i.e. periods of high volatility are priced as bad times by the aggregate SDF.¹⁷ Formally, this means that the risk-adjusted expected market volatility is higher than the objective (i.e., statistical) expected volatility: $E_t^Q [\sigma_m^2] > E_t [\sigma_m^2]$. Equivalently, the price of a market variance swap, $E_t^Q [\sigma_m^2]$, is higher than its expected payoff, $E_t [\sigma_m^2]$. The empirical counterpart of $E_t^Q [\sigma_m^2]$ is the squared VIX index, which is constructed from options to replicate the payoff of a variance swap on the S&P 500. Proposition 2 thus predicts that liquidity providers have negative betas with respect to VIX squared.

Since the price of the variance swap must converge to its payoff as $t \rightarrow 1$, it is expected to drift down under objective expectations:

$$E_0 \left[E_t^Q [\sigma_m^2] \right] < E_0^Q [\sigma_m^2] \quad \text{for } t > 0. \quad (15)$$

This downward drift is the insurance premium the buyer of the variance swap pays to hedge positive shocks to market volatility. In the literature this premium is called the variance risk premium. Since liquidity providers in our model are effectively short variance swaps, the premium they charge—the liquidity premium—reflects the variance risk premium.

¹⁷A simple SDF that prices market variance shocks is: $\Lambda_1 = \exp(-\gamma \sigma_m^2) / E_0 [\exp(-\gamma \sigma_m^2)]$, where $\gamma < 0$ controls the price of market volatility risk and Λ_1 is normalized so that $E_0 [\Lambda_1] = 1$. Since $\gamma < 0$, high-variance states are high-marginal utility states, so investors are willing to pay a premium to hedge them.

The liquidity premium for each asset is given by its market volatility beta multiplied by the variance risk premium. Summing up across assets, the expected payoff of liquidity providers is given by the market volatility beta of their portfolios multiplied by the variance risk premium:

Proposition 3. *The expected payoff of liquidity providers' portfolios from date 0 to date τ is*

$$E_0 \left[\sum_{i=1}^N -\Delta p_{i,\tau} x_i \right] = \left(\sum_{i=1}^N \beta_{i,\sigma_m} \right) \left(E_0 \left[E_\tau^Q \left[\sigma_m^2 \right] \right] - E_0^Q \left[\sigma_m^2 \right] \right) > 0. \quad (16)$$

Thus, the liquidity premium is positive and proportional to the variance risk premium.

Proposition 3 shows that the expected return to liquidity provision—the liquidity premium—is positive and determined by the market volatility risk of liquidity providers' portfolios, i.e. their market volatility beta. Liquidity providers do not earn positive returns because they are constrained or under-diversified as commonly assumed in the literature, but because they are exposed to systematic volatility risk. Of course, liquidity providers could use other markets, such as the variance swap market, to hedge out this volatility risk, but then they would have to hand over the premium they are earning for liquidity provision to the seller of the variance swap. Thus, a liquidity premium emerges here in a perfectly integrated market.

2.1 Market segmentation and inventory costs

We now extend the model to incorporate market segmentation in the form of inventory costs, which are common in the literature. As in Nagel (2012), inventory costs lead liquidity providers to demand additional compensation for holding inventory in a given asset. This compensation is over and above the risk premium on the asset due to its covariance with the aggregate SDF. Thus, in contrast to our main model, the presence of inventory costs means that liquidity providers are partly segmented from the rest of the economy. This could be due to asymmetric information, moral hazard, imperfect competition, or other frictions.

We assume that inventory costs are quadratic so that the marginal cost of holding $-x_i$

units of asset i from date 0 to date t is $-t\gamma_{i,t}x_i$. The pricing condition (3) becomes

$$p_{i,t} = E_t^Q[p_{i,1}] + (1-t)\gamma_{i,t}x_i \text{ for } t \in \{0, \tau\}. \quad (17)$$

Asset i trades at a discount to its risk-adjusted expected payoff if liquidity providers are long ($-x_i > 0$), and at a premium if they are short ($-x_i < 0$). The discount or premium declines over time as the remaining holding period $(1-t)$ shrinks. As it does so, it causes a reversal in the price of the assets, which compensates liquidity providers for their inventory costs.

Inventory costs can vary both across assets and over time. Nagel (2012) focuses on time variation driven by VIX. We can represent it by setting $\gamma_{i,t} = \gamma E_t^Q[\sigma_m^2]$. Collin-Dufresne and Daniel (2014) focus on specialized liquidity providers who are under-diversified and hence bear idiosyncratic risk. We can capture it with $\gamma_{i,t} = \gamma E_t^Q[\sigma_{v,i}^2]$. Our main model is nested by $\gamma_{i,t} = 0$.

The solution to the extended model follows that of the main model. We focus on the key results that will let us tell them apart empirically. The first result extends Lemma 3 by characterizing liquidity providers' exposure to volatility shocks when we further allow for inventory cost shocks:

Lemma 3'. *Let $\Delta\gamma_{i,\tau} = \gamma_{i,\tau} - \gamma_{i,0}$ be the shock to asset i 's inventory cost at date τ . Then the change in the value of liquidity providers' position in asset i from date 0 to τ is*

$$-\Delta p_{i,\tau}x_i = -x_i^2 \left[\frac{\phi}{\sigma_{x,i}^2} \left(E_\tau^Q[\sigma_{v,i}^2] - E_0^Q[\sigma_{v,i}^2] \right) - \tau\gamma_{i,0} + (1-\tau)\Delta\gamma_{i,\tau} \right], \quad (18)$$

Lemma 3' shows that inventory costs affect liquidity providers' returns in two ways. First, the value of their position in asset i drifts up over time at the rate $\gamma_{i,0}x_i^2$ to compensate for the cost of carrying the inventory. As expected, this drift is over and above the compensation for volatility risk.

The second effect is through shocks to the inventory cost at date τ , $\Delta\gamma_{i,\tau}$. A positive inventory cost shock reduces the value of liquidity providers' position on impact. This allows the position to appreciate going forward, and in doing so recoup the higher inventory cost. By date 1 ($\tau \rightarrow 1$), all inventory costs have been paid out and the position recovers to its fundamental value. Hence, inventory cost shocks have a transitory impact

that dissipates by the end of liquidity providers' holding period.

Inventory cost shocks can affect liquidity providers' market volatility betas if they are correlated with market volatility shocks. However, Lemma 3' shows that their impact would be transitory: when market volatility rises, higher inventory costs would push liquidity providers' positions below their fundamental values, but they would then recover as the holding period elapses. We formalize this result with the following proposition, which extends Proposition 2 of the main model:

Proposition 2'. *The market volatility beta of liquidity providers' position in asset i from date 0 to date τ is given by*

$$\beta_{i,\sigma_m}^{0 \rightarrow \tau} = -x_i^2 k \left[\frac{\phi}{\sigma_{x,i}^2} + (1 - \tau) \beta_{\gamma_i, \sigma_m} \right], \quad (19)$$

where $\beta_{\gamma_i, \sigma_m}$ is the market volatility beta of asset i 's inventory cost $\gamma_{i,t}$. Thus, the contribution of inventory cost shocks to liquidity providers' betas decreases with τ and goes to 0 as $\tau \rightarrow 1$.

Proposition 2' shows that if inventory costs and market volatility covary ($\beta_{\gamma_i, \sigma_m} > 0$), then market volatility shocks have a bigger effect on liquidity providers' positions on impact, but their effect is unchanged over the full holding period. Thus, liquidity providers' market volatility betas become more negative, but the additional impact is transitory. By contrast, the impact of market volatility shocks under our main mechanism (first term in Eq. (19)) is permanent. This is because private information risk affects the fundamental value of liquidity providers' positions, not just their cost of carry during the holding period. Proposition 2' thus allows us to test the importance of inventory costs versus private information risk by looking at the persistence of the impact of volatility shocks on liquidity providers' portfolios.

We can further test the importance of the two mechanisms by contrasting their predictions for liquidity providers' expected return, i.e. the liquidity premium. The following proposition, which extends Proposition 3, shows how:

Proposition 3'. *The expected payoff of liquidity providers' portfolio from date 0 to τ is:*

$$E_0 \left[\sum_{i=1}^N -\Delta p_{i,\tau} x_i \right] = \left(\sum_{i=1}^N \beta_{i,\sigma_m}^{0 \rightarrow \tau} \right) \left(E_0 \left[E_\tau^Q \left[\sigma_m^2 \right] \right] - E_0^Q \left[\sigma_m^2 \right] \right) + \tau \gamma_{i,0}. \quad (20)$$

The liquidity premium exceeds the premium for the variance risk of liquidity providers' portfolio by the inventory cost.

Proposition 3' shows that inventory costs raise the liquidity premium over and above the risk premium liquidity providers earn for their volatility risk exposure. The reason is that liquidity providers are segmented from other investors, so the average return they earn exceeds the fair premium for covariance with the aggregate SDF. In contrast, in our main model markets are fully integrated and the two are equal. We can therefore test our main model by seeing if the compensation for liquidity provision is equal to the associated variance risk premium.

2.2 Summary of predictions

We briefly summarize the key predictions of our model. As Lemma 1 shows, we can test them using reversals to proxy for liquidity providers' portfolios. We do so in Section 4.

Prediction 1. *Liquidity providers hold a portfolio of reversals: they buy stocks whose price has fallen and sell stocks whose price has risen. The magnitude of their position is proportional to the stock's price change normalized by its volatility and weighted by its order flow variance.*

Prediction 1 follows from Lemma 1. Since we do not observe liquidity providers' portfolios directly, we use this prediction to construct reversal portfolios as an empirical proxy.¹⁸ Lemma 1 further tells us to weight the reversal portfolios by the variance of order flow ($\sigma_{x,i}^2$). Since we do not observe it, we weight by dollar volume which, like order flow variance, is unsigned. We test the model on the resulting portfolios. We describe the portfolio construction in detail in Section 3.

Prediction 2. *Reversal portfolios have negative market volatility betas. The more extreme reversal portfolios have more negative betas.*

Prediction 2 follows from Proposition 2, which shows that liquidity providers have negative market volatility betas. The betas become larger for more extreme reversal portfolios (made up of stocks with larger price changes) because they are associated with greater liquidity provision. We measure the betas by regressing reversal returns on changes in the squared VIX index as implied by Proposition 2.

¹⁸The extended model in Internet Appendix IA.2 shows that reversals become a noisy but still unbiased proxy for liquidity providers' portfolios in the presence of public news.

Prediction 3. *The variance risk premium predicts reversal returns in the time series. Across reversal portfolios, the predictive coefficient is proportional to the market volatility beta.*

Prediction 3 follows from Propositions 3. The variance risk premium from 0 to τ is given by $E_0 \left[E_\tau^Q [\sigma_m^2] \right] - E_0^Q [\sigma_m^2]$. Thus, by Eq. (16), a reversal portfolio's predictive coefficient is given by its market volatility beta.

Prediction 4. *The impact of market volatility shocks on reversal returns is permanent over liquidity providers' holding period. If instead it is transitory, it reflects the impact of inventory costs or other market frictions.*

Prediction 4 follows from Proposition 2'. In our main model, an increase in market volatility reflects an increase in the amount of private information, which negatively affects the fundamental value of liquidity providers' holdings. Hence, volatility shocks have a permanent impact on reversal returns. In contrast, under the inventory-cost mechanism volatility shocks only affect assets while held by liquidity providers, with no impact on fundamental values. Therefore, Prediction 4 allows us to test the mechanism of our model against the inventory-cost mechanism studied in the literature.

Prediction 5. *The expected return of reversals equals their market volatility beta times the variance risk premium. Any expected return above this amount reflects the impact of inventory costs or other market frictions.*

Prediction 5 follows from Proposition 3'. It allows us to separate the part of the liquidity premium that is compensation for volatility risk from the part that is due to market frictions. Since volatility risk is traded directly in option markets, we can test our main model against the inventory-cost mechanism by comparing the expected returns of reversals to the product of their market volatility betas and the variance risk premium from options.

3 Data and summary statistics

In this section we describe our sample and how we construct the reversal portfolios. Additional details are in Appendix IA.1.

Sample selection: Our main data is from CRSP. We restrict the sample to ordinary common shares and exclude penny stocks and micro-caps. The sample is daily from April 9, 2001 to May 31, 2020. The start date corresponds to “decimalization,” the transition from fractional to decimal pricing on the New York Stock Exchange and NASDAQ. As Bessembinder (2003) shows, 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 rather than risk exposures, hence we exclude this period from the analysis.

Portfolio formation: We construct a set of reversal portfolios following Prediction 1. Each day, we first sort stocks into quintiles by market capitalization. We form ten decile portfolios within each size quintile by sorting stocks by their beta-adjusted standardized return. We weight the portfolios by dollar volume as a proxy for order flow variance (see Prediction 1). We follow Nagel (2012) and hold each portfolio for five trading days. This horizon is also consistent with the evidence in Hendershott and Seasholes (2007) that NYSE specialists (a prime example of liquidity providers) earn most of their returns within five days of entering in a position.

Reversal strategies: We form long-short reversal strategies that buy the low-return portfolio and sell the high-return portfolio within each size quintile. In particular, the Lo–Hi reversal strategy buys the first normalized return decile (“Lo”) and sells the tenth normalized return decile (“Hi”) in a given size quintile.

Aggregate factors: The CBOE VIX index is a model-free measure of the implied volatility of the S&P 500 at a 30-day horizon (as proposed by Britten-Jones and Neuberger, 2000). The squared VIX therefore maps closely to the risk-adjusted expected market variance in our model, $E_t^Q [\sigma_m^2]$. We use data on S&P 500 index options from OptionMetrics to calculate the return to holding the VIX basket, which we refer to as the VIX return (this data ends on December 31, 2019). We use the VIX return to restrict the price of volatility risk to the one that prevails in option markets as implied by Prediction 5.

3.1 Summary statistics

Table 1 presents summary statistics for the reversal strategies. Each panel contains a five-by-five table focusing on a given characteristic. Each row of the table represents a size

quintile (Small, 2, 3, 4, and Big) and each column represents a long-short portfolio formed across the return deciles (Lo–Hi, 2–9, 3–8, 4–7, and 5–6). Each long-short reversal strategy contains 125 stocks on average.

Panel A of Table 1 looks at market capitalizations. The average stock in the largest quintile is worth \$81.75 billion, almost three orders of magnitude larger than the smallest quintile and two orders larger than the middle one. The largest stocks account for 95.42% 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 0.14%.

Panel B of Table 1 looks at idiosyncratic volatility, which we used to normalize the returns of individual stocks before sorting them into portfolios. Here we report its average value within each reversal strategy. Given our normalization, idiosyncratic volatility is relatively flat across return deciles. At the same time, it varies significantly across size quintiles: the largest stocks have idiosyncratic volatility between 1.71% and 2.01% while for the smallest stocks it is between 3.74% and 5.96%. This reflects the well-known fact that volatility is decreasing in size (Campbell et al., 2001).

Panel C of Table 1 shows the illiquidity measure of Amihud (2002), which is calculated as the absolute value of a stock’s return divided by its dollar volume (multiplied by 10^6 for readability). This illiquidity is measured on the portfolio formation date and averaged across all stocks in the portfolio and over time for the portfolio itself. As expected, illiquidity is strongly decreasing in size: the largest stocks have illiquidity that is four orders of magnitude smaller than for the smallest stocks.

Panel D of Table 1 shows sorting-day returns, i.e., the sorting-day return of the long leg of the portfolio minus that of the short leg (without normalizing). The sorting-day returns are negative by construction. Since small stocks are more volatile than large stocks, their sorting-day returns are substantially larger in magnitude. The average sorting-day return of the Lo–Hi strategy is -13.55% for the smallest stocks versus -5.20% for the largest ones. As expected, the sorting-day returns decline in magnitude as we move toward the inner deciles, approaching zero for the 5–6 strategies. The Lo–Hi strategy thus carries the strongest reversal signal and captures the most intensive liquidity provision.

Panels E and F of Table 1 look at turnover. Panel E shows average turnover over the 60 days prior to portfolio formation, while Panel F shows turnover on the day of portfolio formation. From Panel E, average turnover is largely flat across both return deciles and

size quintiles. By contrast, Panel F shows that sorting-day turnover is significantly higher (by about 40%) for the Lo–Hi strategy than the 5–6 strategy, with a monotonically decreasing pattern in between. This result is consistent with the model, where high order flow induces large price changes as liquidity providers filter out the information it contains. Volume, which unlike order flow is unsigned, is therefore increasing in the magnitude of price changes. Panels E and F thus validate the use of volume as an unsigned proxy for order flow. By combining it with returns, which are signed, our reversal strategies capture the portfolios of liquidity providers.

Table 2 provides summary statistics on the VIX return and related measures. From Panel A, the VIX return averages -1.53% per day. This number is in line with estimates of the variance premium from the literature (e.g. Carr and Wu, 2008; Bollerslev, Tauchen and Zhou, 2009; Drechsler and Yaron, 2010). It reflects the very large price investors are willing to pay to hedge variance risk.

Unlike the VIX return, changes in VIX and VIX-squared have a mean of zero. This is because the VIX basket is rebalanced each day to keep its maturity constant. Thus, the change in VIX is not a return and does not carry the variance premium. This is why we compute the VIX return, which is the percentage change in the price of a fixed VIX basket.

The VIX return is also significantly more volatile than changes in VIX and VIX-squared (its standard deviation is 17.86% versus 1.82% and 1.53% , respectively). The VIX return is also right-skewed, as seen from the difference between its mean and median (-1.53% and -5.08%) and 1^{st} and 99^{th} percentiles (-25.33% and 71.29%). The reason is that the VIX return is a percentage change while the changes in VIX and VIX-squared are simple first differences. Consistent with this, the percentage change in VIX-squared (bottom row) has a standard deviation of 16.28% , which is similar to that of the VIX return.

Panel B of Table 2 reports results from regressions of the VIX return on VIX changes, VIX-squared changes, and their percentage counterparts. These regressions allow us to calculate the VIX return premium per unit of beta, i.e. the price of variance risk in option markets. We will use this price of risk in Section 4.6 to test whether it can explain the returns to liquidity provision in stocks.

From columns (1) and (2), the VIX return has a beta of 7.939 with respect to VIX changes and 7.404 with respect to VIX-squared changes. Thus, the implied price of risk is -0.207% ($= -1.53/7.404$) per unit of VIX-squared beta per day. This translates to -1.03%

at the five-day horizon of our reversal strategies.

The VIX return is highly correlated with VIX changes (75%) and VIX-squared changes (53%), as reflected in the R^2 in Panel B columns (1) and (2). It is even more correlated with their percentage change counterparts (86% and 89%, respectively). The explanation for the difference is again the fact that the VIX return is a percentage change. In addition, the VIX return incorporates the day's realized variance (the strategy's "dividend"), while VIX and VIX-squared changes reflect only changes in expected future variance (again because the basket is rebalanced). Overall, the table shows that innovations in VIX and VIX-squared capture most of the variation in the VIX return.

4 Empirical Results

4.1 Average returns

Table 3 shows the post-formation returns of the reversal strategies at the five-day horizon. From Panel A, the Lo–Hi strategy among the largest stocks delivers a five-day return of 20 bps. From Panel B, this return is highly significant, with a t statistic of 4.47 (using Newey-West standard errors with five lags to account for the overlap in the returns). This return is about 10% per year, which is economically large. From Panel C, the portfolio's standard deviation is 2.84% and from panel D, the annual Sharpe ratio is 0.5, slightly higher than the market Sharpe ratio in our sample (0.41). Consistent with the model, the reversal strategy returns decline as we move from the Lo–Hi strategy toward the inner deciles, reaching near zero for the 5–6 strategy. Intuitively, the inner deciles reflect less intensive liquidity provision and therefore earn lower premiums.

The returns increase as we move from large stocks toward small stocks (Avramov, Chordia and Goyal, 2006, find the same result). For the smallest stocks, the Lo–Hi strategy delivers a five-day return of 78 bps, which is highly significant. The strategy's volatility is higher, 6.18%, but nevertheless the Sharpe ratio is also higher, 0.9.¹⁹ Thus, providing liquidity for very small stocks earns higher returns.

Panel E of Table 3 reports the CAPM alphas of the reversal strategies. They are ob-

¹⁹The difference is partly due to bid-ask bounce (Roll, 1984), which can be viewed as part of the return to liquidity provision. It does not impact large stocks whose bid-ask spreads are very small.

tained from the time-series regressions:

$$R_{t,t+5}^p = \alpha^p + \beta^p R_{t,t+5}^M + \epsilon_{t,t+5}^p, \quad (21)$$

where $R_{t,t+5}^p$ is the excess return of portfolio p from t to $t + 5$ and $R_{t,t+5}^M$ is the cumulative excess market return over the same period. Panel F shows the associated t statistics.

The results show that the CAPM cannot explain the returns of the reversal strategies. In all cases, the CAPM alphas are close to the raw average returns and as significant. The Lo–Hi strategy for the largest stocks has an alpha of 17 bps, only slightly lower than its 20-bps average return. The associated t -statistic is 3.90, strongly rejecting the 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 a robust reversal premium and is consistent with liquidity being expensive on average.

4.2 Volatility betas

We now test Prediction 2, which predicts that reversal strategies are exposed to volatility risk. We run the following regressions:

$$R_{t,t+5}^p = \alpha^p + \beta^{p,VIX} \Delta VIX_{t,t+5}^2 + \beta^{p,M} R_{t,t+5}^M + \epsilon_{t,t+5}^p. \quad (22)$$

These regressions expand (21) to include the change in VIX-squared included alongside the market return. This is the appropriate factor according to our model, the empirical counterpart to the change in the risk-adjusted market variance in Proposition 2.

Table 4 presents the results. Panel A reports the betas from a specification with only VIX-squared changes. Focusing on the Lo–Hi strategy first, the VIX betas are uniformly negative and highly statistically significant, consistent with Prediction 2. They are also very similar across size quintiles, both in terms of magnitude and statistical significance. Also consistent with Prediction 2, the betas decline steadily as we move toward the inner decile strategies. They are still mostly significant for the 2–9 strategies but only about half as large; they are very close to zero and insignificant for the 5–6 strategies.

The estimated beta for the large-stock Lo–Hi strategy is -0.20 . This means that the strategy loses 20 bps, an amount equal to its average premium, when VIX-squared rises

by one point. From Table 2, the standard deviation of VIX-squared changes is 1.53 points per day, which works out to 3.42 points per five days. Hence, a one-standard deviation increase in VIX-squared over the holding period wipes out over three times the strategy's average return. Thus, the volatility risk betas of the Lo-Hi reversal strategies are economically large.

Panel B of Table 4 adds the market return as a control. This addresses a potential concern that the negative VIX exposure reflects market risk instead of volatility risk. The table shows that this is not the case. The VIX-squared betas remain very similar to those in Panel A. The strong correlation between the market and VIX lowers the t statistics somewhat but most remain highly significant. The beta of the large-stock Lo-Hi strategy remains identical, -0.20 , and highly significant.

Panel C of Table 4 looks directly at the market betas from the bivariate regression in Panel B. The market betas of the Lo-Hi strategies are very small and insignificant. The market beta of the large-stock Lo-Hi strategy is almost exactly zero. This shows that the strategies are neutral with respect to market risk, which is unsurprising given their long-short construction.

Stocks in our model acquire a volatility beta endogenously when they enter liquidity providers' portfolios and become exposed to private information risk. We therefore want to make sure that our reversal portfolios are not simply picking up stocks that happen to have high volatility betas (or betas to some other factor that is correlated with volatility). If that was the case, then the stocks in the reversal portfolios would exhibit higher volatility betas even before the portfolios are formed.

Figure 3 shows that this is not the case. It plots the volatility betas of the reversal portfolios before and after portfolio formation. For the period before formation, we compute cumulative returns from up to ten days before the formation date to the formation date. For the period after formation we compute them from the formation date to up to ten days after. We then calculate the betas by regressing these cumulative returns on the change in VIX squared over the same period. Panel A focuses on the Lo-Hi strategy among top quintile of stocks, while Panel B includes the remaining quintiles. Both panels show that volatility betas are zero prior to portfolio formation. Thus, there is no sign of a missing factor among the stocks in the reversal portfolios. The betas become sharply negative as soon as the portfolios are formed. They remain negative and continue to build over the

first few days after portfolio formation. This is consistent with private information risk lasting several days. Overall, the results of Table 4 and Figure 3 show that the reversal strategies have large negative volatility betas in line with Prediction 2 of the model.

4.3 Return predictability with VIX

Volatility risk is strongly time-varying and increasing in volatility itself (Singleton, 2006; Broadie, Chernov and Johannes, 2007). Consistent with this, the variance risk premium is larger in more volatile times (Todorov, 2010; Andersen, Fusari and Todorov, 2015). Since our model predicts that reversal returns are explained by volatility risk, this time variation in the compensation for volatility risk should generate predictability in reversal returns. This is the idea underlying Prediction 3, which we test next.

Table 5 shows results from predictive regressions of the reversal strategy returns on the level of VIX squared as of the portfolio formation date. Panel A reports the predictive loadings (times 100 for legibility). Focusing first on the Lo–Hi strategies, the loadings are positive and large across all size quintiles, ranging from 5.70 to 9.81. From Panel B, they are all highly statistically significant. For the largest stocks, the coefficient is 9.09, hence a one-point increase in VIX squared predicts a 9-bps higher return over the next five days. This number is large relative to the strategy’s 17-bps average return. The predictive loadings decline steadily as we move toward the inner decile strategies. Thus, they display the same pattern as the volatility risk betas in Table 4. This is consistent with Prediction 3 of the model. Intuitively, portfolios that are more exposed to volatility risk have a risk premium that co-moves more strongly with the variance risk premium.

Panel C of Table 5 shows the regression R^2 . Focusing again on the Lo–Hi strategies, the R^2 is lowest for the smallest quintile (0.85%) and highest for the largest quintile (3.45%). The latter is extremely high given the strategy’s five-day horizon.²⁰ Thus, aggregate volatility risk has the highest explanatory power for the large-cap reversal strategies. Overall, the predictability results in Table 5 extend the main finding of Nagel (2012) to large-cap stocks and support the predictions of our model.²¹

²⁰Following Campbell and Thompson (2007), it is about six times the strategy’s squared five-day Sharpe ratio. Thus, an investor using VIX to time the large-cap Lo–Hi strategy would see a six-fold increase in expected returns relative to a buy-and-hold investor.

²¹In Internet Appendix IA.2 we replace VIX squared with an estimate of the variance risk premium, the difference between VIX squared and predicted variance of the S&P 500. We estimate the predicted

4.4 Persistence of the impact of volatility shocks

Prediction 4 of our model says that the impact of volatility shocks on liquidity providers' portfolios is permanent over their holding period. This is because volatility shocks reflect changes in the amount of private information about fundamental values. By contrast, under the inventory cost mechanism the impact of volatility shocks, if any, is transitory and fully dissipates by the end of the holding period. This is because inventory costs only affect assets during their time in liquidity providers' portfolios. The persistence of the impact of volatility shocks thus gives a sharp test of our theory, which we perform next.

We begin by looking at the persistence of the returns of the reversal portfolios themselves. Panel A of Figure 4 plots these returns at horizons up to ten days after portfolio formation. The returns build steadily and begin to level off after about five days. They are highest for the Lo–Hi strategy and decline toward the inner deciles.²² Panel B of Figure 4 shows that the predictive loadings on VIX squared follow the same pattern, both across horizons and across portfolios. Panels C and D find the same for the volatility betas, whether we control for the market return (Panel D) or not (Panel C). The fact that the returns, predictive loadings, and betas all show the same pattern over time and cross-sectionally supports the view that the returns to reversals reflect compensation for volatility risk.

Figure 5 tests Prediction 4 by plotting the impulse responses of the reversal strategy returns to a VIX-squared shock one day after portfolio formation:

$$R_{t,t+h}^p = \alpha_{p,h} + \delta_{VIX,h}^p \Delta VIX_{t,t+1}^2 + \epsilon_{t,t+h}^p, \quad \text{for } h = 1, 2, \dots, 20. \quad (23)$$

If the impact of VIX shocks is persistent, the response coefficients $\delta_{VIX,h}^p$ should be flat across horizon h . If it is transitory, they should shrink toward zero.

Panel A of Figure 5 looks at the large-cap Lo–Hi reversal strategy. A one-point increase in VIX squared on day one leads to a drop of 17 bps, about equal to the strategy's

variance by projecting the realized variance of the S&P 500 on its lag and the lagged VIX squared. We find that this estimated variance risk premium predicts reversal returns as well as VIX squared does. This is not surprising since the two measures are highly correlated, as is the case with most such estimates. For simplicity, and to avoid using an estimated regressor, we use VIX squared as our main predictor.

²²The steady pattern shows that the returns are not driven by bid-ask bounce, which can only affect the first day of the holding period. Instead, the figure shows a stable premium paid over time.

five-day return. The impact remains flat over time and settles at 19 bps after twenty trading days. The 95% confidence bands lie comfortably away from zero at horizons well beyond the holding periods associated with liquidity provision. This result is not consistent with the transitory effect predicted by the inventory cost mechanism. Instead, it points to a permanent impact on asset values, as predicted by our model. The results therefore support Prediction 4 of our model.

Panel B Figure 5 includes the other size quintiles. Quintiles three and four show the same pattern as the largest quintile, with a permanent effect over twenty days. This is not the case for the smallest two quintiles, however. Their returns tend to revert back toward zero after the initial drop following the shock. This pattern is consistent with the predicted impact of an increase in inventory costs, suggesting that inventory costs or other market frictions play a role in the case of the smallest stocks.

Table 6 formalizes the results of Figure 5. Panel A shows that after five days the Lo–Hi strategies in each size quintile have similar negative response coefficients to a VIX shock on day one. By day ten, however, the response coefficients of the smallest two quintiles have reverted to zero, while those of the three larger quintiles are unchanged through day twenty.

Overall, Figure 5 and Table 6 support the prediction of our model that volatility shocks have a persistent impact on liquidity providers’ portfolios. The only exception are the smallest stocks, where we find evidence of market segmentation. This helps to explain why these stocks have such abnormally high reversal returns in Table 3 despite their similar volatility betas in Table 4. The difference between small and large stocks is natural because small stocks are thinly traded and rely on specialized intermediaries for liquidity provision.²³ Among large stocks, where this is not a problem, we find significant and persistent volatility risk as predicted by our model.

4.5 Fama-MacBeth asset pricing tests

We now test Prediction 5 of the model, which says that if there is no market segmentation the expected returns of the reversal strategies should equal the product of their volatility

²³Hasbrouck (1988) finds that inventory costs play a role for low-volume stocks but not high-volume stocks. This can reconcile our results with Nagel (2012), which places a large weight on small stocks by focusing on an equal-weighted reversal portfolio. Consistent with our results for small stocks, Nagel (2012) interprets the predictability of this portfolio by VIX through the lens of an inventory cost model.

betas and the price of variance risk. In this section we treat the price of variance risk as a free parameter and recover it from the cross section of reversal strategies. In the next section we sharpen the analysis by testing whether this recovered price is consistent with the price that prevails in option markets so that the markets for liquidity and volatility are integrated.

We run standard two-stage Fama-MacBeth regressions to test whether the volatility betas can explain the average returns of the reversal strategies. The first stage estimates the betas as in Section 4.2 (Eq. (22)). The betas then enter the second-stage regression:

$$R_{t,t+5}^p = \lambda_{0,t} + \lambda_{t,VIX}\beta_{VIX}^p + \lambda_{t,M}\beta_M^p + \epsilon_{t,t+5}^p. \quad (24)$$

This regression provides time series of the estimated factor premia $\lambda_{t,VIX}$ and $\lambda_{t,M}$. Following the Fama-MacBeth methodology, we use this time series variation to calculate the standard errors of the factor premia, again with a Newey-West correction for the overlap in the returns. To assess overall performance, we report root-mean-squared pricing errors for all portfolios and for different combinations of portfolios. We also report the associated p -values, which test whether the pricing errors are jointly equal to zero.

The results of the Fama-Macbeth regressions are presented in Tables 7 and 8. Table 7 shows the estimated factor premia and pricing errors. The first column does not include any factors, hence it provides statistics for the raw returns as a benchmark. From the constant, the average portfolio return is 20.5 bps. The overall r.m.s.e. is similar, 24 bps. The r.m.s.e. of the long-short reversal strategies is also similar, 22.8 bps. Excluding the smallest two quintiles or value-weighting reduces it to 12.3 and 12.4 bps, respectively, because small stocks have larger reversal returns. These returns are still highly significant as seen from the low p -values. They are also sizable given their five-day horizon and equal to 6.2% annualized.

The last row of Table 7 gives the pricing error of an overall liquidity provider portfolio, which weighs the reversal strategies by the product of their dollar volume and normalized return, as implied by Lemma 1. The liquidity provider portfolio puts more weight on the Lo-Hi reversal strategies because they capture more intensive liquidity provision. Its return is therefore somewhat higher than the value-weighted one, 16.8 bps over the five-day horizon, or 8.4% annualized.

Column 2 of Table 7 reports the specification with the market return as the sole factor (i.e., the CAPM). The market premium is positive, 32.5 bps, which works out to 16.25% per year. This number is more than double the equity premium in the sample, hence it is somewhat implausible. The high premium is required because the portfolios exhibit very small differences in market betas. It achieves a cross-sectional R^2 of 15.8% and r.m.s.e. of 11.2 bps, about half the raw value in column 1. However, the CAPM does very little to price the long-short-reversal strategies: their r.m.s.e. is 20.1 bps, only slightly lower than the raw value. When we exclude small stocks or value-weight, the pricing error is 9.4 bps, about one quarter smaller than the raw value. The pricing error of the liquidity provider portfolio is 12.7 bps, also one quarter smaller. As in column 1, all of the p -values are zero to within two decimal points, hence the CAPM is rejected statistically.

Overall, the CAPM can explain about half of the level of returns among the full set of portfolios but requires an implausibly high risk premium to do so. Even then, it is unable to explain most of the differences in returns between the portfolios, which are captured by the long-short reversal strategies. Since the reversal strategies proxy for the returns to liquidity provision, this means that the CAPM is unable to explain the liquidity premium in the cross section of stocks even with an implausible high risk premium.

Column 3 uses the change in VIX squared as the sole factor. The estimated price of risk is -0.570 and highly significant. In contrast to the CAPM, this magnitude is if anything on the low side, as we will see in the next section. The cross-sectional R^2 increases to 26%, the all-portfolio r.m.s.e drops slightly to 10.7 bps, and that of the reversal strategies to 18.3 bps. The impact is larger when we exclude small stocks or value-weight: the pricing errors drop to 6.5 and 6.2 bps, respectively, which is about half the corresponding raw values. The pricing error of the liquidity provider portfolio also drops in half, to 8.4 bps. The p -values remain close to zero, hence this model is also statistically rejected.

The one-factor volatility risk model is thus better able to price the test assets than the CAPM, but it does not do so fully. The reason is that there is a tension between the overall level of returns, i.e. the equity premium, and the differences between them, the liquidity premium. The equity premium requires a relatively low price of volatility risk, while the liquidity premium requires a larger one. The univariate model thus settles on a medium-sized premium, which limits its ability to fully capture the cross section. Note however that our model does not predict that volatility risk should explain the equity premium. It

therefore makes sense to use volatility risk in conjunction with market risk in a two-factor asset pricing model.

The two-factor model is estimated in column 4. The market premium drops to 15.2 bps, which is very close to the equity premium in our sample. This helps the model to price the level of returns among the portfolios. The price of variance risk doubles to -1.079% and is highly significant. This allows the model to do a much better job in explaining the cross section. The cross-sectional R^2 increases to 37.9%, the overall r.m.s.e. drops to 9.7 bps, and the r.m.s.e. of the reversal strategies drops to 15.8 bps. These pricing errors are still jointly significant with p -values close to zero. The reason is the small stock portfolios, whose returns are unusually high. When we remove them, the pricing error drops to just 4 bps, over two thirds smaller than the raw value. The p -value rises to 0.02, hence the model cannot be rejected at the 1% significance level. Value-weighting drives the pricing error down further, to just 0.6 bps, a 95% drop from the raw value. Its p -value rises to 0.79, hence the pricing error is insignificant. The same happens with the liquidity provider portfolio, whose pricing error is just 0.8 bps (0.4% annualized) with a p -value of 0.78. The two-factor model thus prices the liquidity provider portfolio almost perfectly and hence explains the liquidity premium in the cross section of stocks.

Table 8 shows the pricing errors of the individual long-short reversal strategies. Panel A looks at the CAPM specification from column 2 in Table 7. The Fama-McBeth pricing errors are almost identical to those from the time series regressions in Table 3. In particular, all of the Lo-Hi reversal strategies retain their statistically and economically significant pricing errors. The pricing error of the large-cap Lo-Hi strategy is 14 bps (t -statistic of 3.43), only slightly smaller than the 17 bps in Table 3. Thus, even if we treat the price of market risk as a free parameter, the CAPM cannot explain the returns to liquidity provision in stocks.

Panel B of Table 8 shows the pricing errors of the one-factor volatility risk model (column 3 of Table 7). The pricing error of the large-cap Lo-Hi reversal strategy drops to 8 bps. Only the smallest quintiles retain significant pricing errors, including the 68 bps error of the smallest quintile (t -statistic of 7.45).

Panel C looks at the two-factor model, which combines the market and volatility factors (column 4 of Table 7). The pricing error of the large-cap Lo-Hi strategy is fully explained: it drops to -2 bps with a t -statistic of -0.51 . The pricing errors of the middle

two quintiles are also eliminated. Only the smallest two quintiles retain significant pricing errors (58 bps for the smallest quintile versus a raw return of 78 bps). This confirms the finding in Table 7 that the overall r.m.s.e. of the two-factor model remains significant because of the very high reversal returns among the smallest stocks. Overall, Table 8 shows that volatility risk, in combination with the market factor, explains the returns to liquidity provision among all but the smallest stocks.

Figure 6 visualizes the results by plotting the average returns of the reversal strategies against their predicted values from the Fama-MacBeth regressions. Each marker shape and color combination represents a different size quintile. Within it are five data points corresponding to the five long-short reversal strategies across deciles: Lo–Hi, 2–9, and so on. Also shown is the overall liquidity provider portfolio as a hollow black circle.

Panel A shows that the CAPM cannot explain the returns of the reversal strategies. The average returns along the vertical axis display wide variation but the predicted returns along the horizontal axis are confined to a very narrow range. Moreover, the predicted returns are all close to zero, hence the pricing errors are similar to the raw returns.

Panel B shows the one-factor volatility risk model. Here the spread in predicted returns is substantially larger, reflecting the fact that volatility risk does a better job in explaining the cross section of reversals. This is especially true for the three largest size quintiles. Yet, since the model requires a somewhat low price of volatility risk to capture the level of returns, it ends up under-predicting differences in returns as captured by the long-short reversal strategies. This is why a number of them lie above the 45-degree line (i.e., their average returns are higher than their predicted returns).

Panel C of Figure 6 shows that the two-factor model captures the returns of the reversal strategies well. The range of predicted returns is wide and most of the strategies straddle the 45-degree line. The liquidity provider portfolio is almost exactly on the line (recall its pricing error is just 0.6 bps). Only the outer-decile small-stock strategies lie significantly above it. While their predicted returns are fairly high, their realized returns are even higher. The returns of small-stock reversals thus reflect additional compensation above and beyond their volatility risk exposure. This suggests that liquidity provision in small stocks is segmented from the rest of the market, perhaps due to inventory frictions. By contrast, liquidity provision is well integrated among medium and large stocks, which make up over 99% of the market by value. These results still leave the possibility that liq-

liquidity provision in larger stocks is segmented from the broader market. We turn to this question next.

4.6 Option-implied price of volatility risk

We now test whether the price of volatility risk needed to price the reversal strategies is in line with the one that prevails in option markets, where volatility risk is traded directly. Answering this question sharpens our test of Prediction 5. It also sheds light on the broader question of whether the returns to liquidity provision reflect market segmentation frictions or more widely shared economic risks.

Option markets are a natural place to measure the price of volatility risk. As discussed in Section 3, the squared VIX is the price of a basket of options whose payoff replicates the realized variance of the S&P 500. The average return on this basket therefore measures the price of variance risk. Notice that this return is not the same as the innovation in VIX squared because the VIX basket is rebalanced each day. To solve this problem, we track the price of each day's VIX basket to the following day. The percentage change in its price is the VIX return, which we denote R^{VIX} .

Table 2, which we discussed in Section 3.1, regresses the VIX return on our volatility risk factor, ΔVIX^2 . This gives us the option-implied price of risk per unit of ΔVIX^2 beta. From column 2, this price of risk is -1.03% over five days. We similarly obtain the price of market risk from the average market return over a five-day horizon in our sample, which is 16 bps. These prices of risk are reported in the top half of Table 9. Also shown are Newey-West standard errors based on the time series variation of the VIX return and market return.²⁴ The option-implied price of volatility risk and the market risk premium are both highly statistically significant.

The next step is to multiply these restricted prices of risk by the portfolios' betas to obtain their predicted returns. The portfolios' pricing errors are obtained by subtracting these predicted returns from their average returns. As in all prior tests, we use the time series variation in the pricing errors to estimate their standard errors. The bottom half of Table 9 reports the pricing errors for the same combinations of portfolios as in Table 7.

The first column of Table 9 shows the raw returns as before. In the second column, the

²⁴Note that unlike in Table 7 there is no cross-sectional R^2 since this is not a cross-sectional regression.

CAPM lowers the r.m.s.e. of the full set of portfolios from 24 bps to 11.6 bps. However, it has almost no impact on the long-short reversal strategies whose r.m.s.e. is 21.2 bps versus raw returns of 22.8 bps. The same is true when we exclude small stocks or value-weight, hence the lack of explanatory power is not due to small stocks. In the case of the overall liquidity provider portfolio, the pricing error ticks down from 16.8 bps to 14.6 bps, a one-tenth reduction. Thus, while the CAPM with a restricted price of risk does fairly well on the level of returns (the equity premium), it has almost no explanatory power for the differences in returns, i.e. the liquidity premium.

Column 3 shows that the opposite is true of the one-factor volatility risk model. This model does worse on the level of returns, driving the overall r.m.s.e. up to 76 bps, but much better on the difference in returns, driving the long-short reversal strategy r.m.s.e. down to 15.8 bps, a one-third reduction. The model's explanatory power increases substantially when we exclude small stocks: the r.m.s.e. drops from 12.3 to 3.8 bps, a two-thirds reduction. Value-weighting reduces the pricing error even more: from 12.4 to 1.2 bps, a 90% reduction. Finally, the pricing error of the liquidity provider portfolio drops from 16.8 bps to 2.5 bps, also a 90% reduction. The associated p -values are above conventional cutoffs, so the model cannot be rejected on larger stocks.

Column 4 looks at the two-factor model with volatility risk and market risk. This model combines the ability of the CAPM to explain the level of returns with the ability of the volatility risk model to explain the cross-sectional differences in returns. The overall r.m.s.e. falls by about half, the r.m.s.e. of the long-short reversal strategies falls by a third, and excluding small stocks makes it fall by two thirds (from 12.3 to 4 bps). As in column 3, value-weighting brings about a 90% drop in pricing errors (from 12.4 to 1.1 bps). The pricing error of the liquidity provider portfolio also drops by 90%, from 16.8 to 1.5 bps, and becomes insignificant. Thus, consistent with Prediction 5, using an option-implied price of volatility risk explains the returns to liquidity provision among larger stocks.

Table 10 reports the pricing errors of the individual reversal strategies. Panel A is very similar to Panel A of Table 8 and again shows that the CAPM cannot explain the returns of the Lo-Hi reversal strategies. By contrast, Panel B shows that the model with a restricted price of volatility risk fully explains the Lo-Hi returns among larger stocks: The pricing error of the Lo-Hi strategy for the largest quintile drops from 17 to -1 bps and becomes insignificant. The pricing errors of quintiles three and four are similarly eliminated. Panel

C shows the same pattern for the two-factor model.

While volatility risk is able to fully explain the reversal returns among larger stocks, it falls short on the smallest stocks: the Lo–Hi reversal strategy for the smallest quintile has a pricing error of 59 bps in Panel C, which is down only 25% from its raw return. Similarly, the Lo–Hi reversal strategy for the second smallest quintile has a pricing error of 23 bps, down 40% from the raw return. Both remain highly significant. While both strategies have similar volatility betas to those of the larger quintiles, their returns are much larger. Thus, liquidity provision among small stocks earns abnormal returns. This is again consistent with market segmentation, as discussed in Section 2.1.

Finally, Figure 7 depicts average versus predicted returns with the restricted prices of risk. Similar to Figure 6, the CAPM (Panel A) has no ability to fit the reversal strategy returns, while the one-factor volatility risk model (Panel B) and the two-factor volatility risk plus market risk model (Panel C) capture these returns along the 45-degree line. Only the Lo–Hi strategies of the two smallest quintiles and the 2–9 strategy of the smallest quintile lie significantly away from the line. The liquidity provider portfolio, which captures the overall returns to liquidity provision, is priced almost perfectly.

Overall, the results with an option-implied price of volatility risk show that it carries the same price of risk in the markets for liquidity (reversals) and volatility (options). This price of risk explains the overall return to liquidity provision in stocks. This finding supports Prediction 5, which is a distinguishing implication of our model.²⁵ By contrast, models that rely on market segmentation predict that liquidity provision earns abnormal returns. We find evidence for this among the smallest stocks. This evidence is consistent with the result in Section 4.4 that the volatility betas of small-stock reversals are transitory, as predicted by the extended model with inventory cost shocks. By contrast, the betas of medium- and large-stock reversals are permanent, as predicted by the private information mechanism of our main model in which markets are integrated. The results in this section further support this interpretation, indicating that the liquidity premium outside small stocks is explained by the variance risk premium.

²⁵In Internet Appendix IA.5, we find that the same result holds in a new cross section of reversals inspired by the model. Instead of quintiles by size, we form quintiles by the co-movement of a stock's idiosyncratic volatility with market volatility. This volatility co-movement proxies for the strength of the relationship between the amount of private information in a stock and market volatility shocks. As our model implies, high volatility co-movement stocks have larger reversals, their reversals have more negative volatility betas, and these betas explain their reversal returns with an option-implied price of volatility risk.

5 Conclusion

Our results provide a new perspective on the risks and returns to liquidity provision in financial markets. Under this perspective, the price of liquidity reflects the cost of hedging the volatility risk embedded in liquidity provision. This volatility risk stems from the exposure of liquidity providers to uncertainty about the amount of private information they face. A spike in volatility reveals greater private information than liquidity providers priced in, triggering losses on both sides of their portfolios. Consistent with this view, our empirical results show that short-term reversals, which mimic the portfolios of liquidity providers, are exposed to significant and persistent volatility risk. Moreover, this volatility risk exposure explains the returns to reversals, and hence liquidity provision, both over time and among all but the smallest stocks.

More broadly, we find that liquidity is priced based on the broadly shared economic risk of systematic volatility shocks, as opposed to the narrow risk of liquidity providers' financial constraints. Just how broad is this risk? The literature on the variance risk premium emphasizes the fundamental macroeconomic risks faced by a representative agent. Yet it is also possible that the variance risk premium itself is a reflection of the economic centrality of the financial sector, which is in the business of liquidity provision. This possibility is intriguing, as it promises to further integrate our understanding of asset pricing and financial intermediation.

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.
- 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.
- Andersen, Torben G, Nicola Fusari, and Viktor Todorov.** 2015. "The risk premia embedded in index options." *Journal of Financial Economics*, 117(3): 558–584.
- 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.
- 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): 747–777.
- 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.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes.** 2007. "Model specification and risk premia: Evidence from futures options." *Journal of Finance*, 62(3): 1453–1490.
- 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.
- Chung, Juliet.** 2020a. "Covid-19 Caused Chaos for Investors in 2020. These Hedge Funds Earned Billions." *The Walls Street Journal*.
- Chung, Juliet.** 2020b. "This Hedge Fund Saw Risks of Coronavirus Early. Now It's Up 36%." *The Walls Street Journal*.
- Collin-Dufresne, Pierre, and Kent Daniel.** 2014. "Liquidity and return reversals." *Working paper*.
- Collin-Dufresne, Pierre, and Vyacheslav Fos.** 2015. "Do prices reveal the presence of informed trading?" *The Journal of Finance*, 70(4): 1555–1582.
- 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.
- Easley, David, and Maureen O'Hara.** 2004. "Information and the cost of capital." *The journal of finance*, 59(4): 1553–1583.
- 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.
- Geanakoplos, John.** 2003. "Liquidity, default, and crashes."
- Gertler, Mark, and Nobuhiro Kiyotaki.** 2010. "Financial intermediation and credit policy in business cycle analysis." In *Handbook of monetary economics*. Vol. 3, 547–599. Elsevier.
- 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.
- 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.** 1988. “Trades, quotes, inventories, and information.” *Journal of financial economics*, 22(2): 229–252.
- 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.
- Hendershott, Terrence, and Mark S Seasholes.** 2007. “Market maker inventories and stock prices.” *American Economic Review*, 97(2): 210–214.
- 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.
- Holmström, Bengt, and Jean Tirole.** 1998. “Private and public supply of liquidity.” *Journal of political Economy*, 106(1): 1–40.
- Kiyotaki, Nobuhiro, and John Moore.** 1997. “Credit cycles.” *Journal of Political Economy*, 105(2).
- 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.
- Lo, Andrew W, and A Craig MacKinlay.** 1990. “When are contrarian profits due to stock market overreaction?” *The review of financial studies*, 3(2): 175–205.
- 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, Luboš, and Robert F Stambaugh.** 2003. “Liquidity risk and expected stock returns.” *Journal of Political economy*, 111(3): 642–685.
- Rampini, Adriano A, and S Viswanathan.** 2019. “Financial intermediary capital.” *The Review of Economic Studies*, 86(1): 413–455.
- 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.
- Singleton, Kenneth J.** 2006. *Empirical Dynamic Asset Pricing*. Princeton University Press.
- 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.

Todorov, Viktor. 2010. “Variance risk-premium dynamics: The role of jumps.” *The Review of Financial Studies*, 23(1): 345–383.

Table 1: Summary statistics

This table shows summary statistics for the reversal strategies. Each day, stocks are first sorted into quintiles by market capitalization and then deciles by normalized beta-adjusted return. The normalized return is calculated using a 60-day rolling window. The portfolios are weighted by average dollar volume over that window. Stocks with share price in the bottom 20% and stocks with an earnings announcement on the portfolio formation day or the prior day are excluded. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Market cap						Panel B: Idiosyncratic volatility					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.12	0.12	0.12	0.12	0.12	Small	3.74	4.22	4.77	5.50	5.96
2	0.36	0.36	0.36	0.36	0.36	2	3.41	3.72	3.94	4.09	4.17
3	0.89	0.90	0.89	0.89	0.89	3	2.98	3.20	3.34	3.44	3.48
4	2.49	2.49	2.48	2.48	2.48	4	2.47	2.67	2.79	2.87	2.90
Big	81.75	81.21	81.08	78.77	79.26	Big	1.71	1.85	1.94	1.99	2.01

Panel C: Amihud illiquidity						Panel D: Sorting-day return					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	171.66	105.00	55.79	33.03	17.10	Small	-13.55	-7.12	-4.57	-2.59	-0.84
2	7.31	4.42	2.93	1.94	1.60	2	-10.80	-5.57	-3.52	-1.98	-0.64
3	0.94	0.65	0.48	0.37	0.32	3	-8.98	-4.63	-2.96	-1.67	-0.54
4	0.19	0.14	0.11	0.09	0.08	4	-7.39	-3.99	-2.57	-1.46	-0.47
Big	0.02	0.01	0.01	0.08	0.01	Big	-5.20	-2.98	-1.95	-1.12	-0.36

Panel E: Average turnover						Panel F: Sorting-day turnover					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	1.14	1.33	1.47	1.54	1.53	Small	2.59	1.90	1.83	1.82	1.71
2	1.47	1.54	1.63	1.65	1.68	2	2.42	1.58	1.44	1.40	1.37
3	1.74	1.80	1.84	1.87	1.88	3	2.61	1.74	1.58	1.52	1.49
4	1.78	1.86	1.91	1.94	1.96	4	2.55	1.82	1.67	1.61	1.59
Big	1.16	1.23	1.28	1.30	1.30	Big	1.50	1.20	1.14	1.10	1.08

Table 2: Summary statistics: the VIX return

This table shows summary statistics for VIX^2 changes and the VIX return (Panel A) and regressions of the VIX return on VIX^2 changes (Panel B). VIX^2 is the squared VIX index divided by 100. It represents the price of a basket of options whose payoff replicates the variance of the S&P 500 over the following 30 calendar days. The VIX return, R^{VIX} , is the excess return on this basket of options. It is not equal to the percentage change in VIX^2 because the VIX^2 basket changes each day (to keep the horizon constant). The VIX return is the percentage change in the price of a given basket from one day to the next (minus the risk-free rate). The sample is from April 9, 2001 to December 31, 2019 (the latest date for which OptionMetrics data is available).

Panel A: Summary statistics

	Mean	St. Dev.	1 st	5 th	Median	95 th	99 th
R^{VIX}	-1.53	17.86	-25.33	-17.34	-5.08	26.85	71.29
ΔVIX	-0.00	1.82	-4.58	-2.25	-0.08	2.48	5.48
ΔVIX^2	-0.00	1.53	-3.65	-1.14	-0.02	1.18	3.89
$\Delta VIX, \%$	0.26	7.47	-15.55	-9.62	-0.53	12.40	25.13
$\Delta VIX^2, \%$	1.07	16.28	-28.69	-18.32	-1.05	26.35	56.57

Panel B: Regressions

	R^{VIX}			
	(1)	(2)	(3)	(4)
ΔVIX	7.939*** (0.102)			
ΔVIX^2		7.404*** (0.172)		
$\Delta VIX, \%$			2.080*** (0.018)	
$\Delta VIX^2, \%$				0.991*** (0.008)
Constant	-1.498*** (0.172)	-1.515*** (0.220)	-2.013*** (0.135)	-2.523*** (0.121)
Obs.	4,711	4,711	4,711	4,711
R^2	0.562	0.282	0.730	0.783

Table 3: Reversal strategy returns

Average returns, standard deviations, Sharpe ratios, and CAPM alphas of the five-day reversal strategies. Each day, stocks are first sorted into quintiles by market capitalization and then deciles by normalized beta-adjusted return. The normalized return is calculated using a 60-day rolling window. The portfolios are weighted by average dollar volume over that window. Returns and standard deviations are over five days. Sharpe ratios are annualized. The t -statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Average returns						Panel B: t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.78	0.46	0.20	-0.08	-0.10	Small	8.39	5.84	2.21	-0.98	-1.30
2	0.40	0.15	0.06	0.08	-0.02	2	6.37	3.04	1.34	1.74	-0.37
3	0.15	0.12	0.04	0.04	-0.02	3	2.79	2.85	1.05	0.98	-0.63
4	0.16	0.19	0.11	0.12	0.02	4	3.35	4.73	3.38	3.92	0.65
Big	0.20	0.16	0.18	0.08	0.00	Big	4.47	4.66	6.17	3.54	0.17

Panel C: Standard deviations						Panel D: Sharpe ratios					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	6.18	5.49	5.60	5.43	5.41	Small	0.90	0.60	0.25	-0.10	-0.14
2	3.91	3.35	3.24	3.08	3.10	2	0.72	0.32	0.14	0.18	-0.04
3	3.31	2.71	2.67	2.65	2.20	3	0.32	0.30	0.11	0.10	-0.07
4	3.01	2.52	2.16	1.95	1.88	4	0.37	0.53	0.35	0.43	0.07
Big	2.84	2.29	1.96	1.70	1.52	Big	0.50	0.48	0.64	0.35	0.02

Panel E: CAPM alphas						Panel F: CAPM alpha t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.74	0.44	0.16	-0.10	-0.10	Small	7.95	5.62	1.76	-1.23	-1.28
2	0.38	0.12	0.05	0.07	-0.02	2	6.16	2.46	1.09	1.55	-0.50
3	0.12	0.10	0.03	0.03	-0.03	3	2.37	2.42	0.66	0.71	-0.81
4	0.13	0.17	0.10	0.11	0.02	4	2.81	4.38	3.12	3.69	0.87
Big	0.17	0.14	0.16	0.07	-0.00	Big	3.90	4.23	5.60	3.14	-0.15

Table 4: Volatility betas of the reversal strategies

The table shows the betas of the five-day reversal strategies on changes in VIX^2 . The betas are estimated by running the regressions

$$R_{t,t+5}^p = \alpha_p + \beta_{VIX}^p \Delta VIX_{t,t+5}^2 + \beta_M^p R_{t,t+5}^M + \epsilon_{t,t+5}^p,$$

where $R_{t,t+5}^p$ is the cumulative excess return on reversal strategy portfolio p from the portfolio formation date t to $t+5$, $R_{t,t+5}^M$ is the excess return on the market portfolio, and $\Delta VIX_{t,t+5}^2$ is the change in the squared VIX index from date t to $t+5$. Panel A omits the market return while Panels B and C include it. Panel C reports the market betas β_M^p . The t -statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: ΔVIX^2 betas

β_{VIX}^p						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	-0.18	-0.05	-0.15	-0.02	-0.05	Small	-3.48	-0.82	-1.91	-0.27	-0.80
2	-0.18	-0.10	-0.08	-0.10	-0.03	2	-3.45	-2.45	-2.06	-3.80	-1.24
3	-0.19	-0.11	-0.11	-0.05	0.00	3	-3.89	-3.29	-4.55	-2.13	0.01
4	-0.17	-0.13	-0.06	-0.06	0.01	4	-3.98	-3.73	-2.23	-2.25	0.32
Big	-0.20	-0.11	-0.13	-0.07	-0.04	Big	-4.27	-3.31	-4.10	-3.30	-1.72

Panel B: ΔVIX^2 betas (controlling for R^M)

β_{VIX}^p						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	-0.17	-0.02	-0.11	0.10	-0.03	Small	-2.56	-0.21	-1.04	1.34	-0.35
2	-0.16	-0.03	-0.06	-0.08	-0.06	2	-1.84	-0.62	-0.91	-2.07	-1.69
3	-0.18	-0.11	-0.10	-0.03	-0.00	3	-2.54	-1.99	-2.88	-0.95	-0.18
4	-0.15	-0.12	-0.04	-0.06	-0.00	4	-2.18	-2.71	-0.87	-1.46	-0.10
Big	-0.20	-0.09	-0.14	-0.08	-0.03	Big	-3.12	-1.94	-3.28	-2.77	-1.04

Panel C: Market betas

β_M^p						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.02	0.04	0.06	0.16	0.03	Small	0.38	0.56	1.02	3.04	0.59
2	0.04	0.10	0.04	0.04	-0.04	2	0.53	2.56	0.93	1.03	-1.33
3	0.01	0.01	0.01	0.03	-0.00	3	0.25	0.17	0.30	1.00	-0.18
4	0.03	0.01	0.04	-0.00	-0.01	4	0.49	0.36	1.47	-0.02	-0.55
Big	0.00	0.03	-0.01	-0.01	0.01	Big	0.03	1.10	-0.49	-0.31	0.22

Table 5: Predicting reversal returns with VIX

Results from predictability regressions of the reversal strategy returns on VIX. The Lo–Hi strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within a given size quintile. The remaining strategies are constructed analogously for the inner normalized return deciles. Each strategy is held for five trading days. The predictability regressions are

$$R_{t,t+5}^p = a_r^p + b_r^p VIX_t^2 + \epsilon_{r,t+5}^p,$$

where $R_{t,t+5}^p$ is the cumulative excess return on portfolio p from the portfolio formation date t to $t + 5$ and VIX_t^2 is the squared VIX index on the portfolio formation date. The predictive loadings b_r^p are multiplied by 100 for legibility. The t -statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Predictive loadings						Panel B: t -statistics					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	9.81	7.69	9.51	−0.48	−2.60	Small	3.44	2.89	3.21	−0.19	−0.95
2	7.71	4.34	3.70	1.68	−0.34	2	3.00	2.92	2.62	1.18	−0.25
3	6.72	3.13	3.00	2.65	0.44	3	2.98	1.98	2.59	1.85	0.63
4	5.70	5.28	3.01	3.41	−0.54	4	2.94	3.34	2.51	2.76	−0.72
Big	9.09	5.51	3.53	3.02	0.34	Big	4.65	4.13	2.98	3.00	0.47

Panel C: R^2					
	Lo–Hi	2–9	3–8	4–7	5–6
Small	0.85	0.66	0.97	0.00	0.08
2	1.30	0.56	0.44	0.10	0.00
3	1.38	0.45	0.43	0.34	0.01
4	1.21	1.47	0.65	1.03	0.03
Big	3.45	1.94	1.09	1.06	0.02

Table 6: Volatility risk persistence of the reversal strategies

The table shows the impulse response coefficients of the reversal strategies at different horizons on VIX^2 shocks one day after portfolio formation. The response coefficients are estimated by running the regressions

$$R_{t,t+h}^p = \alpha_p + \delta_{VIX,h}^p \Delta VIX_{t,t+1}^2 + \epsilon_{t,t+5}^p,$$

where $R_{t,t+h}^p$ is the cumulative excess return on reversal strategy portfolio p from the portfolio formation date t to $t+h$, and $\Delta VIX_{t,t+1}^2$ is the change in the squared VIX index from date t to $t+1$. The panels show the coefficients $\delta_{VIX,h}^p$. Panel A uses $h = 5$ days, Panel B uses $h = 10$ days, and Panel C uses $h = 20$ days. The t -statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: 5 days

$\delta_{VIX,5}^p$						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	-0.10	-0.13	-0.15	0.01	-0.13	Small	-0.87	-1.46	-1.57	0.16	-1.23
2	-0.20	-0.05	-0.03	-0.17	-0.09	2	-2.38	-1.00	-0.55	-3.34	-1.71
3	-0.18	-0.07	-0.10	-0.09	-0.03	3	-3.04	-1.85	-1.68	-2.00	-0.70
4	-0.20	-0.12	-0.09	-0.05	0.03	4	-3.88	-2.48	-2.13	-1.57	1.19
Big	-0.19	-0.08	-0.10	-0.03	-0.02	Big	-3.12	-1.44	-2.81	-0.57	-0.62

Panel B: 10 days

$\delta_{VIX,10}^p$						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.03	-0.14	-0.39	0.05	-0.14	Small	0.21	-1.40	-3.42	0.44	-0.88
2	-0.01	-0.03	0.03	-0.08	-0.02	2	-0.09	-0.55	0.53	-1.11	-0.57
3	-0.14	-0.13	-0.07	-0.07	-0.07	3	-1.88	-2.21	-1.38	-1.37	-1.87
4	-0.17	-0.12	-0.12	-0.03	0.00	4	-3.13	-2.36	-3.04	-0.65	0.08
Big	-0.21	-0.10	-0.08	-0.08	0.02	Big	-3.60	-2.37	-2.56	-1.90	0.89

Panel C: 20 days

$\delta_{VIX,20}^p$						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.02	-0.28	-0.14	-0.00	-0.10	Small	0.26	-2.68	-1.48	-0.00	-0.60
2	0.01	-0.06	0.10	-0.14	-0.09	2	0.04	-0.95	1.06	-1.32	-0.95
3	-0.22	-0.12	-0.09	-0.14	-0.05	3	-3.85	-2.61	-1.81	-3.33	-1.09
4	-0.23	-0.12	-0.12	0.03	0.09	4	-3.42	-1.98	-3.44	0.35	3.53
Big	-0.19	-0.07	-0.05	-0.11	0.00	Big	-3.86	-1.48	-1.30	-1.82	0.03

Table 7: Fama-Macbeth asset pricing tests

The table shows results from Fama-Macbeth regressions of the reversal portfolios. The first-stage regressions are $R_{t,t+5}^p = \alpha_p + \beta_{VIX}^p \Delta VIX_{t,t+5}^2 + \beta_M^p R_{t,t+5}^M + \epsilon_{t,t+5}^p$, where $R_{t,t+5}^p$ is the cumulative excess return on portfolio p from t to $t + 5$, $\Delta VIX_{t,t+5}^2$ is the change in the squared VIX, and $R_{t,t+5}^M$ is the excess market return. The second-stage regressions are

$$R_{t,t+5}^p = \lambda_{0,t} + \lambda_{t,VIX} \beta_{VIX}^p + \lambda_{t,M} \beta_M^p + e_{t,t+5}^p.$$

The table reports the time-series averages of the premia, λ_{VIX} and λ_M , and constant, λ_0 . Column (1) reports raw returns, column (2) adds in R^M , column (3) replaces it with ΔVIX^2 , and column (4) includes both. Standard errors are Newey-West with five lags to account for the overlap in returns. Also shown are the second-stage R^2 , the root-mean-squared error (r.m.s.e.) among (i) all portfolios, (ii) the long-short reversal strategies, and (iii) the long-short reversal strategies excluding the smallest two quintiles, as well as the pricing errors of (iv) the value-weighted long-short reversal strategy, and (v) the liquidity provider portfolio. The liquidity provider portfolio is double-weighted by volume and the normalized sorting-day return. The sample is from April 9, 2001 to May 31, 2020.

	(1)	(2)	(3)	(4)
β_M		0.325*** (0.122)		0.152 (0.125)
β_{VIX}			-0.570*** (0.156)	-1.079*** (0.230)
Constant	0.205** (0.101)	-0.434*** (0.115)	-0.527*** (0.097)	-0.251** (0.127)
N	50	50	50	50
R^2	0.000	0.158	0.260	0.379
<i>(i) All portfolios:</i>				
R.m.s.e.	0.240	0.112	0.107	0.097
p-value	0.00	0.00	0.00	0.00
<i>(ii) Long-short reversal strategies:</i>				
R.m.s.e.	0.228	0.201	0.183	0.158
p-value	0.00	0.00	0.00	0.00
<i>(iii) Long-short reversal strategies (ex small stocks):</i>				
R.m.s.e.	0.123	0.094	0.065	0.040
p-value	0.00	0.00	0.01	0.02
<i>(iv) Value-weighted reversal strategy:</i>				
Pricing error	0.124	0.094	0.062	0.006
p-value	0.00	0.00	0.00	0.79
<i>(v) Liquidity provider portfolio:</i>				
Pricing error	0.168	0.127	0.084	0.008
p-value	0.00	0.00	0.00	0.78

Table 8: Fama-Macbeth pricing errors

The table shows the pricing errors of the reversal strategies from Fama-Macbeth regressions (see Table 7 for the factor premia). The Lo-Hi strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within a given size quintile. The remaining strategies are constructed analogously for the inner normalized return deciles. Each strategy is held for five trading days. Panel A reports the pricing errors when the market return is used as the only factor. Panel B replaces it with the change in the squared VIX index. Panel C includes both factors. The reported t statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Market

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.72	0.43	0.13	-0.12	-0.11	Small	7.87	5.52	1.50	-1.43	-1.38
2	0.35	0.10	0.04	0.06	-0.02	2	5.91	2.06	0.78	1.23	-0.51
3	0.10	0.08	0.01	0.02	-0.03	3	1.95	2.18	0.30	0.47	-0.81
4	0.11	0.15	0.09	0.10	0.03	4	2.42	4.03	2.83	3.50	0.94
Big	0.14	0.12	0.15	0.07	-0.01	Big	3.43	3.86	5.26	2.88	-0.39

Panel B: ΔVIX^2

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.67	0.43	0.11	-0.09	-0.13	Small	7.45	5.49	1.25	-1.10	-1.62
2	0.29	0.09	0.02	0.02	-0.03	2	5.08	1.98	0.34	0.52	-0.77
3	0.04	0.05	-0.02	0.01	-0.02	3	0.84	1.36	-0.50	0.18	-0.63
4	0.06	0.11	0.07	0.08	0.02	4	1.40	3.06	2.32	2.80	0.78
Big	0.08	0.09	0.10	0.04	-0.02	Big	1.95	2.99	3.65	1.88	-0.84

Panel C: Market and ΔVIX^2

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.58	0.42	0.05	-0.00	-0.13	Small	6.70	5.43	0.60	-0.05	-1.58
2	0.22	0.09	-0.01	-0.00	-0.08	2	4.08	1.85	-0.12	-0.05	-1.80
3	-0.05	-0.01	-0.07	-0.01	-0.03	3	-1.11	-0.17	-1.80	-0.16	-0.96
4	-0.02	0.05	0.06	0.05	0.02	4	-0.45	1.32	2.13	1.77	0.80
Big	-0.02	0.06	0.03	0.00	-0.04	Big	-0.51	1.80	0.86	0.02	-1.65

Table 9: Option-implied price of volatility risk

The table shows pricing results for the five-day reversal portfolios using an option-implied price of volatility risk. The option-implied price of volatility risk is the one that prices the VIX return (see Table 2). The restricted price of market risk is the one that prices the market return. To obtain the pricing errors of the reversal portfolios, we multiply their betas by the restricted prices of risk and subtract the resulting predicted returns from the average returns. The table reports the restricted prices of risk with standard errors based on the time series variation of the VIX return and market return. Column (1) reports raw returns and includes a constant, column (2) adds in the market return, column (3) replaces it with the change in the squared VIX, and column (4) includes both. Standard errors are Newey-West with five lags to account for the overlap in returns. Also shown are the root-mean-squared error (r.m.s.e.) among (i) all portfolios, (ii) the long-short reversal strategies, and (iii) the long-short reversal strategies excluding the smallest two quintiles, as well as the pricing errors of (iv) the value-weighted long-short reversal strategy, and (v) the liquidity provider portfolio. The liquidity provider portfolio is double-weighted by volume and the normalized sorting-day return. Below each pricing error is its associated p value. The sample is from April 9, 2001 to May 31, 2020.

β_M		0.160** (0.081)		0.160** (0.081)
β_{VIX}			-1.032*** (0.173)	-1.032*** (0.173)
Constant	0.205** (0.101)			
N	50	50	50	50
<i>(i) All portfolios:</i>				
R.m.s.e.	0.240	0.116	0.760	0.117
p -value	0.00	0.00	0.00	0.00
	Constant	Market	VIX	Market+VIX
<i>(ii) Long-short reversal strategies:</i>				
R.m.s.e.	0.228	0.212	0.158	0.159
p -value	0.00	0.00	0.00	0.00
<i>(iii) Long-short reversal strategies (ex small stocks):</i>				
R.m.s.e.	0.123	0.107	0.038	0.040
p -value	0.00	0.00	0.05	0.02
<i>(iv) Value-weighted reversal strategy:</i>				
Pricing error	0.124	0.108	0.012	0.011
p -value	0.00	0.00	0.59	0.62
<i>(v) Liquidity provider portfolio:</i>				
Pricing error	0.168	0.146	0.015	0.015
p -value	0.00	0.00	0.62	0.62

Table 10: Pricing errors with an option-implied price of volatility risk

The table shows the pricing errors of the five-day reversal strategies using an option-implied price of volatility. The option-implied price of volatility risk is the one that prices the VIX return (see Table 2). The restricted price of market risk is the one that prices the market return. To obtain the pricing errors of the reversal strategies, we multiply their betas by the restricted prices of risk and subtract the resulting predicted returns from the average returns. The Lo-Hi strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within a given size quintile. The remaining strategies are constructed analogously for the inner normalized return deciles. Each strategy is held for five trading days. Panel A reports the pricing errors when the market return is used as the only factor. Panel B replaces it with the change in the squared VIX index. Panel C includes both factors. The reported t statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Market

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.74	0.44	0.16	-0.10	-0.10	Small	8.02	5.66	1.76	-1.24	-1.28
2	0.38	0.12	0.05	0.07	-0.02	2	6.04	2.41	1.07	1.54	-0.50
3	0.12	0.10	0.02	0.03	-0.03	3	2.29	2.43	0.64	0.71	-0.81
4	0.13	0.17	0.10	0.11	0.02	4	2.76	4.25	3.09	3.71	0.88
Big	0.17	0.14	0.16	0.07	-0.00	Big	3.82	4.14	5.58	3.14	-0.16

Panel B: ΔVIX^2

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.59	0.41	0.04	-0.10	-0.15	Small	6.35	5.19	0.44	-1.20	-1.89
2	0.21	0.05	-0.02	-0.02	-0.05	2	3.36	0.94	-0.48	-0.46	-1.11
3	-0.05	-0.00	-0.07	-0.02	-0.02	3	-0.85	-0.03	-1.75	-0.47	-0.63
4	-0.02	0.05	0.04	0.06	0.02	4	-0.40	1.33	1.26	1.85	0.88
Big	-0.01	0.04	0.05	0.01	-0.04	Big	-0.20	1.29	1.64	0.43	-1.66


Panel C: Market and ΔVIX^2

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.59	0.42	0.06	-0.01	-0.13	Small	6.37	5.42	0.63	-0.12	-1.58
2	0.23	0.09	-0.00	0.00	-0.08	2	3.69	1.78	-0.06	0.02	-1.79
3	-0.04	-0.00	-0.07	-0.00	-0.03	3	-0.84	-0.04	-1.72	-0.12	-0.95
4	-0.01	0.05	0.06	0.06	0.02	4	-0.25	1.37	2.04	1.85	0.80
Big	-0.01	0.06	0.03	0.00	-0.03	Big	-0.32	1.76	1.22	0.18	-1.59

Figure 1: Reversal strategy returns and volatility risk

The figure shows the average return and volatility risk of the reversal strategy against the VIX index. The reversal strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within the top size quintile. The portfolios are formed each day and held for five trading days. Panel A plots the annualized average return of the strategy over a 60-day window. Panel B plots its volatility risk estimated over the same window. The volatility risk of the reversal strategy is the annualized standard deviation of changes in squared VIX times the strategy's beta with respect to these changes, i.e. $\sigma(\beta_{VIX^2} \Delta VIX^2)$. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Reversal strategy average return



Panel B: Reversal strategy volatility risk





Figure 2: VIX, idiosyncratic volatility and order flow-induced volatility

The figure plots idiosyncratic volatility, order flow-induced volatility, and VIX. Idiosyncratic volatility is calculated as follows: Each day, we compute the beta-adjusted returns of all stocks (the betas use 60-day rolling windows). We then square these returns and value-weight them across stocks. We then take the annualized sum over the following 30 calendar days (to match the horizon of VIX), and take the square root. Order flow-induced volatility is calculated as follows: Each day, we run a regression of five-minute returns (from TAQ) on the signed squared-root of order flow. We then compute the daily fitted values of this regression, square them and value-weight them across stocks. We again take their annualized sum over the following 30 calendar days and take the square root. The sample is from April 9, 2001 (September 10, 2003 for order-flow induced volatility) to May 31, 2020.

Panel A: Idiosyncratic volatility and VIX



Panel B: Order flow-induced volatility and VIX




Figure 3: Volatility betas before and after portfolio formation

The figure shows the market volatility betas of the Lo–Hi reversal strategies before and after portfolio formation. The Lo–Hi strategies go long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within a size quintile. Once the portfolios are formed, we compute cumulative returns before and after portfolio formation. For the period before formation, we compute cumulative returns from up to ten days before the formation date to the formation date, i.e. we construct the cumulative return R_h^p from h to 0 for $h = -10, \dots, -1$. For the period after formation, we compute them from the formation date to up to ten days after, i.e. we construct R_h^p from 0 to h for $h = 1, \dots, 10$. The portfolio formation date 0 itself is excluded. We then regress these returns on the change in VIX squared over the same period, $\Delta VIX_{0,h}^2$:

$$R_h^p = \alpha_h^p + \beta_h^p \Delta VIX_{0,h}^2 + \epsilon_{0,h}^p \quad \text{for } h = -10, \dots, 10.$$

Panel A shows the estimated coefficients β_h^p for the top size quintile together with a 95% confidence interval. Panel B shows β_h^p for all five size quintiles. The sample is from April 9, 2001 to May 31, 2020.

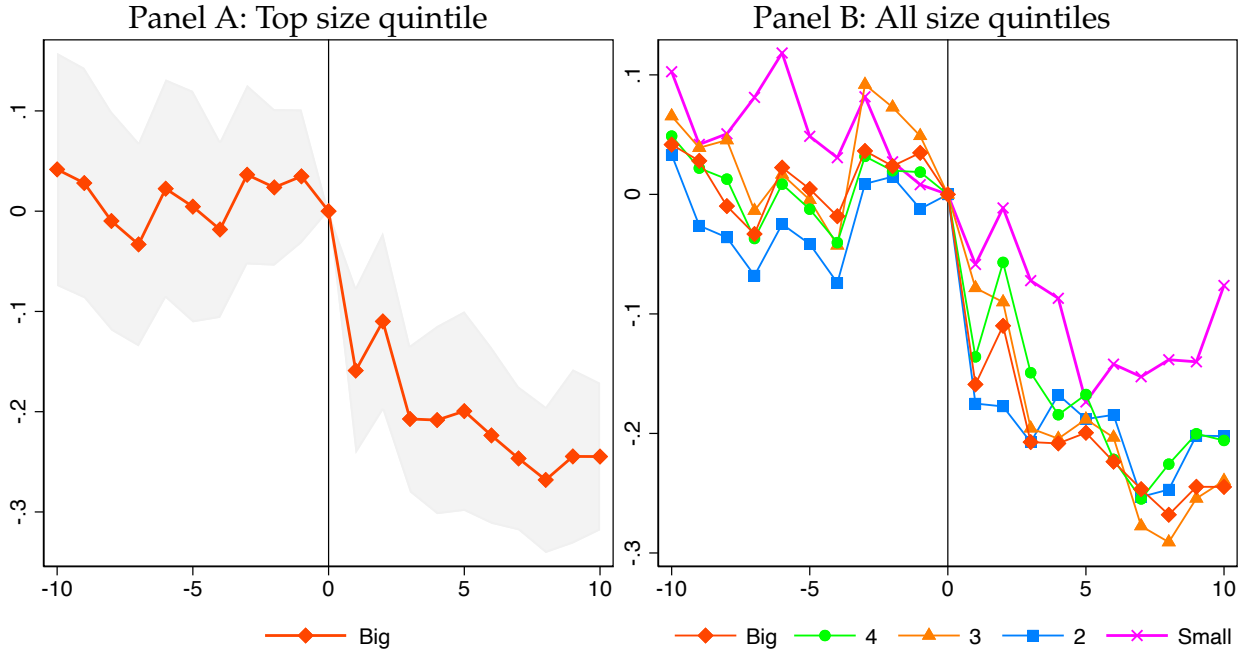


Figure 4: Returns, predictive loadings, and volatility betas by horizon

The figure shows average returns (Panel A), predictive loadings by VIX^2 (Panel B), volatility risk betas (Panel C), and volatility risk betas controlling for the market return (Panel D) for the reversal strategies within the largest size quintile. The Lo-Hi strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio (the remaining strategies are constructed analogously). Each strategy is held for up to ten trading days. The horizontal axis shows the holding period. On the vertical axis, Panel A plots average returns and Panel B plots the predictive loading on VIX^2 . The sample is from April 9, 2001 to May 31, 2020.

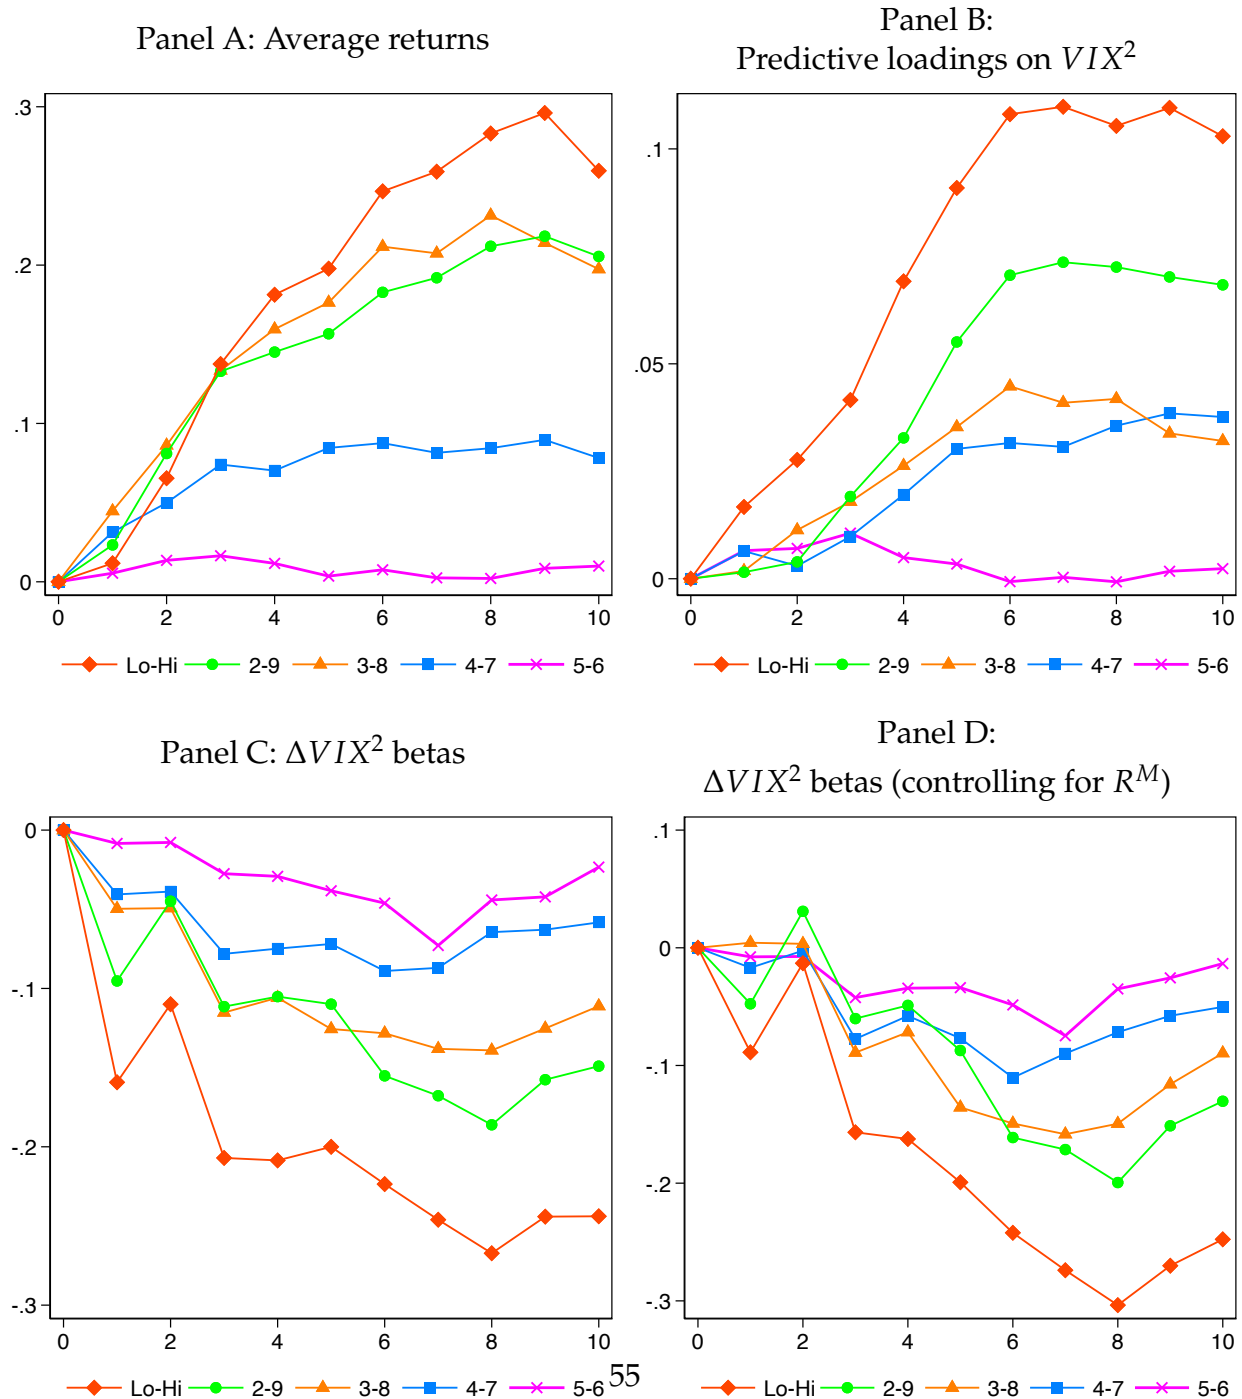


Figure 5: Volatility risk persistence of the reversal strategies

The figure shows the response coefficients of reversal strategies to VIX^2 shocks one day after portfolio formation. The response coefficients are estimated by running the regressions

$$R_{t,t+h}^p = \alpha_{p,h} + \delta_{VIX,h}^p \Delta VIX_{t,t+1}^2 + \epsilon_{t,t+h}^p$$

where $R_{t,t+h}^p$ is the cumulative excess return on portfolio p from the portfolio formation date t to $t+h$ and $\Delta VIX_{t,t+1}^2$ is the change in the squared VIX from t to $t+1$. 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. Panel A focuses on the top size quintile. Gray shading indicates 95% confidence interval. Panel B includes the smaller size quintiles. Each strategy is held for up to twenty trading days. The horizontal axis shows the holding period. The sample is from April 9, 2001 to May 31, 2020.




Figure 6: Fama-MacBeth regressions: average and predicted returns

The figure shows the average returns of the reversal strategies against their predicted returns from Fama-Macbeth regressions (see Tables 7 and 8). Each observation corresponds to a different reversal strategy. Colors and markers denote different size quintiles. Within each quintile, there are five long-short reversal strategies (Lo–Hi, 2–9, 3–8, 4–7, and 5–6). Each strategy is held for five trading days. Also shown is the liquidity provider (LP) portfolio, which weighs the reversal strategies by their normalized return and volume. Panel A uses the market return as the only factor. Panel B replaces it with the change in the squared VIX index. Panel C includes both factors. The sample is from April 9, 2001 to May 31, 2020.

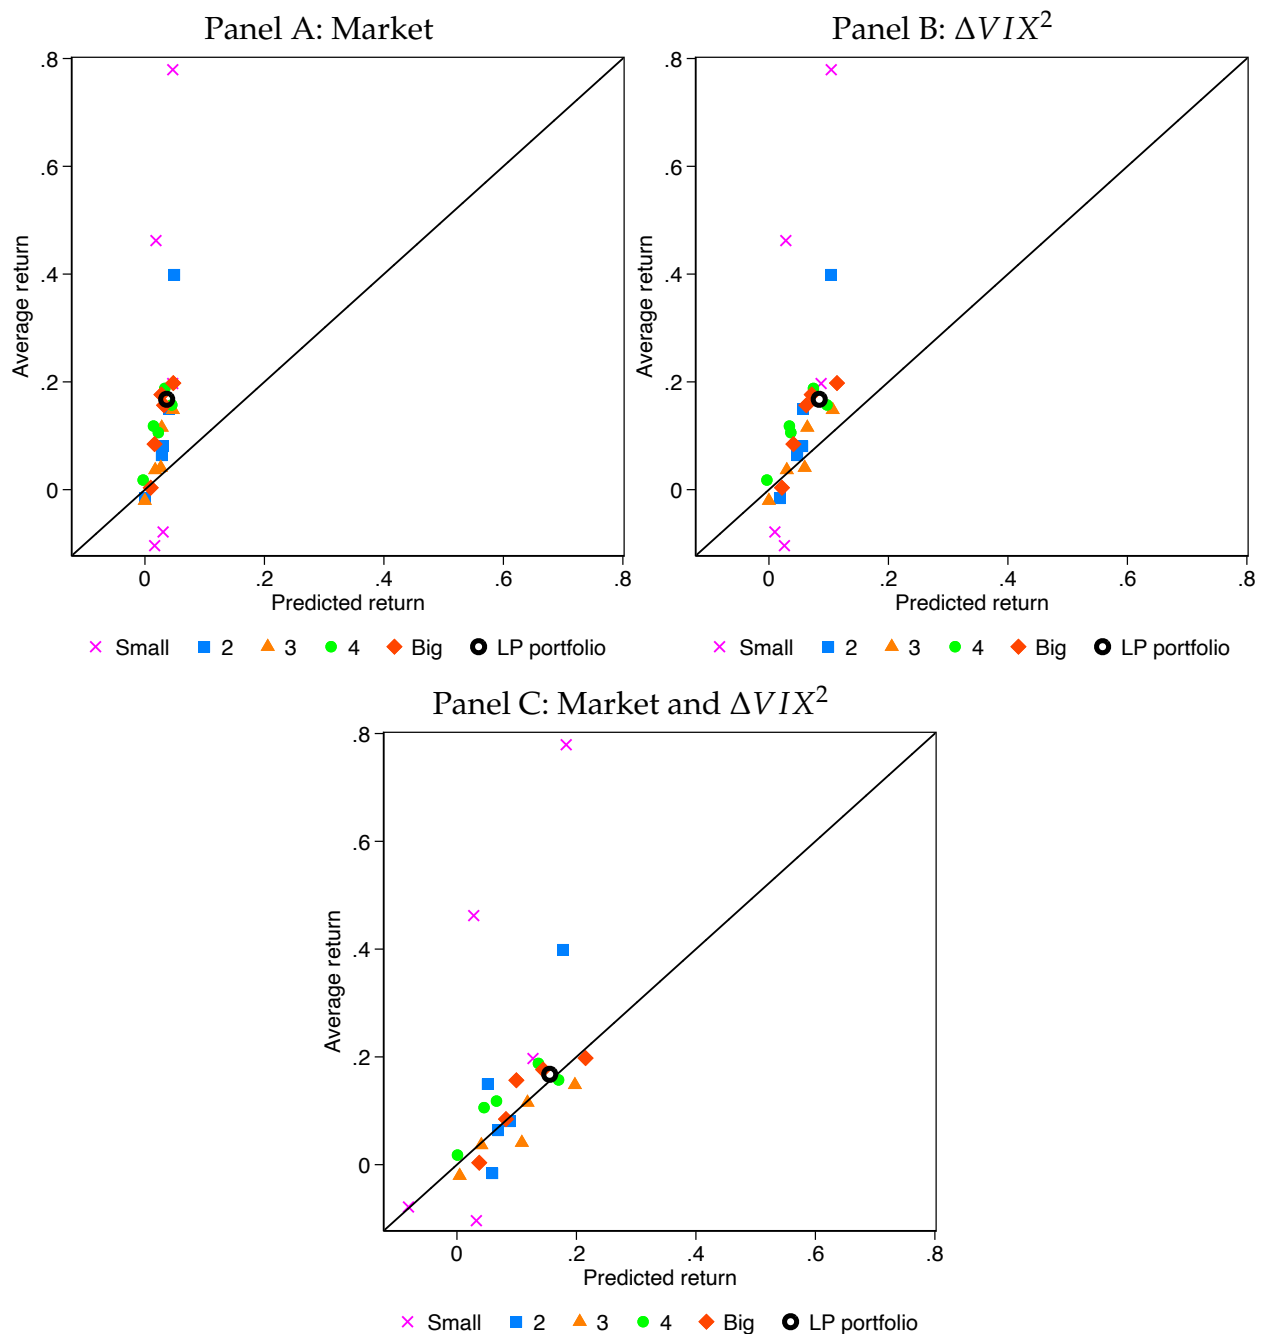
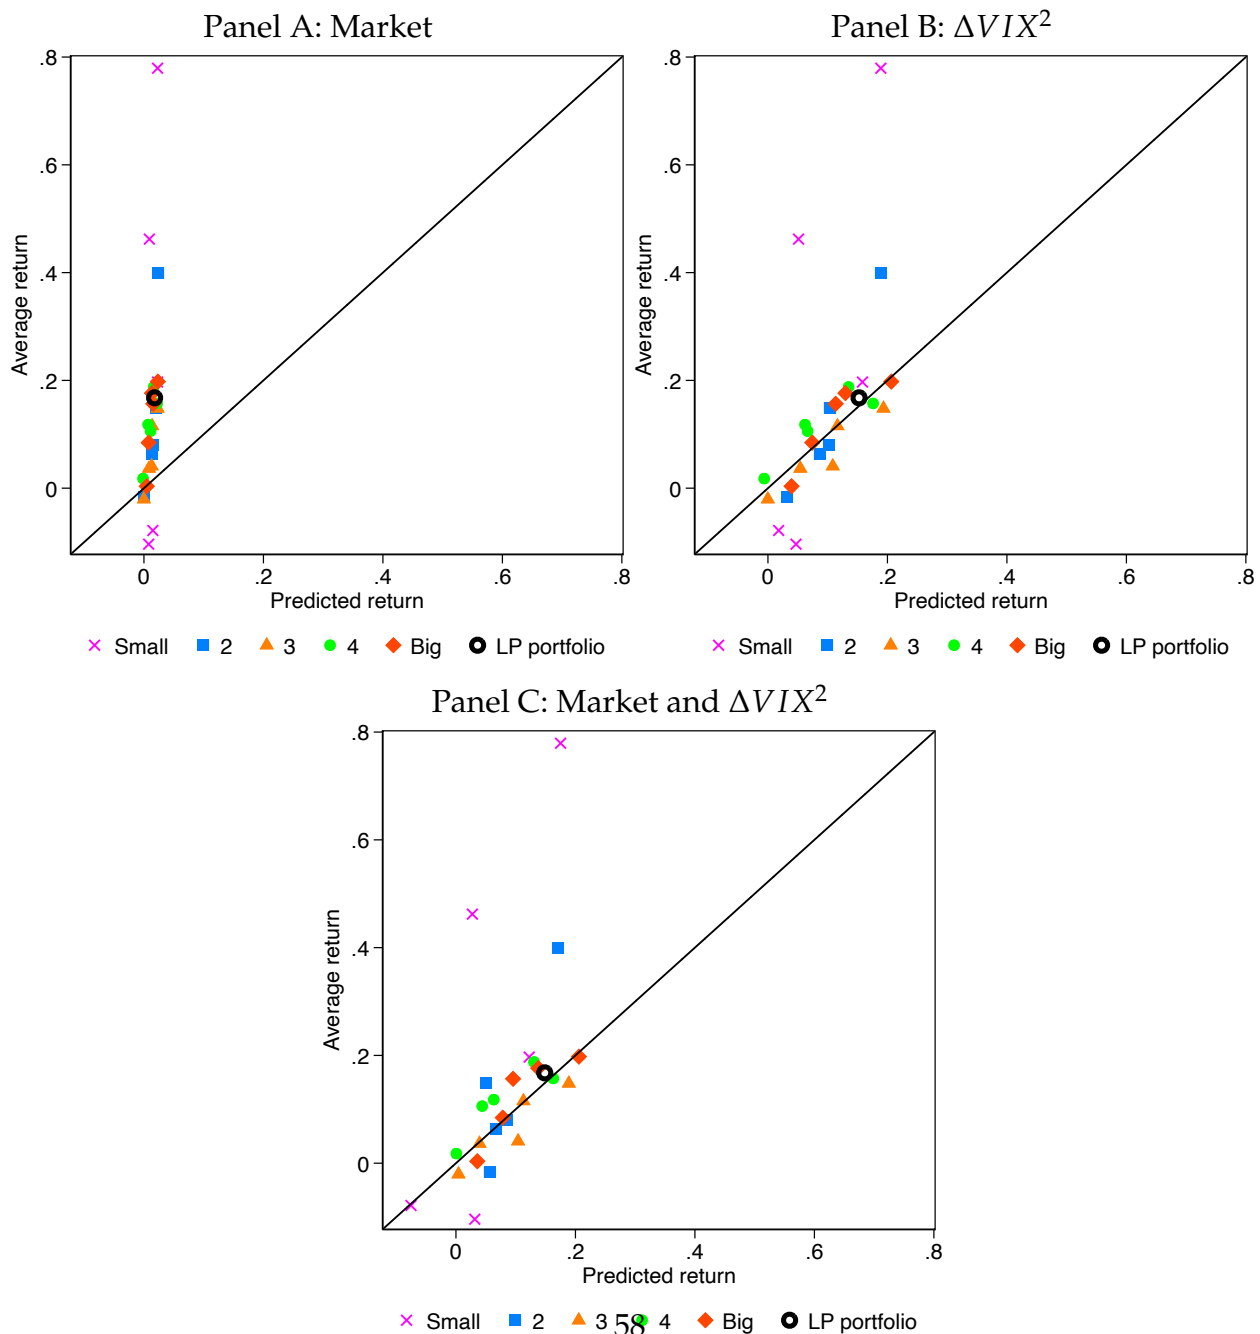


Figure 7: Option-implied price of volatility risk

The figure shows the average returns of the reversal strategies against their predicted returns using a restricted price of volatility risk obtained from option markets. The restricted price of volatility risk is the one that prices the VIX return. To calculate the predicted returns of the reversal strategies, we multiply their ΔVIX^2 betas by the restricted price of volatility risk. Each observation corresponds to a given reversal strategy. Different colors and markers distinguish the five size quintiles. Within each quintile, there are five long-short reversal strategies (Lo-Hi, 2-9, 3-8, 4-7, and 5-6). Each strategy is held for five trading days. Panel A uses the market return as the only factor. Panel B replaces it with the change in the squared VIX index. Panel C includes both factors. The sample is from April 9, 2001 to May 31, 2020.



Internet Appendix for

“Liquidity and Volatility”

IA.1 Data

This section provides additional details on how we construct our sample.

Stock screens: Our main data is from CRSP. We restrict the sample to ordinary common shares (share codes 10 and 11) listed on NYSE, NASDAQ, and AMEX. We exclude penny stocks and micro-caps by removing observations with a share price below the 20th percentile.²⁶ We also exclude stocks that are within one day of an earnings announcement as in Collin-Dufresne and Daniel (2014).²⁷

Portfolio formation: We form ten decile portfolios within each size quintile by sorting stocks by their normalized return (see Prediction 1). To compute a stock’s normalized return, we first market-beta-adjust it and then normalize it by dividing it by its standard deviation over a 60-day rolling window. Market-beta adjusting the returns removes the influence of market movements on the composition of the reversal portfolios. Since market movements are a form of public news, this makes the portfolios better proxies for the returns to liquidity provision, as discussed in Section IA.2.

Within each decile portfolio we weight stocks by their dollar volume. As discussed in Prediction 1, this further improves the mapping to our model. Since volume is unsigned, it serves as a rough proxy for the variance of order flow, which captures the scale of liquidity provision in a stock. We compute dollar volume over the same 60-day rolling window as the standard deviation.

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 one-month T-Bill). We obtain the VIX index from CBOE. The VIX index is a model-free measure of the implied volatility of the S&P 500 at a 30-day horizon (as proposed by Britten-Jones and Neuberger, 2000). Specifically, the squared VIX index is the price of a basket of options whose payoff replicates the realized variance of the S&P 500 over the following 30 calendar days. It therefore maps closely to the risk-adjusted expected market variance in our model, $E_t^Q [\sigma_m^2]$. In particular, following Propositions 2 and 3, and as discussed in Prediction 2, we use changes in the squared VIX (divided by 100 for legibility) as our variance risk factor.

The VIX return: To measure the premium for VIX exposure, it is necessary to calculate the VIX return from underlying options data rather than simply use the percentage change in the VIX index because the VIX basket changes each day (to keep its horizon constant). The percentage change in VIX is therefore not an investable trading strategy.

²⁶We use a relative cutoff because many stocks (e.g., Citigroup) fell below \$5 during the 2008 financial crisis. Our results are robust to an absolute cutoff of \$5, as is commonly used.

²⁷Earnings announcements are public-news events and as shown in Internet Appendix IA.2 this introduces measurement error when we use reversals as a proxy for the returns to liquidity provision. Consistent with this, the literature on the post-earnings announcement drift shows that earnings announcements do not produce reversals but continuation (e.g. Bernard and Thomas, 1989).

We solve this problem by calculating the return to holding the same VIX basket from one day to the next. This is an investable strategy and hence a valid return; we call it the VIX return.

To construct the VIX return, we first replicate the VIX index by reconstructing the VIX basket from the OptionMetrics data and following the methodology provided by the CBOE.²⁸ The replication is very accurate: our replicated VIX has a 99.83% correlation with the official VIX published by CBOE. The VIX return is the daily percentage change in the price of the basket used to construct the replicated VIX.

IA.2 Public news

In the main version of the model only informed traders receive signals about final payoffs ahead of time. This is common in the literature (e.g., Kyle, 1985) but unrealistic in practice as there are also many instances of public news, such as earnings announcements. In this section, we therefore expand the model to incorporate public news.

Prices at the final date 1 are now given by

$$p_{i,1} = \bar{v}_i + v_i + u_i, \quad (\text{IA.1})$$

where u_i is the component of the final payout about which there is public news, and is independent of v_i . Since the news about u_i is received by all market participants, they share the same time series of expectations of its value, $E_t[u_i]$. We allow public news to incorporate systematic risk that is priced. The risk pricing is captured by the risk-adjusted expectation of u_i , $E_t^Q[u_i] = E_t\left[\frac{\Lambda_T}{\Lambda_t}u_i\right]$. We normalize $E^Q[u_i] = 0$ so that the price of asset i prior to date 0 continues to be \bar{v}_i . Proposition IA.1 extends the results in the baseline model to account for public news.

Proposition IA.1. *The model with public news implies the following results:*

i. *The price of asset i on date $t \in \{0, \tau\}$ is given by*

$$p_{i,t} = \bar{v}_i + E_t^Q[u_i] + \frac{\phi E_t^Q[\sigma_{v,i}^2]}{\sigma_{x,i}^2} x_i. \quad (\text{IA.2})$$

ii. *The position of liquidity providers in asset i , $-x_i$, is proportional to the date-0 decline in the price of the asset net of the public news component:*

$$-x_i = -\frac{\sigma_{x,i}^2}{\phi} \left(\frac{\Delta p_{i,0} - E_0^Q[u_i]}{E_0^Q[\sigma_{v,i}^2]} \right). \quad (\text{IA.3})$$

²⁸See the CBOE white paper at <https://www.cboe.com/micro/vix/vixwhite.pdf>. Since October 2014, the CBOE uses weekly options as well as the traditional monthly ones. We follow their approach exactly.

iii. The change in the value of liquidity providers' position in asset i between dates 0 and τ is

$$-\Delta p_{i,\tau} x_i = -\frac{\phi x_i^2}{\sigma_{x,i}^2} \left(E_\tau^Q [\sigma_{v,i}^2] - E_0^Q [\sigma_{v,i}^2] \right) - x_i \left(E_\tau^Q [u_i] - E_0^Q [u_i] \right). \quad (\text{IA.4})$$

As the number of assets in the liquidity provider's portfolio becomes large ($N \rightarrow \infty$), $\sum_{i=1}^N x_i \left(E_\tau^Q [u_i] - E_0^Q [u_i] \right) \rightarrow 0$, i.e., the public news component is diversified out of the value of the liquidity providers' portfolios.

iv. The beta of liquidity providers' position in asset i to expected market volatility $E_\tau [\sigma_m^2]$ is

$$\beta_{i,\sigma_m} = -\frac{\phi k x_i^2}{\sigma_{x,i}^2} - x_i \beta_{u_i,\sigma_m} = -\left(\frac{\Delta p_{i,0} - E_0^Q [u_i]}{E_0^Q [\sigma_{v,i}^2]} \right)^2 \frac{k \sigma_{x,i}^2}{\phi} - x_i \beta_{u_i,\sigma_m}, \quad (\text{IA.5})$$

where β_{u_i,σ_m} is the market volatility beta of the public news component of asset i . As the number of assets N in the liquidity providers' portfolio grows large ($N \rightarrow \infty$),

$$\sum_{i=1}^N -x_i \beta_{u_i,\sigma_m} \rightarrow 0, \quad (\text{IA.6})$$

i.e., public news does not affect the portfolio's market volatility beta.

v. The expected payoff on liquidity providers' portfolios from date 0 to τ is:

$$E_0 \left[\sum_{i=1}^N -\Delta p_{i,\tau} x_i \right] = \left(\sum_{i=1}^N \beta_{i,\sigma_m} \right) \left(E_0 \left[E_\tau^Q [\sigma_m^2] \right] - E_0^Q [\sigma_m^2] \right) > 0. \quad (\text{IA.7})$$

Thus, the liquidity premium is positive and proportional to the variance risk premium.

Part (i) of Proposition IA.1 shows that prices now reflect public news in addition to the private news contained in order flow. Yet, public news does not change order flow because it is observable by everyone. Part (ii) shows that the positions of liquidity providers are proportional to price changes net of this public news component. To the extent that we cannot separately identify the public news component, it makes reversals a noisy proxy for liquidity providers' portfolios. To address this issue in our empirical analysis we exclude earnings announcements, which are the most prominent example of firm-level public news.

Part (iii) of Proposition IA.1 shows that liquidity providers' position in each asset is given by a short variance swap as before, together with a directional bet on the public news component. The directional bet stems from the long or short position that liquidity providers have in the asset. However, since they are equally likely to be long or short, their average position is zero and hence these directional bets tend to cancel out at the portfolio level. Thus, by the law of large numbers liquidity providers' exposure to public news becomes negligible as the number of assets in the portfolio grows large: $(\sum_{i=1}^N x_i (E_\tau^Q [u_i] - E_0^Q [u_i]) \rightarrow 0)$. In contrast, the liquidity provider has a negative ex-

posure to market volatility in every position, so this exposure does *not* cancel out at all. Thus, there is no change in the model's central prediction, that liquidity providers have an unambiguously negative market volatility beta.

Indeed, part (iv) shows that for a large number of assets, the market volatility beta of the liquidity provider's portfolio is exactly the same function of order flow as in the model without public news. Part (v) then shows that the liquidity premium charged by liquidity providers is, as before, the product of their portfolio's market volatility beta and the market variance risk premium. Therefore, although public news introduces measurement error into our tests, all of the key predictions of the model remain the same.

IA.3 Proofs

Proof of Proposition 1. Liquidity providers' information set is $\mathcal{F}_0 = \{x_i, \sigma_{x,i}^2 : i = 1, \dots, N\}$ on date 0 and $\mathcal{F}_\tau = \{E_\tau[\sigma_m^2], \mathcal{F}_0\}$ on date τ . From (3), the price of asset i is

$$p_{i,t} = E_t^Q[p_{i,1}] = \bar{v}_i + E^Q[v_i | \mathcal{F}_t]. \quad (\text{IA.8})$$

Using the law of iterated expectations to condition on the value of $\sigma_{v,i}^2$, and then using the joint normality of y_i and z_i given $\sigma_{v,i}^2$, we can use standard joint normal filtering to give the liquidity provider's updated expectation of v_i in terms of the net demand x_i he observes:

$$E^Q[v_i | \mathcal{F}_t] = E^Q\left[E[v_i | \sigma_{v,i}^2, \mathcal{F}_t] \middle| \mathcal{F}_t\right] = E^Q\left[\frac{\phi \sigma_{v,i}^2}{\sigma_{x,i}^2} x_i \middle| \mathcal{F}_t\right] = \frac{\phi E_t^Q[\sigma_{v,i}^2]}{\sigma_{x,i}^2} x_i. \quad (\text{IA.9})$$

The last equality follows from the fact that the liquidity provider knows x_i and $\sigma_{x,i}^2$. The first equality uses the fact that, given $\sigma_{v,i}^2$, v_i is idiosyncratic and is therefore orthogonal to the aggregate stochastic discount factor Λ_t corresponding to the Q measure. Substituting (IA.9) into (IA.8) gives (8). \square

Proof of Lemma 1. The result is obtained by re-arranging (8) as applied to date 0, solving for x_i , and substituting $\Delta p_{i,1} = p_{i,0} - \bar{v}_i$. \square

Proof of Lemma 2. The result is obtained by applying Proposition 1 to dates τ and 0 and taking the difference. \square

Proof of Lemma 3. The result is obtained by multiplying (10) by $-x_i$. \square

Proof of Proposition 2. The first equality is the definition of a beta. The second equality

substitutes in Lemma 3 as follows:

$$\beta_{i,\sigma_m} = \frac{\text{Cov} \left(-\Delta p_{i,\tau} x_i, E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)}{\text{Var} \left(E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)} \quad (\text{IA.10})$$

$$= \frac{\text{Cov} \left(-\frac{\phi x_i^2}{\sigma_{x,i}^2} \left(E_\tau^Q [\sigma_{v,i}^2] - E_0^Q [\sigma_{v,i}^2] \right), E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)}{\text{Var} \left(E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)} \quad (\text{IA.11})$$

$$= -\frac{\phi x_i^2}{\sigma_{x,i}^2} \frac{\text{Cov} \left(E_\tau^Q [k\sigma_m^2 + \varepsilon_v + \varsigma_{v,i}^2] - E_0^Q [k\sigma_m^2 + \varepsilon_v + \varsigma_{v,i}^2], E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)}{\text{Var} \left(E_\tau^Q [\sigma_m^2] - E_0^Q [\sigma_m^2] \right)} \quad (\text{IA.12})$$

$$= -\frac{\phi k x_i^2}{\sigma_{x,i}^2}. \quad (\text{IA.13})$$

This gives (13). The second to last equality uses the factor structure of volatility as given by (12). The last equality uses the fact that $\varsigma_{v,i}^2$ is idiosyncratic and hence uncorrelated with σ_m^2 . Finally, substitute (9) into (13) to get to get (14). \square

Proof of Proposition 3. The first step is to write the realized payoff of liquidity providers by average across assets in Lemma 3: The expected payoff of liquidity providers' portfolios from date 0 to date t is:

$$\sum_{i=1}^N -\Delta p_{i,1} x_i = \sum_{i=1}^N -\frac{\phi x_i^2}{\sigma_{x,i}^2} \left(E_\tau^Q [\sigma_{v,i}^2] - E_0^Q [\sigma_{v,i}^2] \right). \quad (\text{IA.14})$$

To get the expected payoff, take the realized payoff's date-0 expected value under the objective measure:

$$E_0 \left[\sum_{i=1}^N -\Delta p_{i,1} x_i \right] = E_0 \left[\sum_{i=1}^N -\frac{\phi x_i^2}{\sigma_{x,i}^2} \left(E_\tau^Q [\sigma_{v,i}^2] - E_0^Q [\sigma_{v,i}^2] \right) \right] \quad (\text{IA.15})$$

$$= \sum_{i=1}^N -\frac{\phi x_i^2}{\sigma_{x,i}^2} E_0 \left[E_\tau^Q [k\sigma_m^2 + \varepsilon_v + \varsigma_{v,i}^2] - E_0^Q [k\sigma_m^2 + \varepsilon_v + \varsigma_{v,i}^2] \right] \quad (\text{IA.16})$$

$$= \sum_{i=1}^N -\frac{\phi k x_i^2}{\sigma_{x,i}^2} \left(E_0 \left[E_\tau^Q [\sigma_m^2] \right] - E_0^Q [\sigma_m^2] \right). \quad (\text{IA.17})$$

The last equality uses the fact that $\varsigma_{v,i}^2$ and ε_v are orthogonal to our pricing measure and hence $E_t^Q [\varsigma_{v,i}^2 + \varepsilon_v]$ is a martingale under the objective measure. \square

Proof of Lemma 3'. Re-arranging (17) as applied to $t = 0$, the position of liquidity providers

in asset i is

$$-x_i = -\frac{\Delta p_{i,0}}{\gamma_{i,0} + \frac{\phi E_0^Q[\sigma_{v,i}^2]}{\sigma_{x,i}^2}}. \quad (\text{IA.18})$$

The result follows from differencing (17) between dates τ and 0 and multiplying by $-x_i$. \square

Proof of Proposition 2'. The results follow from taking the covariance of the right side of (18) with $E_t^Q[\sigma_m^2] - E_0^Q[\sigma_m^2]$ and dividing by $\text{Var}\left(E_t^Q[\sigma_m^2] - E_0^Q[\sigma_m^2]\right)$ for $t \in \{\tau, 1\}$. In the cases $\gamma_{i,t} = \gamma E_t^Q[\sigma_m^2]$ or $\gamma_{i,t} = \gamma E_t^Q[\sigma_{v,i}^2]$ the term $\beta_{\gamma_i, \sigma_m}$ simplifies to γ . \square

Proof of Proposition 3'. This result follows by taking the expected value of (18) as applied to $t \in \{\tau, 1\}$ and simplifying using Proposition 2'. \square

Proof of Proposition IA.1. The proof largely follows the proofs of Propositions 1–3 and Lemmas 1–3:

- i.* The information set of liquidity providers is $\mathcal{F}_0 = \{x_i, \sigma_{x,i}^2, E_0^Q[u_i] : i = 1, \dots, N\}$ on date 0 and $\mathcal{F}_\tau = \{E_\tau^Q[u_i], \sigma_m^2, \mathcal{F}_0\}$ on date τ . From (3), the price of asset i is

$$p_{i,t} = E_t^Q[p_{i,1}] = \bar{v}_i + E_t^Q[u_i] + E^Q[v_i | \mathcal{F}_t]. \quad (\text{IA.19})$$

Applying the law of iterated expectations as in (IA.9),

$$E^Q[v_i | \mathcal{F}_t] = E^Q\left[E[v_i | \sigma_{v,i}^2, \mathcal{F}_t] | \mathcal{F}_t\right] = E^Q\left[\frac{\phi \sigma_{v,i}^2}{\sigma_{x,i}^2} x_i | \mathcal{F}_t\right] = \frac{\phi E_t^Q[\sigma_{v,i}^2]}{\sigma_{x,i}^2} x_i. \quad (\text{IA.20})$$

The last equality uses the fact that v_i is idiosyncratic and hence independent of $E_t^Q[u_i]$ conditional on $\sigma_{v,i}^2$. Plugging (IA.20) into (IA.19) gives (IA.2).

- ii.* Re-arranging (IA.2) to solve for $-x_i$ gives (IA.3).
- iii.* Differencing (IA.2) between dates τ and 0 and multiplying by $-x_i$ gives (IA.4).
- iv.* The market volatility beta of asset i 's public news component is given by

$$\beta_{u_i, \sigma_m} = \frac{\text{Cov}\left(E_\tau^Q[u_i] - E_0^Q[u_i], E_\tau^Q[\sigma_m^2] - E_0^Q[\sigma_m^2]\right)}{\text{Var}\left(E_\tau^Q[\sigma_m^2] - E_0^Q[\sigma_m^2]\right)}. \quad (\text{IA.21})$$

Taking the covariance of the right side of (IA.4) with $E_\tau^Q[\sigma_m^2] - E_0^Q[\sigma_m^2]$ and dividing by $\text{Var}\left(E_\tau^Q[\sigma_m^2] - E_0^Q[\sigma_m^2]\right)$ gives (IA.5). Next, note that since v_i and z_i are

idiosyncratic, x_i is independent of β_{u_i, σ_m} and so $\sum_{i=1}^N x_i \beta_{u_i, \sigma_m} = 0$.

v. This result follows the proof of Proposition 3.

□

IA.4 Idiosyncratic volatility and market volatility

Table IA.1 formally assesses the relationship between idiosyncratic and market volatility shown in Figure 2. In order to provide a benchmark, column (1) regresses realized market volatility over 21 trading days on VIX as of the start of the period. The R^2 is 55.2%, showing that VIX is a powerful predictor of market volatility, as expected. The coefficient is slightly less than one, 0.924, reflecting the fact that some of the variation in VIX is driven by changes in the variance premium rather than expected variance.

Column (2) replaces market volatility with idiosyncratic volatility from Figure 2. The R^2 is almost identical, 55.5%, and the coefficient is 0.969. Thus, VIX is as good at predicting idiosyncratic volatility as it is market volatility. This is remarkable given the fact that VIX is constructed to predict market volatility. In column (3), we replace VIX with the contemporaneous realized market volatility. The R^2 is even higher at 84.4%, implying a correlation of 92%. Thus, realized idiosyncratic volatility and market volatility move practically in lockstep.

Columns (4) to (6) repeat columns (1) to (3) in terms of variances instead of volatilities. The same results emerge: market variance and idiosyncratic variance are almost perfectly correlated (93%), and VIX squared is an equally powerful predictor of both.

IA.5 Volatility co-movement portfolio sorts

In this section we extend our empirical analysis by forming portfolios based on differences in volatility co-movement between stocks. Our model predicts that reversal returns and volatility risk betas should be larger for stocks with greater volatility co-movement. A rise in VIX induces a greater increase in the volatility of these stocks, which reveals that there is more private information about their values. This increases the risk of providing liquidity in these stocks.²⁹

We estimate the volatility co-movement of each stock by running time series regressions of its idiosyncratic volatility on VIX. Idiosyncratic volatility is computed as the standard deviation of market-adjusted returns over a five-day rolling window (the horizon of our portfolios). We then regress it on VIX as of the start of the window:

$$\hat{\sigma}_{t,t+5}^i = a + k_i VIX_t + \epsilon_{i,t}. \quad (\text{IA.22})$$

To make sure we do not introduce look-ahead bias, we run these regressions on one year of past data at each point in time. This interval is sufficient for creating meaningful variation in ex-post volatility co-movement (see Table IA.3 below). The next step is to sort

²⁹Formally, we can generalize Eq. (5) by allowing for heterogeneous loadings: $\sigma_{v,i}^2 = k_i \sigma_v^2 + \zeta_{v,i}^2$, $k_i > 0$. The volatility betas become $\beta_{i, \sigma_m} = -\phi k_i (x_i^2 / \sigma_x^2)$, hence they are more negative for assets with a larger k_i . It follows from Proposition 3 that the expected reversal return is correspondingly larger.

stocks into quintiles by their coefficients k_i and then deciles by their normalized returns. The portfolios are again weighted by dollar volume. They are thus analogous to our main portfolios in Section 3 but with the k_i 's in place of size.

Table IA.3 presents summary statistics on the volatility co-movement (k -sorted) portfolios. Market capitalization (Panel A), idiosyncratic volatility (Panel B), average turnover (Panel C), and sorting-day returns (Panel D) are fairly similar across the quintiles sorted by the co-movement coefficients k . By construction, the pre-sorting coefficients k are strongly increasing across quintiles (Panel E). Importantly, the substantial spread is also evident in the post-sorting coefficients (Panel F). This validates our empirical approach and allows us to conduct an economically meaningful test of the model. Moreover, consistent with the model and Figure 2, the ex-post co-movement coefficients are all positive, indicating that higher market volatility is associated with higher idiosyncratic volatility.

Table IA.4 shows the average returns of the k -sorted reversal strategies. Panel A shows a strong pattern in average reversal returns across quintiles. As predicted by our model, portfolios of stocks whose idiosyncratic volatility co-moves more strongly with aggregate volatility earn substantially higher reversal returns. Average reversal returns for the Lo–Hi strategies increase monotonically from 13 bps to 29 bps from the first quintile to the fifth. Panels E and F show that the corresponding CAPM alphas are very similar to the raw returns and statistically significant.

Table IA.5 looks at the volatility risk betas. From Panel A, the betas become larger (more negative) as we go from low to high co-movement quintiles. The beta of the Lo–Hi strategy in the first quintile is -0.13 , while that of the fifth quintile is -0.27 . Both are significant (Panel B). The same is true when we control for market exposure (Panels C and D). By contrast, Panels E and F show no pattern in the strategies' market betas, which are economically small. Table IA.5 thus shows that stocks whose idiosyncratic volatility is more sensitive to aggregate volatility have reversal returns that are more exposed to volatility risk, which is consistent with our model.

Table IA.6 shows the results of Fama-Macbeth regressions with an option-implied price of risk. These results are analogous to Table 10 for our main portfolio sorts. From Table IA.6, controlling for volatility risk eliminates the pricing errors of the Lo–Hi strategies. The pricing error of the high co-movement quintile five drops from 26 bps under the CAPM to 1 bps in the one-factor model with volatility risk and -5 bps in the two-factor model. Volatility risk thus again explains the liquidity premium in the cross section of stocks.

Figure IA.1 similarly shows that accounting for volatility risk aligns the average returns of the reversal strategies with their predicted returns. In particular, the overall liquidity provider portfolio lies almost exactly on the 45-degree line. Volatility risk thus again explains the liquidity premium in the cross section of stocks.

Table IA.1: Idiosyncratic volatility, market volatility, and VIX

This table shows the relationship between idiosyncratic volatility, market volatility, and VIX, as well as their variance counterparts. Idiosyncratic volatility and variance are calculated as follows. Each day, we compute the beta-adjusted returns of all stocks using a 60-day rolling window to estimate the betas. We then square these returns and value-weight them across stocks. Idiosyncratic variance is the annualized sum of these squared value-weighted returns over the next 21 trading days (to match the 30-calendar-day horizon of VIX). Idiosyncratic volatility is the square root of idiosyncratic variance. Market variance is the annualized sum of the squared market returns over the next 21 trading days and market volatility is its square root. The sample is from April 9, 2001 to May 31, 2020.

	Mkt. vol.	Idio. vol.		Mkt. var.	Idio. var.	
	(1)	(2)	(3)	(4)	(5)	(6)
VIX	0.924*** (0.012)	0.969*** (0.013)				
Mkt. vol.			0.961*** (0.006)			
VIX ²				0.913*** (0.016)	1.429*** (0.024)	
Mkt. var.						1.418*** (0.008)
Constant	-1.881*** (0.258)	10.555*** (0.269)	13.934*** (0.118)	-0.338*** (0.120)	3.463*** (0.181)	4.550*** (0.076)
Obs.	4,794	4,814	4,794	4,794	4,814	4,794
R ²	0.552	0.555	0.844	0.397	0.417	0.860

Table IA.2: Predicting reversal returns with an estimated VRP

Results from predictability regressions of the reversal strategy returns on an estimated variance risk premium. The Lo–Hi strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within a given size quintile. The remaining strategies are constructed analogously for the inner normalized return deciles. Each strategy is held for five trading days. The predictability regressions are

$$R_{t,t+5}^p = a_r^p + b_r^p \widehat{VRP}_t + \epsilon_{r,t+5}^p$$

where $R_{t,t+5}^p$ is the cumulative excess return on portfolio p from the portfolio formation date t to $t + 5$ and \widehat{VRP}_t is the estimated variance risk premium. The predictive loadings b_r^p are multiplied by 100 for legibility. We construct the estimated VRP as follows. Let $RV_{t,t+21}$ be the 21-day forward realized variance of the S&P 500. We run a regression in logs of this forward realized variance on VIX squared and the 21-day backward realized variance:

$$\log(RV_{t,t+21}) = a + b \log(VIX_t^2) + c \log(RV_{t-21,t}) + \epsilon_{t,t+21}$$

The log specification ensures that the predicted realized variance is positive. The estimated VRP is the difference between VIX squared and exponentiated predicted realized variance:

$$\widehat{VRP}_t = VIX_t^2 - \exp \left\{ \hat{a} + \hat{b} \log(VIX_t^2) + \hat{c} \log(RV_{t-21,t}) + 0.5 \hat{\sigma}^2(\epsilon_{t,t+21}) \right\}$$

The t -statistics are based on Newey–West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Predictive loadings						Panel B: t -statistics					
	Lo–Hi	2–9	3–8	4–7	5–6		Lo–Hi	2–9	3–8	4–7	5–6
Small	4.79	3.99	3.77	−0.31	0.08	Small	4.24	3.77	3.74	−0.26	0.07
2	3.25	1.32	1.48	0.71	−0.22	2	3.11	2.02	2.23	1.17	−0.38
3	2.62	1.37	1.40	1.22	0.18	3	2.44	2.00	2.83	2.06	0.57
4	2.30	2.06	1.35	1.60	−0.20	4	2.67	2.64	2.70	2.92	−0.61
Big	4.01	2.35	1.68	1.13	0.16	Big	4.15	3.81	2.82	2.19	0.53

Panel C: R^2					
	Lo–Hi	2–9	3–8	4–7	5–6
Small	0.95	0.83	0.72	0.01	0.00
2	1.10	0.24	0.33	0.08	0.01
3	0.99	0.41	0.44	0.34	0.01
4	0.92	1.05	0.62	1.06	0.02
Big	3.15	1.66	1.15	0.69	0.02

Table IA.3: Summary statistics: k -sorted portfolios

This table shows summary statistics for reversal strategies formed by quintiles of k , the sensitivity of a stock's idiosyncratic volatility to market volatility. It is estimated by running

$$\sigma_{i,t} = k_{0,i} + k_i VIX_t + \epsilon_{i,t}^\sigma,$$

where $\sigma_{i,t}$ is the idiosyncratic volatility of stock i , measured from stock i 's beta-adjusted returns over five trading days following date t . The regression is estimated using a one-year rolling window. Each day, stocks are first sorted into quintiles by k and then deciles by normalized beta-adjusted return. The normalized return is calculated using a 60-day rolling window. The portfolios are weighted by average dollar volume over that window. Post-sorting k 's are estimated by regressing the weighted average idiosyncratic volatility of the stocks in each portfolio (taken over the five trading days following portfolio formation) on VIX. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Market cap						Panel B: Idiosyncratic volatility					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	52.01	52.22	48.50	47.92	45.88	Low k	1.89	2.03	2.20	2.33	2.43
2	57.50	56.73	54.90	54.32	53.72	2	1.54	1.62	1.69	1.74	1.77
3	55.88	56.47	54.91	53.79	54.01	3	1.70	1.78	1.85	1.89	1.90
4	47.55	49.10	47.96	46.95	46.72	4	2.08	2.18	2.25	2.28	2.30
High k	30.74	34.21	33.94	34.75	32.90	High k	2.95	3.07	3.19	3.25	3.30

Panel C: Average turnover						Panel D: Sorting-day return					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	1.30	1.37	1.43	1.47	1.50	Low k	-6.00	-3.18	-2.03	-1.14	-0.37
2	0.96	0.99	1.03	1.05	1.07	2	-4.87	-2.67	-1.71	-0.97	-0.32
3	1.04	1.08	1.11	1.12	1.13	3	-5.27	-2.89	-1.85	-1.05	-0.34
4	1.34	1.39	1.42	1.43	1.44	4	-6.23	-3.41	-2.17	-1.22	-0.39
High k	2.00	2.09	2.13	2.18	2.19	High k	-8.57	-4.55	-2.88	-1.62	-0.53

Panel E: Pre-sorting k						Panel F: Post-sorting k					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	-0.39	-0.41	-0.44	-0.48	-0.49	Low k	0.49	0.45	0.45	0.45	0.46
2	0.26	0.26	0.26	0.26	0.26	2	0.79	0.74	0.73	0.71	0.72
3	0.68	0.68	0.68	0.68	0.68	3	1.17	1.07	1.04	1.05	1.04
4	1.20	1.20	1.20	1.20	1.20	4	1.68	1.50	1.45	1.47	1.45
High k	2.39	2.41	2.44	2.45	2.46	High k	2.51	2.06	2.04	2.08	2.06

Table IA.4: Reversal strategy returns: k -sorted portfolios

Average returns, standard deviations, Sharpe ratios, and CAPM alphas of the five-day reversal strategies formed by quintiles of k , the sensitivity of a stock's idiosyncratic volatility to market volatility. Each day, stocks are first sorted into quintiles by k and then deciles by normalized beta-adjusted return. The normalized return is calculated using a 60-day rolling window. The portfolios are weighted by average dollar volume over that window. The t -statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Average returns						Panel B: t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	0.13	0.13	0.10	0.03	0.04	Low k	2.46	3.17	2.57	0.71	1.18
2	0.16	0.08	0.06	0.07	0.02	2	3.62	2.53	1.97	2.66	1.01
3	0.16	0.16	0.16	0.08	0.01	3	3.68	4.52	5.33	2.79	0.40
4	0.18	0.24	0.17	0.13	0.04	4	3.29	5.43	4.62	3.63	1.21
High k	0.29	0.30	0.18	0.14	-0.05	High k	3.00	5.01	3.45	2.91	-1.11

Panel C: Standard deviations						Panel D: Sharpe ratios					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	3.34	2.88	2.55	2.69	2.45	Low k	0.28	0.32	0.27	0.07	0.12
2	2.66	2.14	1.97	1.88	1.70	2	0.42	0.26	0.20	0.26	0.10
3	2.91	2.40	2.06	1.96	1.82	3	0.39	0.47	0.54	0.29	0.04
4	3.55	2.99	2.58	2.46	2.43	4	0.37	0.56	0.47	0.38	0.12
High k	5.62	3.99	3.65	3.40	3.42	High k	0.37	0.54	0.34	0.29	-0.11

Panel E: CAPM alphas						Panel F: CAPM alpha t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	0.10	0.11	0.08	0.02	0.04	Low k	2.00	2.73	2.14	0.47	1.21
2	0.13	0.06	0.04	0.06	0.02	2	2.99	2.04	1.41	2.50	0.82
3	0.14	0.14	0.15	0.07	0.01	3	3.25	3.96	4.94	2.42	0.26
4	0.15	0.22	0.16	0.12	0.04	4	2.72	4.99	4.26	3.41	1.09
High k	0.26	0.29	0.16	0.13	-0.06	High k	2.67	4.74	3.17	2.70	-1.22

Table IA.5: Volatility risk of the k -sorted reversal strategies

The table shows the betas with respect to changes in VIX^2 of the five-day reversal strategies by quintiles of k , the sensitivity of a stock's idiosyncratic volatility to market volatility. The betas are estimated by running the regressions

$$R_{t,t+5}^p = \alpha_p + \beta_{VIX}^p \Delta VIX_{t,t+5}^2 + \beta_M^p R_{t,t+5}^M + \epsilon_{t,t+5}^p,$$

where $R_{t,t+5}^p$ is the cumulative excess return on reversal strategy portfolio p from the portfolio formation date t to $t + 5$, $R_{t,t+5}^M$ is the excess return on the market portfolio, and $\Delta VIX_{t,t+5}^2$ is the change in the squared VIX index from date t to $t + 5$. Panel A omits the market return while Panels B and C include it. Panel C reports the market betas β_M^p . The t -statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: ΔVIX^2 betas

β_{VIX}^p						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	-0.13	-0.10	-0.11	-0.08	0.00	Low k	-3.76	-3.65	-4.46	-1.88	0.02
2	-0.19	-0.09	-0.13	-0.01	-0.00	2	-4.46	-3.30	-4.95	-0.45	-0.01
3	-0.21	-0.15	-0.09	-0.06	-0.04	3	-4.91	-4.37	-3.91	-1.25	-2.54
4	-0.22	-0.15	-0.12	-0.05	-0.04	4	-4.15	-3.34	-5.35	-1.82	-1.22
High k	-0.27	-0.04	-0.07	-0.04	-0.01	High k	-3.77	-0.75	-1.41	-1.05	-0.45

Panel B: ΔVIX^2 betas (controlling for R^M)

β_{VIX}^p						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	-0.07	-0.07	-0.09	-0.09	-0.02	Low k	-1.42	-2.14	-2.84	-1.37	-0.59
2	-0.15	-0.06	-0.13	0.02	0.01	2	-2.78	-1.49	-3.29	0.68	0.57
3	-0.20	-0.12	-0.09	-0.05	-0.05	3	-3.23	-2.49	-2.80	-0.82	-2.36
4	-0.17	-0.17	-0.13	-0.01	-0.05	4	-2.46	-2.75	-4.06	-0.33	-1.28
High k	-0.34	-0.03	-0.07	-0.01	0.02	High k	-2.93	-0.35	-1.32	-0.15	0.35

Panel C: Market betas

β_M^p						t -statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Small	0.10	0.04	0.03	-0.01	-0.03	Small	2.28	1.19	1.19	-0.18	-1.08
2	0.05	0.05	0.00	0.04	0.02	2	1.41	1.32	0.09	2.00	0.76
3	0.02	0.04	-0.01	0.01	-0.01	3	0.39	1.30	-0.20	0.31	-0.74
4	0.08	-0.02	-0.01	0.06	-0.02	4	1.70	-0.60	-0.21	1.84	-0.64
Big	-0.10	0.02	-0.01	0.05	0.05	Big	-1.05	0.46	-0.18	1.15	1.15

Table IA.6: Option-implied price of volatility risk: k -sorted portfolios

The table uses a restricted price of volatility risk obtained from option markets to compute the pricing errors of the five-day reversal strategies formed by quintiles of k , the sensitivity of a stock's idiosyncratic volatility to market volatility. The restricted price of volatility risk is the one that prices the VIX return with ΔVIX^2 as a pricing factor. To obtain the pricing errors of the reversal strategies, we multiply their ΔVIX^2 betas by the restricted price of volatility risk and subtract the resulting predicted average return from the actual return of each strategy, which we then average over time. The Lo-Hi strategy goes long the lowest normalized return decile portfolio and short the highest normalized return decile portfolio within a given k quintile. The remaining strategies are constructed analogously for the inner normalized return deciles. Each strategy is held for five trading days. Panel A reports the pricing errors when the market return is used as the only factor. Panel B replaces it with the change in the squared VIXN index. Panel C includes both factors. The reported t statistics are based on Newey-West standard errors with five lags to account for the overlap in the returns. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Market

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	0.10	0.11	0.08	0.02	0.04	Low k	1.96	2.77	2.12	0.47	1.23
2	0.13	0.06	0.04	0.06	0.02	2	2.96	2.03	1.40	2.52	0.84
3	0.13	0.14	0.15	0.07	0.01	3	3.09	3.90	4.97	2.53	0.26
4	0.15	0.22	0.16	0.12	0.04	4	2.63	4.98	4.17	3.29	1.11
High k	0.26	0.29	0.16	0.13	-0.06	High k	2.67	4.83	3.23	2.76	-1.22

Panel B: ΔVIX^2

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	-0.01	0.03	-0.01	-0.06	0.04	Low k	-0.13	0.77	-0.36	-1.42	1.20
2	-0.04	-0.02	-0.08	0.06	0.02	2	-0.83	-0.57	-2.62	2.29	1.00
3	-0.06	0.01	0.06	0.02	-0.03	3	-1.35	0.20	2.15	0.74	-1.27
4	-0.04	0.08	0.04	0.08	0.00	4	-0.79	1.83	1.18	2.19	0.13
High k	0.01	0.26	0.11	0.10	-0.07	High k	0.08	4.33	2.11	2.09	-1.42

Panel C: Market and ΔVIX^2

Pricing errors						t statistics					
	Lo-Hi	2-9	3-8	4-7	5-6		Lo-Hi	2-9	3-8	4-7	5-6
Low k	0.04	0.05	0.00	-0.06	0.03	Low k	0.75	1.17	0.02	-1.49	0.71
2	-0.01	0.01	-0.08	0.08	0.03	2	-0.24	0.17	-2.63	3.20	1.22
3	-0.05	0.03	0.06	0.03	-0.04	3	-1.22	0.84	2.04	0.91	-1.54
4	-0.01	0.06	0.04	0.11	-0.01	4	-0.12	1.39	0.98	2.98	-0.19
High k	-0.05	0.27	0.10	0.12	-0.04	High k	-0.56	4.44	1.90	2.62	-0.90

Figure IA.1: Volatility co-movement portfolios

The figure shows the average returns of the volatility co-movement reversal strategies against their predicted returns using a restricted price of volatility risk from option markets. The reversal strategies are formed by quintiles of k , the sensitivity of a stock's idiosyncratic volatility to market volatility. The restricted price of volatility risk is the one that prices the VIX return. To calculate the predicted returns of the reversal strategies, we multiply their VIX-squared betas by the restricted price of volatility risk (and their market betas by the average market excess return). Each observation corresponds to a reversal strategy. Each color and marker corresponds to a k quintile. Panel A uses the market return as the only factor. Panel B replaces it with VIX-squared changes. Panel C includes both factors. The sample is from April 9, 2001 to May 31, 2020.

