Liquidity and Volatility

Itamar Drechsler,   Alan Moreira,   Alexi Savov[*]

October 2020

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. As aggregate volatility shocks carry a very large premium in option markets, this negative exposure can explain why liquidity provision earns high average returns. We show this by incorporating uncertainty about the amount of private information into an otherwise standard model. 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 their average returns with the same risk price as in option markets, and their predictability by VIX in the time series. 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

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 bear a negative exposure to market volatility. Since market volatility carries a large negative price of risk in option markets (Carr and Wu, 2008), the negative exposure of liquidity providers explains why liquidity provision earns a substantial premium (Krishnamurthy and Vissing-Jorgensen, 2015), why this premium increases with volatility (Nagel, 2012), and why liquidity contracts during crises (Brunnermeier, 2009). By integrating liquidity and volatility, our results provide a new perspective on the risks and returns to liquidity provision.

Why do liquidity providers have a negative exposure to market volatility? The basic problem liquidity providers face is adverse selection by investors with private information. In standard models (e.g. Kyle, 1985), the quantity of private information in the market is constant and known to liquidity providers. It therefore poses no risk to them. While this is a simple assumption, it is clearly unrealistic. Volatility, the observed end-product of the flow of information, is highly time-varying, both at the level of individual assets and for the market as a whole. Moreover, the volatilities of individual assets covary strongly with each other and with market volatility. This suggests that flows of information, both public and private, vary substantially over time and in a synchronous manner.

To capture this co-movement, we extend the model of Kyle (1985) by making the amounts of market and private information uncertain and positively correlated. Because of this correlation, market volatility informs liquidity providers about the amount of private information they face. In particular, an increase in market volatility leads them to estimate a greater amount of private information.

As in Kyle (1985), liquidity providers' estimate of private information determines how much they adjust prices in the direction of the net order flow they receive. If their estimate is too low, they adjust prices too little, and because of adverse selection trade too much with informed investors. They then accumulate long positions at prices that are too high and short positions at prices that are too low.

Now consider what happens when there is a positive shock to market volatility. The shock increases liquidity providers' estimate of private information. Liquidity providers learn that they had not adjusted prices enough and likely traded too much with informed investors. Prices immediately come to reflect this. Liquidity providers see their long

positions fall and their short positions rise, taking losses across the portfolio. The opposite happens when a negative shock hits: liquidity providers' estimate of private information declines, their longs rise and their shorts fall. In sum, liquidity providers have a negative exposure to market volatility.

In existing models of liquidity provision, which focus on inventory risk (e.g., Stoll, 1978; Grossman and Miller, 1988), a liquidity provider's risk is idiosyncratic. To explain the liquidity premium these models assume that some friction prevents the liquidity provider from diversifying this risk away. While this may apply in some cases, particularly for small and thinly-traded assets, it is unlikely to explain the liquidity premium among large and heavily-traded assets.

In our model liquidity providers face a systematic risk (fluctuations in market volatility) that cannot be diversified away. We thus need to specify how this risk is priced by the economy's stochastic discount factor. Consistent with the strong evidence from option prices, we assume that market volatility carries a large negative price of risk. This implies that liquidity providers earn a positive premium for the negative volatility risk exposure they bear. The premiums for liquidity and volatility are thus one and the same.

To test the model, we need to proxy for liquidity providers' portfolios, which are not directly observed. As in Nagel (2012), the model shows that they are mimicked by short-term reversals (e.g., Lehmann, 1990), a trading strategy that buys assets whose price has fallen and shorts assets whose price has risen. The reason for this is that a large price change reflects a large net order flow, requiring the liquidity provider to take a large position. We thus use reversals to proxy for the returns to liquidity provision.

We build reversal portfolios using daily U.S. stock returns from 2001 to 2020 (i.e., the period since "decimalization," when liquidity provision became competitive; see Bessembinder, 2003). Each day we sort stocks into quintiles by market cap 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. This gives us fifty portfolios sorted by size and return. We further construct long-short reversal strategies within each size quintile by buying the low-return deciles and shorting 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.

2

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

Figure 1 plots the average return of this strategy over rolling 60-day forward-looking windows against the VIX index, which measures the risk-adjusted expected volatility of the S&P 500 over the next month. The plot shows that VIX strongly predicts the reversal strategy's subsequent return: the two series have correlation of 46%. A predictive regression shows that a one-point increase in VIX predicts a 9-bps higher reversal return, a large increase relative to the strategy's 17-bps average return. These findings confirm the main finding of Nagel (2012) in a setting with only large stocks.

According to our model, this predictability is due to the VIX also predicting the reversal strategy's systematic volatility risk, which results from its negative beta to expected market volatility. We estimate this volatility risk by running 60-day rolling regressions of the reversal strategy's return on contemporaneous changes in the squared VIX index (we square VIX to convert it into a variance, following the model). We then take the annualized volatility of the regression's fitted value. The volatility of the fitted value gives the reversal strategy's systematic volatility risk, the component of its return that is due to its exposure to market volatility. As predicted by the model, the average beta in this regression is strongly negative: a one-point increase in VIX-squared is associated with a 20-bps contemporaneous drop in the reversal strategy's return.[1]

Panel B of Figure 1 plots the large-cap reversal strategy's systematic volatility risk against VIX. The figure shows that the two series track each other very closely, including during the 2008 financial crisis and the 2020 Covid crisis. This shows that liquidity provision is especially risky during crises. The relationship is also strong in normal times and the overall correlation is 58%. Importantly, the volatility risk of the reversal strategy has the same pattern as its average return in Panel A. Thus, the price of liquidity is highest when providing liquidity is riskiest, as predicted by our model.

Our main analysis tests the model's predictions on the full cross section of reversal

---

[1]This is not due to market beta, which is close to zero. Moreover, the VIX beta is unchanged when we control for the market return.

portfolios. We first show that the long-short 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 also confirm that these returns are not explained by market risk.

Next, we test the model's central prediction that reversal strategies' volatility risk (their VIX betas) can explain their average returns. We again estimate volatility betas by regressing each strategy's return on the contemporaneous change in VIX-squared (we also control for the market return to remove any market beta). We find that the volatility betas are strongly negative; they are largest for the Lo–Hi strategies (-0.2 in the case of large-caps) and converge toward zero for the 5–6 strategies. Thus, consistent with the model, the reversal strategies' volatility betas display the same pattern as their average returns. This is the cross-sectional analog of the time-series result in Figure 1.

We use Fama-MacBeth regressions to formally test whether volatility risk explains the reversal strategies' average returns. The estimated premium for the volatility betas is a large and significant -1.08% per unit of beta per five days. This number is consistent with the highly negative variance risk premium found in the prices of index options. The VIX-squared betas shrink the pricing errors of the portfolios dramatically. The pricing error of the large-cap Lo–Hi strategy drops from 17 bps to -2 bps and becomes statistically insignificant. Similarly, the second-largest quintile's Lo–Hi pricing error shrinks from 13 bps 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, appear to behave quite differently than the rest. Besides this exception, the reversal strategies' average returns are explained well by their volatility betas. In particular, the pricing error of an overall liquidity provider portfolio, which we construct by combining the reversal strategies as implied by the model, drops from 17 bps to 0.8 bps. Volatility risk thus explains the overall liquidity premium in the cross section of stocks.

Since market variance is traded directly in the index options market, we can use index option returns to estimate its price of risk and test if it explains the reversal returns. Following the literature, we do so by using the VIX squared, which equals the price of a basket of options that replicates an S&P 500 variance swap. From the returns on this basket, we estimate that the price per unit of VIX-squared beta in option markets is -1.03%, which is very close to our estimate from the equity reversal strategies.

We set the price of VIX-squared risk to this restricted value (and the premium for market risk to the market's average excess return) and calculate the reversal strategies' pricing errors.[2] The restricted pricing tests perform very similarly to the unrestricted Fama-MacBeth regressions, reflecting the similar price of risk estimates. For instance, the pricing errors of the largest and second-largest quintiles' Lo–Hi portfolios decrease from 0.17 bps and 13 bps to -1 bps. As before, only the pricing error of the smallest stocks' Lo–Hi portfolio remains large and significant. The pricing error of the overall liquidity portfolio drops from 17 bps to 1.5 bps.

Thus, small-stock reversals stand out from the others. While their VIX betas are negative, they are too small to explain their outsized returns. This suggests that undiversified inventory risk may play a role. To analyze this, we extend our model to incorporate inventory costs, as in the prior literature. Moreover, we allow the costs to increase with VIX, as suggested by Nagel (2012). This makes reversal returns sensitive to VIX, giving them a negative VIX beta. However, the model shows that this sensitivity is transitory: the impulse response to a VIX shock must shrink to zero by the time the liquidity provider disposes of the inventory. This is indeed what we find for the small-stock reversal strategy: their initial drop from a positive VIX shock reverts back to zero after a few days.

The pattern is very different for medium and large stocks. Their response to VIX shocks is permanent, with no sign of decay well beyond liquidity providers' holding periods. This result is implied by our model. A VIX shock that increases adverse selection has a permanent impact on liquidity providers' returns because it signals a decrease in the fundamental value of their holdings. Thus, our model explains both the average returns and impulse responses of large- and medium-stock reversals, while those of small stocks require the additional ingredient of inventory risk as in Nagel (2012).

Our final set of results revolves around the model's assumption that the quantities of private and market information fluctuate together over time. Although the quantity of information is not observable, we can proxy for it with the volatility it creates. Therefore, as a test of this assumption we check whether market volatility and average idiosyncratic volatility are positively correlated. We find that they clearly are: the two track each other extremely closely with a correlation of 92%.

---

[2]This is equivalent to including the VIX-squared portfolio and the market as test assets and requiring that they be priced with no error, see, e.g., Constantinides, Jackwerth and Savov (2013).

Second, we create a new set of test assets by using the co-movement of a stock's idiosyncratic volatility with market volatility as a proxy for the strength of the co-movement of its private information with market information. Thus, we sort stocks into quintiles based on the strength of the co-movement of their idiosyncratic volatility with market volatility, and then construct reversal portfolios within the quintiles. Consistent with the model, we find that the reversal portfolios in the high co-movement quintiles have much more negative market volatility betas and higher average returns than stocks in the low-co-movement quintiles. We then repeat all of our pricing tests on these "co-movement-sorted" portfolios. We find that their volatility betas do a very good job explaining their average returns, with an estimated price of risk that is again close to the one estimated from options prices. Thus, consistent with our model, the variance risk premium explains the liquidity premium in the cross section of co-movement portfolios.

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 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 make the amount of asymmetric information unknown to liquidity providers and correlated across assets.[3] This exposes liquidity providers to undiversifiable volatility risk.

The presence of undiversifiable risk allows our model to generate a liquidity premium without relying on financial frictions. This contrasts with inventory models, which as-

---

[3]Consistent with such correlation, Chordia, Roll and Subrahmanyam (2000) and Hasbrouck and Seppi (2001) find strong co-movement in illiquidity across assets. Similarly, Campbell et al. (2001) and Herskovic et al. (2016) find strong co-movement in volatilities across assets and with the market, as do we.

sume imperfect diversification (e.g., Stoll, 1978; Grossman and Miller, 1988; Duffie, 2010; Nagel, 2012). It also contrasts with the literature on intermediary asset pricing, which assumes segmentation between financial institutions and outside investors. This 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). As emphasized by Holmström and Tirole (1998), in these models liquidity is only scarce at the level of the aggregate claim. In our model there is no segmentation and liquidity is scarce at the level of the individual asset due to the co-movement in the amounts of asymmetric information across assets. This leads to the unique prediction that the premium for liquidity is explained by the variance risk premium with the same price of risk that prevails in other markets.

The empirical literature finds a liquidity premium in the form of higher average returns for illiquid assets (Amihud and Mendelson, 1986; Brennan and Subrahmanyam, 1996; Easley and O'Hara, 2004). Chordia, Sarkar and Subrahmanyam (2004) show that aggregate liquidity is decreasing in volatility. Amihud (2002) shows that an asset's liquidity impacts its expected return. Pástor and Stambaugh (2003) show that aggregate liquidity is a priced factor. Acharya and Pedersen (2005) provide an equilibrium model to capture this fact. Our contribution is to connect the liquidity premium to the variance risk premium both empirically and theoretically.

Starting with Lehmann (1990) and Lo and MacKinlay (1990), the literature has used short-term reversals to proxy for the returns to liquidity provision. Hameed, Kang and Viswanathan (2010) show that these returns are high following stock market downturns. Nagel (2012) shows that they are increasing in VIX. Collin-Dufresne and Daniel (2014) interpret them as due to slow-moving capital as in Duffie (2010). Our contribution is to explain reversal returns with the variance risk premium.

The literature on the variance risk premium focuses on option markets. Carr and Wu (2008); Todorov (2009) and Bollerslev and Todorov (2011) show that investors pay a large premium to hedge aggregate volatility risk. Bollerslev, Tauchen and Zhou (2009) and Drechsler and Yaron (2010) show that the variance risk premium predicts aggregate stock returns. Manela and Moreira (2017) extend this result back to 1890. Bao, Pan and Wang

(2011) and Longstaff et al. (2011) find parallel results for bonds. Drechsler and Yaron (2010), Drechsler (2013), and Dew-Becker et al. (2017) provide models that explain the behavior of the variance risk premium based on macroeconomic risk. Our contribution is to apply the insights of this literature to the pricing of liquidity.

2 Model

We present a model similar to 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 $i = 1, \ldots N$ 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 in all assets. Trading takes place on dates 0 and τ. Final payoffs are realized on date 1 and are given by

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

where $\overline{v}_i$ 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 orthogonal to the mechanism we are studying. The value of $\overline{v}_i$ is known by everyone ahead of time. Since v_i is idiosyncratic, $\overline{v}_i$ is the price of the asset before any trading takes place. Thus, $\sigma_{v,i}$ is the volatility of asset i's price over the whole time period.

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

$$y_i = \phi v_i. \tag{2}$$

The parameter ϕ controls how aggressively informed traders trade in the direction of their private signal v_i.[4] The uninformed liquidity traders demand z_i units of the asset, where $z_i \sim N\left(0, \sigma_{z,i}^2\right)$ is uncorrelated across assets.

We assume liquidity providers share risks frictionlessly with the rest of the economy, i.e. there is no market segmentation. As a result, they price date-1 payoffs using the

---

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

aggregate stochastic discount factor (SDF) Λ_1. The price of asset i is:

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

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.[5]

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:

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

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

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

where $k > 0$ is the loading of idiosyncratic variance on the common factor.[6] The common factor creates variation in the total amount of private information across assets, which makes their idiosyncratic volatilities co-move, as is the case empirically. We further assume that the total amount of private information covaries positively with the amount of information about the aggregate market σ_m^2:[7]

$$\sigma_v^2 = \sigma_m^2 + \varepsilon_v,\qquad(6)$$

where ε_v is orthogonal to σ_m^2 and all other shocks. As with σ_v^2, market participants do not

---

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

[6]In Section 4.8 we further allow stocks to have different loadings k_i on the common factor.

[7]We can write the aggregate market 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.

know σ_m^2 and must form expectations about it, $E_\tau[\sigma_m^2]$, at each point in time.

For simplicity, we assume liquidity providers observe the variance of order flow

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

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. The important thing is that liquidity providers cannot use the cross section to infer how much total private information σ_v^2 is in the market. To ensure this holds, we assume there is also a common factor in the quantity of liquidity demand:

$$\sigma_{z,i}^2 = \sigma_z^2 + \varsigma_{z,i}^2. \tag{8}$$

By observing the cross section of order flow, liquidity providers can infer the combination $\phi^2 k \sigma_v^2 + \sigma_z^2$ of the common factors but cannot separate out the amount of private information σ_v^2 from the amount of liquidity demand σ_z^2. They therefore remain uncertain about the amount of private information, and update their estimates when news arrives.

We assume such news arrives at the interim date τ. In particular, public news causes liquidity providers to update their expectations of future market volatility, $E_\tau[\sigma_m^2]$. By equation (6), liquidity providers also revise their estimate of the amount of private information, $E_\tau[\sigma_v^2]$. 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 order flow x_i.

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

$$p_{i,t} = \bar{v}_i + \phi \frac{E_t^Q\left[\sigma_{v,i}^2\right]}{\sigma_{x,i}^2} x_i. \tag{9}$$

All proofs are in Internet Appendix IA.2. As in Kyle (1985), prices are sensitive to order flow x_i because it contains information about fundamental values due to trading by informed 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

10

of prices based on their expectation of this quantity, $E_t^Q[\sigma_{v,i}^2]$. Note that this expectation is risk-adjusted, which is important because $\sigma_{v,i}^2$ can covary with the SDF Λ_1. In fact, it covaries positively with market variance, σ_m^2, which has a large negative price of risk (investors strongly dislike high market volatility). Thus, Proposition 1 generalizes Kyle (1985) by accounting for stochastic volatility and its large price of risk.

Liquidity providers absorb the order flow of other investors, so their portfolio holds $-x_i$ in each asset. If x_i was observable empirically, we could use it to directly analyze the risks liquidity providers face. Unfortunately, in general it is not.[8] Luckily, Proposition 1 shows that we can proxy for x_i using the change in an asset's price on date 0: $\Delta p_{i,0} = p_{i,0} - \bar{v}_i$. We can use it to characterize liquidity providers' portfolios as follows:

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

$$-x_i \;=\; -\frac{1}{\phi}\left(\frac{\sigma_{x,i}^2}{E_0^Q\left[\sigma_{v,i}^2\right]}\right)\Delta p_{i,0} \tag{10}$$

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 reversal strategies are mimicking portfolios for the portfolios of liquidity providers.[9] The next step is to characterize their returns. As Proposition 1 shows, an asset's price is exposed to volatility shocks. 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*

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

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 expected volatility increases the price of an asset that had positive date-0 order flow and decreases the price of an asset that had

---

[8]Volume is observable but has no sign. It also captures gross, not net, trading and there is no clear mapping between the two. Thus, volume is a poor proxy for net order flow.

[9]If there are also public news on date 0, then the proxy becomes imperfect. In Internet Appendix IA.1 we allow for such news and show that the model's main predictions carry through.

negative date-0 order flow. This is because an increase in expected volatility implies that there was more private information in the market than liquidity providers estimated at date 0. To reflect this change in the informativeness of past order flow, liquidity providers adjust prices further in the direction of their date-0 order flow.

Since order flow itself 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*

$$-\Delta p_{i,\tau} x_i = -\frac{\phi x_i^2}{\sigma_{x,i}^2} \left(E_\tau^Q \left[\sigma_{v,i}^2 \right] - E_0^Q \left[\sigma_{v,i}^2 \right] \right). \tag{12}$$

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 creation in terms of two trading strategies: reversals and variance swaps.[10] These strategies seem unrelated, as reversals are a bet against private information and variance swaps are a bet on volatility. Yet, they are tightly connected in our model, since the level of volatility is determined by the flow of information, which includes the private information liquidity providers are betting against. Thus, liquidity providers have a built-in negative exposure to volatility risk.

Since liquidity providers risk-adjust returns using the aggregate SDF (they are fully diversified), any premium they demand for holding reversal portfolios can only be due to an exposure to undiversifiable risk. As asset payoffs v_i are uncorrelated, one might think there is no such undiversifiable risk here. In this case liquidity providers would compete the cost of liquidity provision down to zero. However, this is not correct. Because liquidity providers are negatively exposed to assets' idiosyncratic volatilities, and idiosyncratic volatilities load on the common factor σ_v^2, liquidity providers are exposed to shocks to this undiversifiable risk factor. Moreover, since σ_v^2 loads strongly on market volatility σ_m^2 (Eq. 6), and market volatility shocks carry a high risk price in the SDF, so do shocks to

---

[10] A 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, $E_t^Q \left[\sigma_{v,i}^2 \right]$.

σ_v^2. Thus, liquidity providers have a large exposure to a highly priced risk factor—market volatility—for which demand a large risk premium.

The following proposition solves for asset i's market volatility beta on date τ:

Proposition 2. *The beta of liquidity providers' position in asset i to risk-adjusted expected market volatility $E_\tau^Q \left[\sigma_m^2 \right]$ is*

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

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

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

Proposition 2 shows that liquidity providers portfolios' are negatively exposed to market volatility risk. Note again that this is the case even though assets' final payoffs are completely uncorrelated with each other and with the market return. Nevertheless, liquidity providers portfolios have unambiguously negative market volatility betas. The reason is that even though assets' final payoffs are idiosyncratic, their volatilities have a common component. Since liquidity providers are effectively short variance swaps (Lemma 3), this common component exposes them to undiversifiable risk.

A large literature in volatility and option pricing documents that market volatility risk commands a very large negative price of risk, i.e. periods of large volatility spikes are priced as bad times by market participants. Formally, this means that risk-adjusted expected market volatility is higher than objective (i.e., statistical) expected volatility: $E_t^Q \left[\sigma_m^2\right] > E_t \left[\sigma_m^2\right]$. Equivalently, the price of a market variance swap, $E_t^Q \left[\sigma_m^2\right]$, is higher than its expected payoff, $E_t \left[\sigma_m^2\right]$. Since both the price and payoff converge to σ_m^2 as $t \to 1$, the price of the variance swap drifts down over time (under objective expectations):

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

This expected decrease in the price of the variance swap is the insurance premium the

buyer of the variance swap pays to hedge positive shocks to market volatility over the intervening period. In the literature this premium is called the variance risk premium. Since providing liquidity in our model is effectively shorting variance swaps, the premium liquidity providers charge—the liquidity premium—reflects the variance risk premium they should earn on the variance swaps.

Since the aggregate SDF prices market volatility shocks, the liquidity premium for an asset is given by its market volatility beta multiplied by the variance risk premium.[11] Summing up across all liquidity providers' positions, we get that 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 on liquidity providers' portfolios from date 0 to date τ is*

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

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 earn positive returns not because they are constrained or under-diversified, as is often 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

We now extend the model to incorporate market segmentation in the form of inventory costs, which are emphasized in the literature.

---

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

As in Nagel (2012), we model inventory costs by assuming that liquidity providers demand compensation for holding inventory of a given asset. This compensation is over and above any risk premium on the asset due to its covariance with the economy's stochastic discount factor. Thus, in contrast to our main model, in which the pricing of an asset is determined only by the economy's SDF, 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: the marginal cost of holding $-x_i$ units of asset i from date 0 to date t is $-t\gamma_{i,t}x_i$. Hence, the pricing condition (3) becomes

$$p_{i,t} = E_t^Q[p_{i,1}] + (1-t)\gamma_{i,t}x_i. \tag{17}$$

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, it compensates liquidity providers for their inventory costs and causes a reversal in the price of the asset equal to the inventory cost.

Inventory costs can vary both across assets and over time. Nagel (2012) focuses on variation over time driven by the VIX index. This is represented here by setting $\gamma_{i,t} = \gamma E_t^Q[\sigma_m^2]$. Collin-Dufresne and Daniel (2014), on the other hand, focus on specialized liquidity providers who are under-diversified and hence bear idiosyncratic risk. This is captured by setting $\gamma_{i,t} = \gamma E_t^Q[\sigma_{v,i}^2]$. Since the main model in this paper has no inventory costs, it is nested by $\gamma_{i,t} = 0$. The following proposition solves the extended model and shows how inventory costs affect the predictions of the main model.

Proposition 4. *The model with inventory costs implies the following:*

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

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

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

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

(19)

Part (i) of Proposition 4 shows that inventory costs induce liquidity providers to take smaller positions. Equivalently, they require a larger initial price change to hold a given position. To the extent that inventory costs are not observed and vary by asset, they introduce measurement error when we proxy for liquidity providers' portfolios using reversals. To address this issue in our empirical analysis, we split the sample by size, as inventory costs are likely to be higher for small stocks.

Part (ii) of Proposition 4 shows that inventory costs affect the returns on liquidity providers' positions in two ways. The first is that the value of their position in asset i drifts up over time at the rate $\gamma_{i,0} x_i^2$ to compensate them for the cost of carrying the inventory. As expected, this drift is over and above any compensation for risk.

The second effect is through shocks to the remaining inventory cost at date τ, $(1-\tau)(\gamma_{i,\tau} - \gamma_{i,0})$. An increase in the remaining inventory cost reduces the value of liquidity providers' position on impact so that it can appreciate going forward and in doing so recoup the higher inventory cost. By date 1 ($\tau \to 1$), all inventory costs have been paid out and the value of the position is back to its fundamental value. Hence, inventory cost shocks have a transitory impact that dissipates by date 1.[12]

Proposition 5. *The market volatility beta of liquidity providers' position in asset i from date 0 to date τ is given by*

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

(20)

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 costs to liquidity providers' volatility betas decreases with τ and goes to 0 as τ → 1.

Proposition 5 shows that inventory costs can make liquidity providers' market volatil-

---

[12]Of course, liquidity providers' payoff net of inventory costs is affected by surprise changes in realized inventory costs, but by date 1 liquidity providers no longer hold any inventory, hence this has no effect on the date-1 value of their positions.

ity betas more negative for small τ, but do not change the volatility beta over the full holding period as $\tau \to 1$. The short-run effect requires inventory costs to be increasing in market volatility, $\gamma > 0$, a condition that is satisfied in the two formulations we considered above. In both cases, a positive shock to market volatility on date τ raises inventory costs and lowers the value of liquidity providers' position, thereby giving it a more negative volatility beta. As τ increases towards 1, however, the remaining inventory cost runs out and the value of the position converges toward its fundamental value. Thus, inventory costs have no impact on volatility betas over the full holding period.

Because inventory costs are proportional to the holding period, even their short-run, transitory contribution to volatility betas is small. In contrast, the negative volatility betas due to uncertain private information are larger and do not decay with τ because they are driven by information about the fundamental value of the asset.

Proposition 5 thus allows us to separate the contributions of private information and inventory costs to the returns to liquidity provision. Private information exposes liquidity providers to volatility shocks, for which they demand compensation equal to their volatility beta multiplied by the variance risk premium. Inventory costs do not create exposure to volatility shocks over the full holding period. In addition, they can account for any excess returns to liquidity provision over and above the fair compensation for volatility risk. We formalize this with the following proposition:

Proposition 6. *The expected payoff of liquidity providers' portfolio from date 0 to τ is:*

$$E_0 \left[\sum_{i=1}^{N} -\Delta p_{i,t} x_i \right] = \left(\sum_{i=1}^{N} \beta_{i,\sigma_m}^{0 \to \tau} \right) \left(E_0 \left[E_\tau^Q \left[\sigma_m^2 \right] \right] - E_0^Q \left[\sigma_m^2 \right] \right) + \tau \gamma_{i,0}. \qquad (21)$$

The liquidity premium exceeds the variance risk premium associated with liquidity providers' portfolio by the inventory cost. As $\tau \to 1$, the portfolio's volatility beta is independent of inventory costs, so inventory costs only enter the expected payoff through the term $\gamma_{i,0}$.

Proposition 6 shows that as $\tau \to 1$, the component of liquidity providers' expected payoff that is due to inventory costs is independent of their volatility betas or the variance risk premium. This component is only the part of the payoff that exceeds its fair compensation for variance risk. The reason is that with inventory costs liquidity providers are segmented from other investors in the economy, so the average return they earn can ex-

ceed the fair premium based on the covariance of their payoffs with the economy's SDF. In contrast, in our main model, markets are fully integrated, so liquidity providers' average returns are given by their market volatility betas. Thus, we can infer the importance of inventory costs by looking at the average returns to liquidity provision in excess of their variance risk premium.

2.2 Summary of empirical predictions

We now summarize the key empirical predictions that emanate from the model. We test these predictions in Section 4.

Prediction 1. *Stocks' idiosyncratic volatilities share a common component that covaries positively with the market portfolio's volatility.*

This prediction is captured by Equation 6. We test it by analyzing the time-series relationship between the average of stocks' idiosyncratic volatilities and the VIX index, which gives the (risk-adjusted) expected market volatility. We test this prediction in Section 4.1, and in Section 4.8 we go further by showing that cross-sectional differences in this relationship line up with differences in the volatility betas and returns of reversal portfolios, as predicted by Eqs. (14) and (16).

Prediction 2. *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 in a stock is proportional to its day-0 price change normalized by its volatility and weighted by its order flow variance.*

Prediction 2 follows from Lemma 1. Since we do not directly observe liquidity providers' portfolios, we use this prediction and data on returns to construct reversal portfolio as an empirical proxy and use it to test the model.[13] Lemma 1 further requires us to weight the reversal portfolios by the variance of order flow ($\sigma_{x,i}^2$). Since we do not observe this variance either, we instead 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 3. *Reversals portfolios have negative market volatility betas; the more extreme portfolios have more negative betas.*

---

[13]The extended model in Internet Appendix IA.1 shows that reversals become a noisy but still unbiased proxy for liquidity providers' portfolios in the presence of public news about fundamentals.

Prediction 3 follows from Proposition 2. We estimate the reversal portfolios' market volatility betas and test if they are negative and larger for the more extreme reversal portfolios, which are associated with greater liquidity provision. We implement this by running regressions of the reversal strategy returns on contemporaneous innovations in the squared VIX index as Proposition 2 implies. We test this prediction in Section 4.3.

Prediction 4. *The expected return of reversals due to their volatility risk equals their market volatility beta times the variance risk premium. Any expected return in excess of this is explained by inventory costs or other market frictions.*

Prediction 4 is given by Proposition 3 for the baseline model and Proposition 6 for the inventory cost generalization. This prediction allows us to separate the part of reversals' expected return that is compensation for volatility risk from whatever remains due to inventory costs or other market frictions. Since volatility risk is traded in the options market, we test if liquidity provision is integrated with the rest of the market by comparing the price of volatility risk we estimate from the cross-section of reversal returns with the one that prevails in the options market. We test this prediction in Sections 4.4 and 4.5.

Prediction 5. *The expected returns of reversals co-move with the variance risk premium, i.e., the variance risk premium predicts reversals' returns in the time series. In the cross section this predictive coefficient is proportional to reversals' volatility betas.*

Prediction 5 is also given by Propositions 3 and Proposition 6. The variance risk premium for the period 0 to τ is given by $E_0\left[E_\tau^Q\left[\sigma_m^2\right]\right] - E_0^Q\left[\sigma_m^2\right]$. Thus, by Eq. 16, a reversal's predictive coefficient is its market volatility beta.[14]

Prediction 6. *Reversal portfolios' betas remain negative and constant over liquidity providers' entire holding period. Equivalently, the impact of market volatility shocks on reversal portfolios' returns does not dissipate over liquidity providers' holding period. In contrast, under a pure inventory cost model market volatility shocks have a transitory impact on reversal returns, so reversals' market volatility betas decay to zero by the end of liquidity providers' holding period.*

Prediction 6 follows from Proposition 5. In our model, a positive market volatility shock is associated with an increase in the quantity of private information in the market, which negatively affects the expected fundamental value of the reversal portfolios.

---

[14]Nagel (2012) finds evidence for this prediction by showing that the return of a reversal portfolio is predicted by VIX. Nagel (2012) interprets this evidence through the lens of an inventory cost model. Our model provides a novel explanation for this finding and relates it to other cross-sectional results.

Hence, volatility shocks have a permanent impact on reversal returns. In contrast, in inventory-cost models volatility shocks only affect the inventory cost of assets while held by liquidity providers, and thus have no impact on their fundamental value. Thus, Prediction 6 allows us to separate the impact of volatility shocks on private information—the mechanism of our model—from the inventory-cost mechanism that has been the focus of the literature. We test this prediction in Section 4.7.

3   Data and summary statistics

In this section we describe our data and how we construct the reversal portfolios.

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.[15] We also exclude stocks that are within one day of an earnings announcement as in Collin-Dufresne and Daniel (2014). Earnings announcements are public-news events and as shown in Internet Appendix IA.1 this introduces measurement error when we use reversals as a proxy for the returns to liquidity provision.[16]

Sample selection: The sample is daily from April 9, 2001 to May 31, 2020, which gives a total of 4,815 trading days. The starting 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 2 of our model. Each day, we first sort stocks into quintiles by market capitalization. We do so because small stocks exhibit higher inventory costs and other forms of market segmentation, allowing us to differentiate our model as discussed in Section 2.1.

We form ten decile portfolios within each size quintile by sorting stocks by their nor-

---

[15]We 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.

[16]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).

malized return. 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.1. In addition, market-beta-adjusting ensures that the returns on the portfolios are not compensation for market risk.

Normalizing the returns by dividing by their standard deviation is also motivated by the model (see Prediction 2). Formally, the amount of liquidity provision in a stock, as proxied by the position of the liquidity provider, is proportional to the stock's price change scaled by its variance (see Lemma 1). Normalizing thus ensures that we are not just picking up more volatile stocks.

Within each decile portfolio we weight stocks by their dollar volume. As discussed in Prediction 2, 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.

We follow Nagel (2012) and hold each portfolio for five trading days in most of our tests. This horizon is also consistent with the evidence in Hendershott and Seasholes (2007) that NYSE specialists (liquidity providers) earn most of their returns within five days of entering in a position. As we will see, the same is true of our reversal portfolios. Note that although liquidity provision in recent years has been dominated by high-frequency trading, lower-frequency liquidity provision remains essential since imbalances between ultimate buyers and sellers often persist for several days. The presence of a reversal premium supports this view.

Reversal strategies: We form long-short reversal strategies that buy the low-return portfolio and sell the high-return portfolios 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. The 2–9, 3–8, 4–7, and 5–6 reversal strategies are constructed analogously for the inner deciles.

Liquidity provider portfolio: In some tests, we report results for an aggregate portfolio that mimics the portfolios of liquidity providers in our model. The portfolio is constructed

21

by weighing the long-short reversal strategies by the product of their normalized return and dollar volume, again as implied by Lemma 1. Thus, the liquidity provider portfolio implements the same weighting across portfolios as within them.

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, $E_t^Q \left[\sigma_m^2 \right]$. In particular, following Propositions 2 and 3, and as discussed in Predictions 1–6, we use changes in the squared VIX (divided by 100 for legibility) as our variance risk factor.

The VIX return: We use data on S&P 500 index options from OptionMetrics to calculate the return to holding the VIX basket (this data ends on December 31, 2019). We use this return to restrict the price of volatility risk as implied by Proposition 3. It is necessary to calculate it 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 thus not an investable trading strategy. 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.[17] 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.

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 given

---

[17]See 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.

size quintile (Small, 2, 3, 4, and Big) and each column represents a given long-short port-folio formed across the normalized return deciles (Lo–Hi, 2–9, 3–8, 4–7, and 5–6). Each reversal strategy, corresponding to an entry in the tables, 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 quintile. 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 was used to normalize the returns of individual stocks before sorting them into portfolios. Here we report its aver-age 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 calcu-lated 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 ex-pected, illiquidity is strongly decreasing in size: the largest stocks have illiquidity that is four orders of magnitude smaller than for the smallest stocks. Liquidity is thus relatively more abundant among the largest 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). By construction, all the sorting-day returns are negative. Since small stocks are more volatile than large stocks, their sorting-day returns are substantially larger in magnitude. For instance, the average sorting-day return of the Lo–Hi strategy is −13.55% for the smallest stocks versus −5.20% for the largest ones. By construction, the sorting-day returns decline in magni-tude as we move toward the inner deciles until they are close to zero for the 5–6 strategies. This reflects the fact that the Lo–Hi strategy carries the strongest reversal signal and there-fore captures the most intensive liquidity provision.

Panels E and F of Table 1 look at turnover. Panel E shows average turnover over the

23

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).[18] 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 illustrates the fact that the VIX basket is re-balanced each day to keep its maturity constant, hence the change in VIX is not a valid return and does not convey the variance premium. This is why we compute the VIX return, which is the percentage change in the price of a given 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 more right-skewed, as shown by the difference between its mean and median (-1.53% versus -5.08%) and 1^{st} and 99^{th} percentiles (-25.33% versus 71.29%). The reason for this 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

---

[18]The variance premium is typically measured as the realized variance of the S&P 500 over 30 calendar days divided by the squared VIX index. Since it can only be computed at 30-day horizons, comparing it to the VIX return requires dividing by 21 (the number of trading days in a typical 30-day period).

calculate the VIX return premium per unit of beta, i.e. the price of risk. We will use this price of risk in Section 4.5 to test whether the price of variance risk in option markets can account for the returns to liquidity provision as implied by our model.

Given the greater volatility of the VIX return, its per-unit price of risk is smaller than its total premium. 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 Idiosyncratic volatility and market volatility

We begin by verifying that idiosyncratic volatility and market volatility are closely related, consistent with Prediction 1. This is shown graphically in Figure 2, which plots an aggregate measure of idiosyncratic volatility against the VIX index. We calculate this measure from the beta-adjusted returns of individual stocks, which takes out the market component to give us a clean measure of idiosyncratic volatility. Each day, we square these returns, value-weight them across stocks, and take the square root. To make idiosyncratic volatility comparable to VIX, which is forward-looking, we average it over the following 21 trading days (about 30 calendar days), and annualize it.

The figure shows that idiosyncratic volatility and VIX are highly related and share many of the same fluctuations. For instance, like VIX, idiosyncratic volatility rose sharply during the 2008 financial crisis and even more so during the 2020 Covid crisis. Thus,

consistent with Prediction 1, the idiosyncratic volatilities of stocks have a large common component that is strongly correlated with market volatility.[19] In our model, this common component represents a systematic risk to liquidity providers.

Our model (Propositions 2 and 3) requires us to use a forward-looking measure of expected risk-adjusted variance. Since idiosyncratic variance is an ex-post realized outcome, the appropriate forward-looking measure for testing the model is the squared VIX.

4.2 Reversal strategy 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 (based on Newey-West standard errors with five lags to correct for overlap in the returns). In terms of magnitude, the 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 resulting annual Sharpe ratio is 0.5, slightly higher than the market Sharpe ratio in our sample (0.41).

As our model predicts (Predictions 3 and 4), 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. This is true among all size quintiles. Intuitively, the inner deciles reflect less intensive liquidity provision and therefore earn lower premiums.

The reversal strategies' 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.[20] Thus, providing liquidity for very small stocks earns higher returns.

Panels E and F of Table 3 report the CAPM alphas and associated t statistics of the reversal strategies. They are obtained from the time-series regressions:

$$R_{t,t+5}^{p} = \alpha^{p} + \beta^{p} R_{t,t+5}^{M} + \epsilon_{t,t+5}^{p}, \tag{22}$$

[19]Table IA.1 in the Internet Appendix presents a formal test. It shows that VIX predicts idiosyncratic volatility at least as well or better than market volatility.

[20]The difference is partly due to bid-ask bounce (Roll, 1984), which can be seen as part of the return to liquidity provision. In any case, it does not impact large stocks, whose bid-ask spreads are very small.

where $R^p_{t,t+5}$ is the cumulative excess return of portfolio p from t to $t+5$ and $R^M_{t,t+5}$ is the cumulative excess market return over the same period. As in Panel B, the t statistics are based on Newey-West standard errors with five lags.

The panels show that the CAPM cannot account for the returns of the reversal strategies. In all cases, the CAPM alphas are about the same as the raw average returns and remain significant whenever the raw returns are significant. For instance, 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.

4.3 Reversal strategy volatility risk

We now test Prediction 3 of our model directly. Specifically, we test whether the reversal strategies are exposed to volatility risk by running the following regressions:

$$R^p_{t,t+5} \;=\; \alpha^p + \beta^{p,VIX}\Delta VIX_{t,t+5} + \beta^{p,M}R^M_{t,t+5} + \epsilon^p_{t,t+5}. \tag{23}$$

These regressions have the same form as (22), but with the change in VIX-squared included alongside the market return. This is the right factor to use according to our model. It is the empirical counterpart to the change in the risk-adjusted market variance in Proposition 2. As before, we compute t-statistics based on Newey-West standard errors with five lags to account for the overlap in the returns.

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 3. They are also very similar across size quintiles, both in terms of magnitude and statistical significance. Also consistent with Prediction 3, 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 an amount equal to its average premium whenever VIX-squared rises by

27

one point. From Table 2, the standard deviation of VIX-squared changes is 1.53 points per day, which works out to about 3.42 points per five days. Hence, a one-standard deviation increase in VIX-squared over the holding period wipes out nearly three and a half times the strategy's average return. This shows that 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 reflect market risk rather than volatility risk. The table shows that this is not the case. The betas with respect to VIX-squared changes 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 as intended by their construction.

Overall, the results of Table 4 show that the reversal strategies have large and robust negative volatility risk betas, confirming Prediction 3 of the model.

4.4 Fama-MacBeth regressions

We now test Prediction 4 of the model, which says that if there is no market segmentation in liquidity provision the expected returns of the reversal strategies should equal the product of their volatility risk betas and the price of variance risk. In this section we treat the price of variance risk as a free parameter that we recover from the cross section of reversal strategies. In the next section we sharpen the analysis by testing whether the price of variance risk in the cross section of reversal strategies and option markets are the same such that these markets are integrated.

We run two-stage Fama-MacBeth regressions to test whether volatility risk exposure can explain the average returns of the reversal strategies. The first stage estimates the betas as in Section 4.3 (see Equation (23)). The betas then enter the second-stage regression:

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

This regression provides time series of estimated factor premia $\lambda_{t,VIX}$ and $\lambda_{t,M}$. Following the Fama-MacBeth methodology, we use the time series variation in the factor premia to calculate their standard errors, again with a Newey-West correction to account for the overlap in the returns. To assess the performance of our asset pricing model, 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 of the portfolios are jointly equal to zero.

The results of the Fama-Macbeth regressions are presented in Tables 5 and 6. Table 5 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 0.205%. The overall r.m.s.e. is similar, 0.24%. The r.m.s.e. of the long-short reversal strategies is also similar, 0.228%. Excluding the smallest two quintiles or value-weighting reduces it to 0.123% bps and 0.124%, 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, equal to 6.2% annualized. The last row gives the pricing error of the 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, 0.168% over the five-day horizon or 8.4% annualized.

Column 2 of Table 5 reports the specification with the market return as the sole factor (i.e., the CAPM). The market premium is positive, 0.325%, 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 only small differences in market betas. It achieves a cross-sectional R^2 of 15.8% and r.m.s.e. of 0.112%, 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 0.201%, only slightly lower than the raw value. When we exclude small stocks or value-weight, the pricing error is 0.094%, about one quarter smaller than the raw value. The pricing error of the liquidity provider portfolio is 0.127%, also one quarter smaller. As in column 1, all of the p-values are zero to within two decimal points, hence the CAPM is also 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 price of risk to do so. Even then, it is unable to explain most of the differences in returns between the portfolios, which are captured by the 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.

Column 3 introduces 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 0.107%, and that of the reversal strategies to 0.183%. The impact is larger when we exclude small stocks or value-weight: the pricing errors drop to 0.065% and 0.062%, respectively, which is about half the corresponding raw values. The pricing error of the liquidity provider portfolio also drops in half, to 0.084%. The p-values remain close to zero, hence this model is also statistically rejected.

Thus, the univariate volatility risk model is 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 therefore settles on a medium-sized price of risk, which limits its ability to fully capture the cross section. Note however that the 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 0.152% (over a five-day period). Although insignificant, this number is very close to the equity premium in the sample. This helps the model price the level of returns among the portfolios. The price of variance risk doubles to -1.079% and is highly significant. As we will see in the next section, this number is very close to the price of variance risk implied by option markets. For now, we see that the model does a significantly better job in explaining the cross section. The cross-sectional R^2 increases to 37.9%, the overall r.m.s.e. drops to 0.097%, and the r.m.s.e. of the reversal strategies drops to 0.158%. These pricing errors are still jointly significant with p-values close to zero. The reason is the small stock port-

folios, whose returns are unusually high. When we remove them, the pricing error drops to just 0.04%, 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 even further, to just 0.006%, a 95% decline 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 0.008% (0.4% annualized), and whose *p*-value is 0.78. The bivariate model thus prices the liquidity provider portfolio almost perfectly and therefore explains the liquidity premium in the cross section of stocks.

Table 6 takes a closer look at the pricing errors of the individual long-short reversal strategies. Panel A shows the CAPM pricing errors, which are based on the estimates in column 2 of Table 5. 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 0.14% (*t*-statistic of 3.43), only slightly smaller than the 0.17% 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 the cross section of stocks.

Panel B of Table 6 shows the pricing errors of the univariate volatility risk model (column 3 of Table 5). The pricing error of the large-cap Lo–Hi reversal strategy drops to 0.08%. Only the smallest quintiles retain significant pricing errors, including the 0.68% pricing error of the smallest quintile (*t*-statistic of 7.45).

Panel C looks at the two-factor model, which combines the market risk and volatility risk factors (column 4 of Table 5). Now the pricing error of the large-cap Lo–Hi strategy is fully explained: it drops to -0.02% 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. The pricing error of the smallest quintile is 0.58% (versus a raw return of 0.78%). This confirms the finding in Table 5 that the overall r.m.s.e. of the two-factor model remains significant because of the high reversal returns among the smallest stocks. Overall, Table 6 shows that volatility risk, in combination with the market factor, explains the returns to liquidity provision among all but the smallest stocks.

Figure 3 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 represents a different size quintile. Within it are five data points corresponding

31

to the five long-short reversal strategies across deciles: Lo–Hi, 2–9, and so on. Also shown is the 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 univariate 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 some reversal strategies lie above the 45-degree line (their average returns are higher than their predicted returns).

Panel C of Figure 3 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 reversal strategies lie symmetrically around the 45-degree line. The liquidity provider portfolio falls almost exactly on the 45-degree line (recall its pricing error is 0.006%). Only the outer-decile small-stock strategies lie significantly above it. While their predicted returns are fairly high, their realized average returns are even higher. Thus, volatility risk cannot fully explain the returns to liquidity provision among small stocks. By contrast, Figure 3 shows that volatility risk can explain the returns to liquidity provision among larger stocks, which are economically more important.

4.5 Option-implied price of volatility risk

The next question we address is whether the price of volatility risk needed to price the reversal strategies is consistent with the one that prevails in option markets, where volatility risk is directly traded. Answering this question sharpens our test of Prediction 4. It also sheds light on the broader question of whether the returns to liquidity provision reflect narrow intermediation frictions or widely shared economic risks.

The natural place to measure the price of volatility risk is in option markets. As discussed in Section 3, the squared VIX is the price of a basket of options whose payoff repli-

cates the realized variance of the S&P 500. Yet we cannot use the change in VIX squared to measure the price of variance risk because the VIX basket is rebalanced each day. To solve this problem, we track the price of a given VIX basket from one day to the next. The percentage change in the price of this basket is the VIX return, which we denote R^{VIX}.

Table 2, which we discussed in Section 3.1, gives us the option-implied price of risk for our volatility risk factor, ΔVIX^2. Based on column 2, this price of volatility risk is -1.03% per unit of beta over a five-day horizon (-51.5% annualized). We similarly obtain the price of market risk from the average market return over a five-day horizon in our sample. The resulting restricted price of market risk is 0.16%. These prices of risk are reported in the top half of Table 7. Also shown are Newey-West standard errors based on the time series variation of the VIX return and market return.[21] The option-implied price of volatility risk is highly statistically significant.

The next step is to multiply the restricted prices of risk by the portfolios' betas in order to obtain their predicted returns. Then to get the pricing errors we subtract the portfolios' predicted returns from their average returns. As in all prior tests, we use the time series variation in these pricing errors to estimate their standard errors. The bottom half of Table 7 reports the pricing errors for the same sets of portfolios as in Table 6.

The first column of Table 7 shows the raw returns as before. In the second column, the CAPM lowers the r.m.s.e. of the full set of portfolios from 0.24% to 0.116%. However, it has almost no impact on the long-short reversal strategies whose r.m.s.e. is 0.212% under the CAPM versus 0.228% in the raw returns. 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 liquidity provider portfolio, the pricing error ticks down from 0.168% to 0.146%, 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 univariate volatility risk model. This model does worse on the level of returns, driving the overall r.m.s.e. up to 0.76%, but much better on the difference in returns, driving the long-short reversal strategy r.m.s.e. down to 0.158%, a one-third reduction. The model's explanatory power increases substantially when we exclude small stocks: the r.m.s.e. drops from 0.123% to 0.038%, a

---

[21]Note that there is no cross-sectional R^2 since we are not running a cross-sectional regression.

two-thirds reduction. Value-weighting reduces the pricing error even more: from 0.124% to 0.012%, a 90% reduction. Finally, the pricing error of the liquidity provider portfolio drops from 0.168% to 0.015%, also a 90% reduction. The associated p-values are above conventional cutoffs. Thus, the model with an option-implied price of volatility risk explains the returns to liquidity provision among larger stocks.

Finally, 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 differences in returns. Thus, 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 0.123% to 0.04%). As in column 3, value-weighting leads to a 90% drop in pricing errors (from 0.124% to 0.011%). The pricing error of the liquidity provider portfolio also drops by 90% (from 0.168% to 0.015%) and becomes insignificant.

Table 8 reports the pricing errors of the individual reversal strategies. Panel A is very similar to Panel A of Table 6 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 0.17% to −0.01% 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 0.59% 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 0.23%, down 40% from the raw return. Both remain highly significant. The reason for this is that while both strategies have volatility betas similar to those of the larger quintiles, their returns are significantly larger. In other words, liquidity provision among small stocks earns abnormal returns. This is consistent with market segmentation in the form of inventory costs, as discussed in Section 2.1.

Finally, Figure 4 depicts average versus predicted returns with the restricted prices of risk. Similar to Figure 3, 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 restricted prices of risk show that volatility risk carries the same price of risk in the market for liquidity (reversals) and the market for variance (options). This price of risk explains the returns to liquidity provision among all but the smallest stocks and overall. This is consistent with Prediction 4, which is a key implication of our model. It shows that the liquidity premium in the cross section of stocks is explained by the variance risk premium.

4.6   Predicting reversals 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 times of high volatility (Todorov, 2010; Andersen, Fusari and Todorov, 2015). Investors thus demand greater compensation for bearing greater volatility risk. 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 5, which we test next. This time series test provides evidence complementary to the cross-sectional tests of the previous section.

Table 9 shows results from predictive regressions of reversal strategy returns on the level of VIX squared as of the portfolio formation date. Panel A reports the predictive loadings (times 100). Focusing first on the Lo–Hi strategies, the loadings are similar across 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, implying that a one-point increase in VIX squared is associated with a 9-bps higher return over the following five days. This number is large relative to the 17-bps average return of the strategy. The predictive loadings decline steadily as we move toward the inner decile strategies. Hence, they display exactly the same pattern as the volatility risk betas in Table 4. This is consistent with Proposition 3 of the model. Intuitively, portfolios that are more exposed to volatility risk also have a risk premium that co-moves more strongly with aggregate volatility risk.

Panel C of Table 9 shows the R^2 coefficients from the predictive regressions. 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.[22] Thus, just as in the pricing results in Section 4.4, aggregate volatility risk has the highest explanatory power for large-cap reversal strategies.

Overall, the predictability results in Table 9 confirm the main finding of Nagel (2012) that VIX predicts reversals, and extend this result to large-cap stocks.

4.7   Reversal strategy dynamics

To further test our model, and in particular Prediction 6, we zoom in on the dynamics of the returns and volatility risk of the reversal strategies. These dynamics are shown in Figures 5 and 6. In each figure, we fix the composition of the portfolios as of the formation date (date 0) and follow their returns over time (i.e., without rebalancing).

Panel A of Figure 5 shows the average returns of the large-cap reversal strategies for horizons up to ten trading days. Returns rise steadily with horizon and begin to level off past the five-day mark. They are highest for the Lo–Hi strategy and decline toward the inner deciles.[23] Panel B of Figure 5 shows that the predictive loadings on VIX squared follow the same steady pattern as the returns. Panels C and D show the same for the volatility risk betas, whether we control for the market return (Panel D) or not (Panel C). The steady pattern across horizon indicates similar exposure to volatility risk throughout the holding period. The fact that predictive loadings and volatility risk betas line up with the average returns builds further support for the pricing results of Section 4.4.

Figure 6 tests Prediction 6 of our model by looking at the persistence of the impact of VIX shocks on the reversal strategies. In our model, VIX shocks have a persistent impact because they reveal information about the fundamental value of the assets in liquidity providers' portfolios. By contrast, a transitory impact would be consistent with a model where VIX shocks increase inventory costs (see Section 2.1). Increased inventory costs reduce the value of liquidity providers' portfolios only temporarily in order to compen-

---

[22]Following 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 reversal strategy would see a six-fold increase in expected returns relative to a buy-and-hold strategy.

[23]The 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.

sate them at a higher rate going forward. The differential predictions of the two models thus provide us with a sharp test. We implement it by plotting the impulse responses of reversal strategy returns to a VIX-squared shock one day after portfolio formation:

$$R^p_{t,t+h} \;=\; \alpha_{p,h} + \beta^p_{VIX,h} \Delta VIX^2_{t,t+1} + \epsilon^p_{t,t+h'} \quad \text{for} \quad h = 1, 2, \ldots, 20. \tag{25}$$

If the impact of VIX shocks is persistent, the betas $\beta^p_{VIX,h}$ should be flat across horizon h. If it is transitory, they should converge towards zero.

Panel A of Figure 6 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 five-day return. The impact remains flat over time and settles at 19 bps after twenty trading days. The 95% confidence bands lie comfortable away from zero at horizons well beyond the holding period associated with liquidity provision. Thus, this is inconsistent with the transient effect predicted by inventory costs. Instead, it points to a permanent impact on value, as predicted by our model. The results therefore support Prediction 6 of our model, and go against the predictions of inventory cost models.

Panel B Figure 6 looks at all size quintiles. The pattern for quintiles three and four is identical to that for the largest quintile. This is not the case for the smallest two quintiles, however. The returns of the smallest quintiles tend to revert back towards zero after the initial drop, consistent with the impact of an increase in inventory costs.

Table 10 formalizes the results of Figure 6. Panel A shows that after five days the Lo–Hi strategies of each size quintile have similar negative betas with respect to day-one VIX shocks. By day ten, however, the smallest two quintiles have reverted to about zero, while the three larger quintiles retain their betas unchanged through day twenty.

Figure 6 and Table 10 thus reinforce the results of our asset pricing tests in Section 4.4. Large-cap reversals have large and persistent volatility risk exposures that explain their returns with the same price of risk as in option markets. By contrast, small stocks exhibit evidence of market segmentation: their returns are abnormally high, and their volatility risk exposures are transitory. The differences between small and large stocks are natural because small stocks are thinly traded and rely on specialized intermediaries for liquidity provision. The inventory constraints of these liquidity providers therefore influence their prices. Our results for small stocks are thus consistent with Nagel (2012), who analyzes

37

reversals without value-weighting and interprets their predictability by VIX through the lens of an inventory cost model. At the same time, when we extend the analysis to large stocks we find that the liquidity premium is explained by volatility risk.

4.8   Volatility co-movement portfolio sorts

In Section 4.1, we showed that idiosyncratic volatility and aggregate volatility co-move strongly, consistent with Prediction 1. In this section we extend this test to the cross section 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.[24]

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:

$$\widehat{\sigma}^i_{t,t+5} \;=\; a + k_i VIX_t + \epsilon_{i,t}. \tag{26}$$

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.2 in the Internet Appendix). The next step is to sort 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.

The results for the volatility co-movement portfolios are summarized in Table 11.[25] Panel A shows a strong pattern in CAPM alphas across quintiles. As predicted by our model, portfolios of stocks whose idiosyncratic volatility co-moves more strongly with market volatility earn substantially higher reversal returns. The alphas of the Lo–Hi

---

[24]Formally, we can generalize Eq. 5 by allowing for heterogeneous loadings: $\sigma^2_{v,i} = k_i \sigma^2_v + \varsigma^2_{v,i}$, $k_i > 0$. The volatility betas become $\beta_{i,\sigma_m} = -\phi k_i \left(x^2_i / \sigma^2_x \right)$, 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.

[25]The Internet Appendix contains a full set of tables analogous to our main results.

38

strategies increase monotonically from 10 bps to 26 bps from the first quintile to the fifth.

Panel B of Table 11 looks at the volatility risk betas (controlling for market beta). 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 -0.07 and insignificant, while the beta of the fifth quintile is -0.34 and highly significant.

Panel C of Table 11 shows the pricing errors from Fama-Macbeth regressions with an option-implied price of risk. 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 -5 bps in the two-factor model with volatility risk. Figure 7 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.

5   Conclusion

Our results provide a new perspective on 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 asymmetric information they face. A spike in volatility reveals greater asymmetric 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 substantial volatility risk. Moreover, this exposure explains the returns to reversals over time and across all but the smallest stocks.

Overall, we find that liquidity is priced based on the broad economic risk of aggregate volatility fluctuations, 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 importance of the financial sector in the economy. This possibility is intriguing, as it promises to further integrate the asset pricing and financial intermediation literatures.

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

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

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

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

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

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

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

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

Collin-Dufresne, Pierre, and Kent Daniel. 2014. "Liquidity and return reversals." *Working paper*.

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, 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).

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

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

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

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, Ľuboš, 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.* Princetion 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.

<table>
<tr><td colspan="6" align="center">Panel A: Market cap</td></tr>
<tr><td></td><td>Lo–Hi</td><td>2–9</td><td>3–8</td><td>4–7</td><td>5–6</td></tr>
<tr><td>Small</td><td>0.12</td><td>0.12</td><td>0.12</td><td>0.12</td><td>0.12</td></tr>
<tr><td>2</td><td>0.36</td><td>0.36</td><td>0.36</td><td>0.36</td><td>0.36</td></tr>
<tr><td>3</td><td>0.89</td><td>0.90</td><td>0.89</td><td>0.89</td><td>0.89</td></tr>
<tr><td>4</td><td>2.49</td><td>2.49</td><td>2.48</td><td>2.48</td><td>2.48</td></tr>
<tr><td>Big</td><td>81.75</td><td>81.21</td><td>81.08</td><td>78.77</td><td>79.26</td></tr>
</table>

	Lo–Hi	2–9	3–8	4–7	5–6
		Panel B: Idiosyncratic volatility			
Small	3.74	4.22	4.77	5.50	5.96
2	3.41	3.72	3.94	4.09	4.17
3	2.98	3.20	3.34	3.44	3.48
4	2.47	2.67	2.79	2.87	2.90
Big	1.71	1.85	1.94	1.99	2.01

Panel C: Amihud illiquidity

	Lo–Hi	2–9	3–8	4–7	5–6
Small	171.66	105.00	55.79	33.03	17.10
2	7.31	4.42	2.93	1.94	1.60
3	0.94	0.65	0.48	0.37	0.32
4	0.19	0.14	0.11	0.09	0.08
Big	0.02	0.01	0.01	0.08	0.01

Panel D: Sorting-day return

	Lo–Hi	2–9	3–8	4–7	5–6
Small	−13.55	−7.12	−4.57	−2.59	−0.84
2	−10.80	−5.57	−3.52	−1.98	−0.64
3	−8.98	−4.63	−2.96	−1.67	−0.54
4	−7.39	−3.99	−2.57	−1.46	−0.47
Big	−5.20	−2.98	−1.95	−1.12	−0.36

Panel E: Average turnover

	Lo–Hi	2–9	3–8	4–7	5–6
Small	1.14	1.33	1.47	1.54	1.53
2	1.47	1.54	1.63	1.65	1.68
3	1.74	1.80	1.84	1.87	1.88
4	1.78	1.86	1.91	1.94	1.96
Big	1.16	1.23	1.28	1.30	1.30

Panel F: Sorting-day turnover

	Lo–Hi	2–9	3–8	4–7	5–6
Small	2.59	1.90	1.83	1.82	1.71
2	2.42	1.58	1.44	1.40	1.37
3	2.61	1.74	1.58	1.52	1.49
4	2.55	1.82	1.67	1.61	1.59
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***	-1.515***	-2.013***	-2.523***
	(0.172)	(0.220)	(0.135)	(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

	Lo–Hi	2–9	3–8	4–7	5–6
Small	0.78	0.46	0.20	−0.08	−0.10
2	0.40	0.15	0.06	0.08	−0.02
3	0.15	0.12	0.04	0.04	−0.02
4	0.16	0.19	0.11	0.12	0.02
Big	0.20	0.16	0.18	0.08	0.00

Panel B: *t*-statistics

	Lo–Hi	2–9	3–8	4–7	5–6
Small	8.39	5.84	2.21	−0.98	−1.30
2	6.37	3.04	1.34	1.74	−0.37
3	2.79	2.85	1.05	0.98	−0.63
4	3.35	4.73	3.38	3.92	0.65
Big	4.47	4.66	6.17	3.54	0.17

Panel C: Standard deviations

	Lo–Hi	2–9	3–8	4–7	5–6
Small	6.18	5.49	5.60	5.43	5.41
2	3.91	3.35	3.24	3.08	3.10
3	3.31	2.71	2.67	2.65	2.20
4	3.01	2.52	2.16	1.95	1.88
Big	2.84	2.29	1.96	1.70	1.52

Panel D: Sharpe ratios

	Lo–Hi	2–9	3–8	4–7	5–6
Small	0.90	0.60	0.25	−0.10	−0.14
2	0.72	0.32	0.14	0.18	−0.04
3	0.32	0.30	0.11	0.10	−0.07
4	0.37	0.53	0.35	0.43	0.07
Big	0.50	0.48	0.64	0.35	0.02

Panel E: CAPM alphas

	Lo–Hi	2–9	3–8	4–7	5–6
Small	0.74	0.44	0.16	−0.10	−0.10
2	0.38	0.12	0.05	0.07	−0.02
3	0.12	0.10	0.03	0.03	−0.03
4	0.13	0.17	0.10	0.11	0.02
Big	0.17	0.14	0.16	0.07	−0.00

Panel F: CAPM alpha *t*-statistics

	Lo–Hi	2–9	3–8	4–7	5–6
Small	7.95	5.62	1.76	−1.23	−1.28
2	6.16	2.46	1.09	1.55	−0.50
3	2.37	2.42	0.66	0.71	−0.81
4	2.81	4.38	3.12	3.69	0.87
Big	3.90	4.23	5.60	3.14	−0.15

Table 4: Volatility risk 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

$$R^p_{t,t+5} = \alpha_p + \beta^p_{VIX}\Delta VIX^2_{t,t+5} + \beta^p_M R^M_{t,t+5} + \epsilon^p_{t,t+5},$$

where $R^p_{t,t+5}$ is the cumulative excess return on reversal strategy portfolio p from the portfolio formation date t to $t+5$, $R^M_{t,t+5}$ is the excess return on the market portfolio, and $\Delta VIX^2_{t,t+5}$ 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^p_M. 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^p_{VIX}	Lo–Hi	2–9	3–8	4–7	5–6	t-statistics	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^p_{VIX}	Lo–Hi	2–9	3–8	4–7	5–6	t-statistics	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^p_M	Lo–Hi	2–9	3–8	4–7	5–6	t-statistics	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: Fama-Macbeth regressions

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

Table 6: Fama-Macbeth pricing errors

The table shows the pricing errors of the reversal strategies from Fama-Macbeth regressions (see Table 5 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 7: 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**		0.160**
		(0.081)		(0.081)
β_{VIX}			−1.032***	−1.032***
			(0.173)	(0.173)
Constant	0.205**			
	(0.101)			
N	50	50	50	50
(i) All portfolios:				
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
(ii) Long-short reversal strategies:				
R.m.s.e.	0.228	0.212	0.158	0.159
p-value	0.00	0.00	0.00	0.00
(iii) Long-short reversal strategies (ex small stocks):				
R.m.s.e.	0.123	0.107	0.038	0.040
p-value	0.00	0.00	0.05	0.02
(iv) Value-weighted reversal strategy:				
Pricing error	0.124	0.108	0.012	0.011
p-value	0.00	0.00	0.59	0.62
(v) Liquidity provider portfolio:				
Pricing error	0.168	0.146	0.015	0.015
p-value	0.00	0.00	0.62	0.62

Table 8: 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

Table 9: 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

$$R^p_{t,t+5} = a^p_r + b^p_r VIX^2_t + \epsilon^p_{r,t+5},$$

where $R^p_{t,t+5}$ is the cumulative excess return on portfolio p from the portfolio formation date t to $t+5$ and VIX^2_t is the squared VIX index on the portfolio formation date. The predictive loadings b^p_r 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

	Lo–Hi	2–9	3–8	4–7	5–6
Small	9.81	7.69	9.51	−0.48	−2.60
2	7.71	4.34	3.70	1.68	−0.34
3	6.72	3.13	3.00	2.65	0.44
4	5.70	5.28	3.01	3.41	−0.54
Big	9.09	5.51	3.53	3.02	0.34

Panel B: t-statistics

	Lo–Hi	2–9	3–8	4–7	5–6
Small	3.44	2.89	3.21	−0.19	−0.95
2	3.00	2.92	2.62	1.18	−0.25
3	2.98	1.98	2.59	1.85	0.63
4	2.94	3.34	2.51	2.76	−0.72
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 10: Persistence of the volatility risk of the reversal strategies

The table shows the betas of the reversal strategies at different horizons on changes in VIX^2 one day after portfolio formation. The betas are estimated by running the regressions

$$R^p_{t,t+h} \;=\; \alpha_p + \beta^p_{VIX,h}\Delta VIX^2_{t,t+1} + \epsilon^p_{t,t+5},$$

where $R^p_{t,t+h}$ is the cumulative excess return on reversal strategy portfolio p from the portfolio formation date t to $t+h$, and $\Delta VIX^2_{t,t+1}$ is the change in the squared VIX index from date t to $t+1$. The panels show the betas $\beta^p_{VIX,h}$. 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

$\beta^p_{VIX,5}$	Lo–Hi	2–9	3–8	4–7	5–6		t-statistics	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

$\beta^p_{VIX,10}$	Lo–Hi	2–9	3–8	4–7	5–6		t-statistics	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

$\beta^p_{VIX,20}$	Lo–Hi	2–9	3–8	4–7	5–6		t-statistics	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 11: Volatility co-movement portfolios

CAPM alphas, volatility risk betas, and pricing errors of the 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 (normalized using a 60-day rolling window). 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). The strategies are held for five days. Panel A shows the CAPM alphas. Panel B shows the volatility risk betas, estimated by regressing the reversal strategy returns on the change in VIX squared and the market return. Panel C shows pricing errors using a restricted price of volatility risk from option markets (the one that prices the VIX return). To obtain the pricing errors, we multiply their volatility risk betas by the restricted price of volatility risk and their market betas by the average excess market return. We then subtract these predicted average return from the actual returns, and average over time. 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: CAPM alphas

| | Alphas | | | | | | 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

Panel B: Volatility risk betas (controlling for R^M)

| | Betas | | | | | | 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: Pricing errors (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 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. $\sigma\left(\beta_{VIX^2}\Delta VIX^2\right)$. The sample is from April 9, 2001 to May 31, 2020.

Panel A: Reversal strategy average return
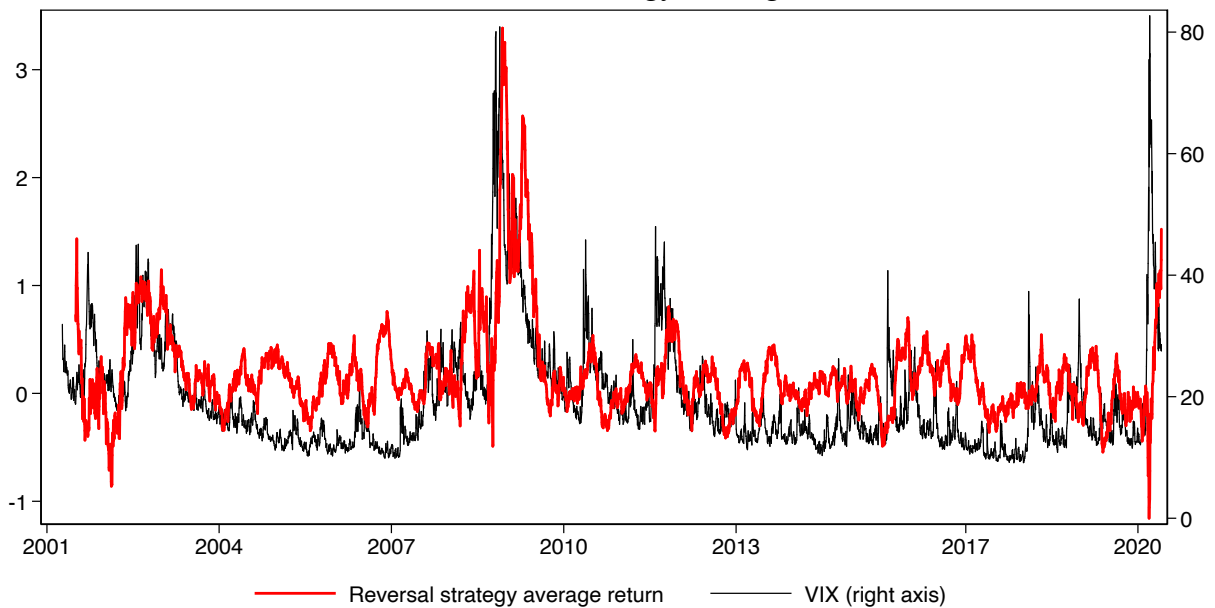
Panel B: Reversal strategy volatility risk

Figure 2: Idiosyncratic volatility and VIX

The figure shows the relationship between idiosyncratic volatility and VIX. Idiosyncratic volatility is 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. We take 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 this sum. The sample is from April 9, 2001 to May 31, 2020.

Figure 3: 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 5 and 6). 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.
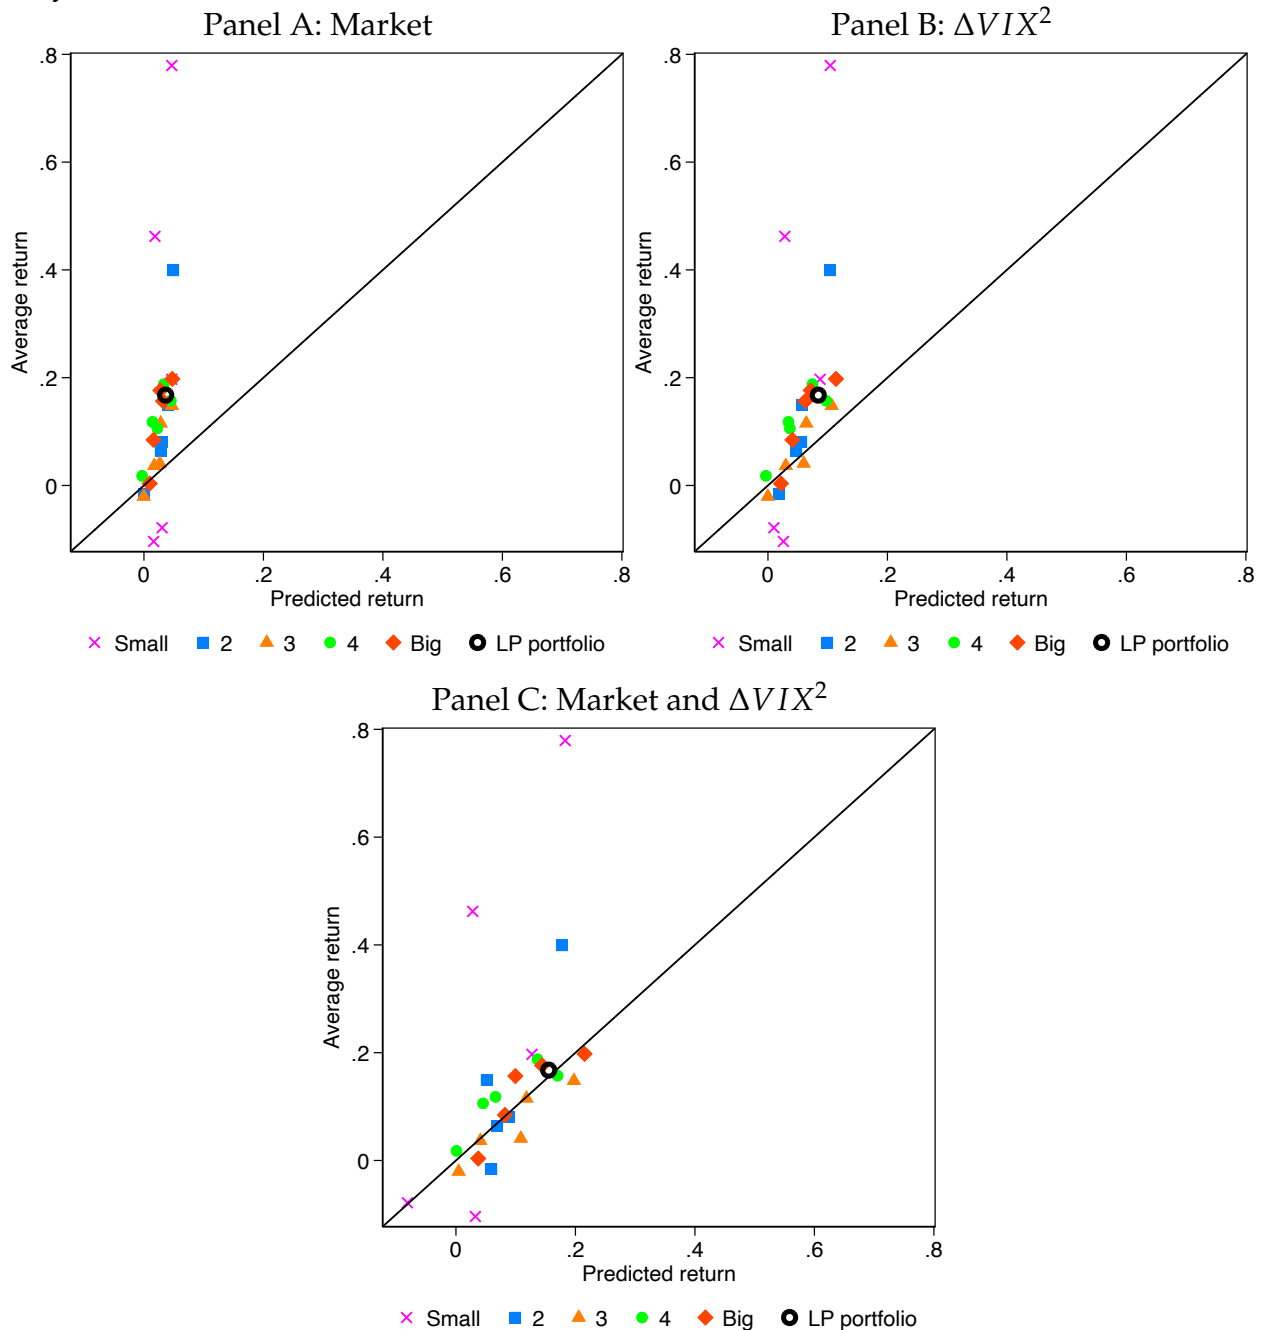
Figure 4: Option-implied price of volatility risk

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.


Panel A: Market

Panel B: ΔVIX^2

Panel C: Market and ΔVIX^2
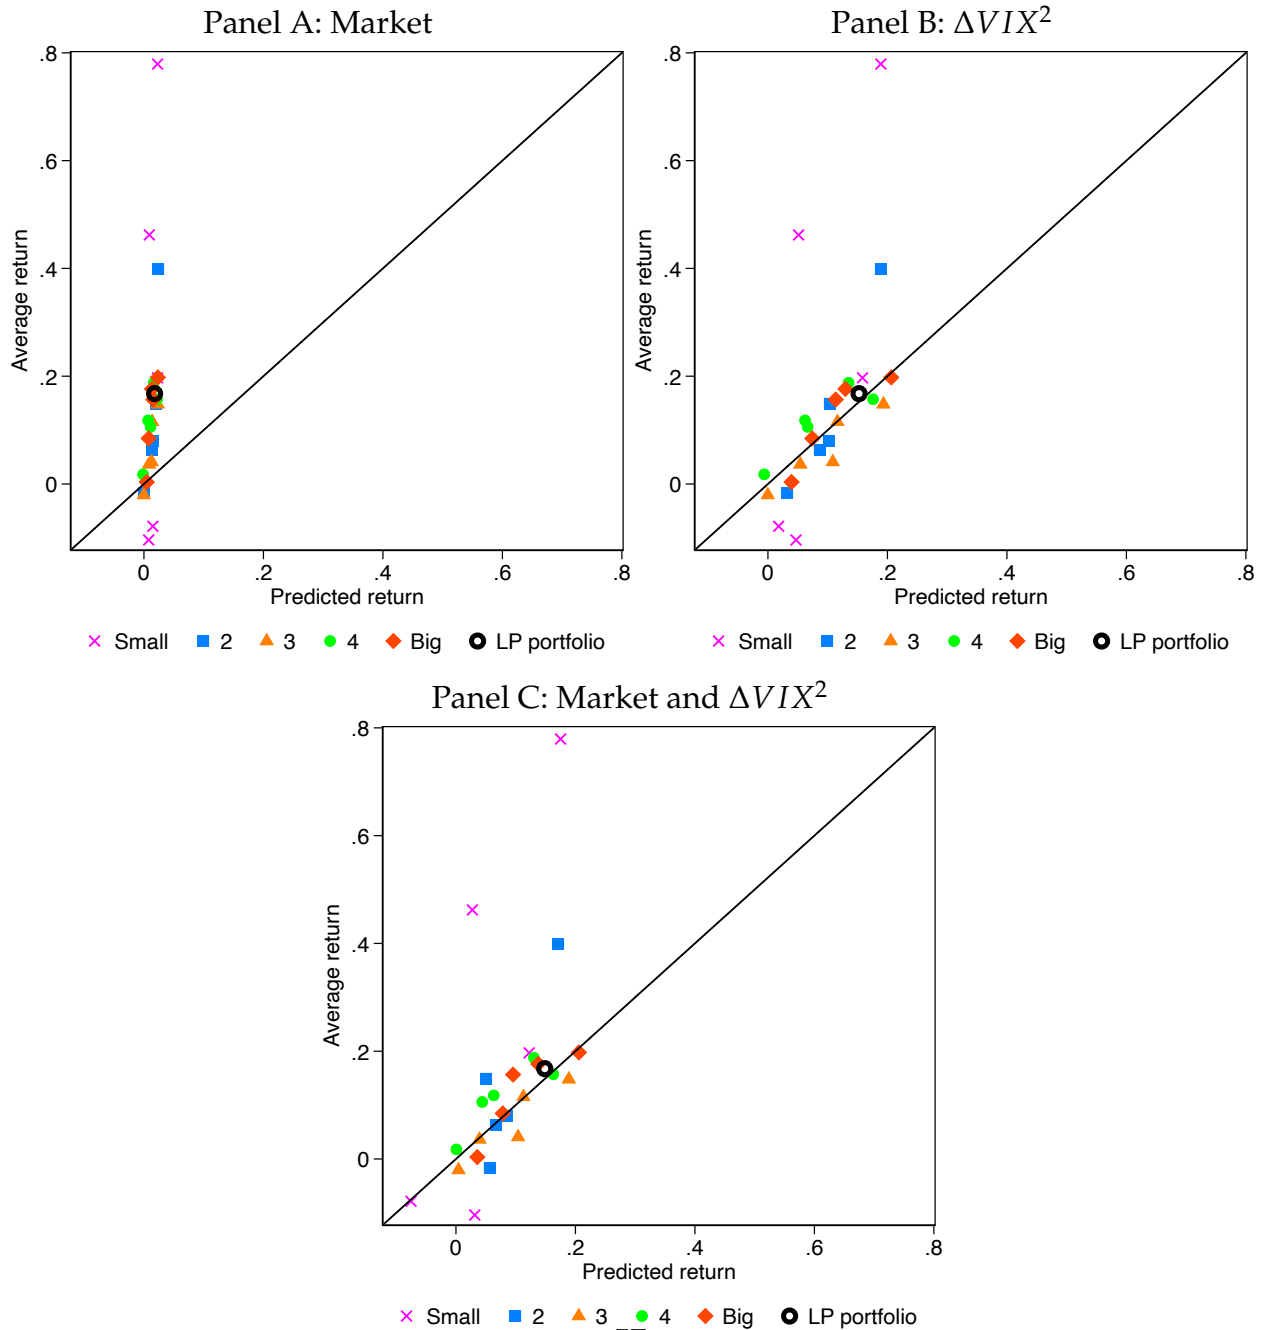
Figure 5: Returns, predictive loadings, and betas by horizon

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.

Panel A: Average returns

Panel B:
Predictive loadings on VIX^2

Panel C: ΔVIX^2 betas

Panel D:
ΔVIX^2 betas (controlling for R^M)
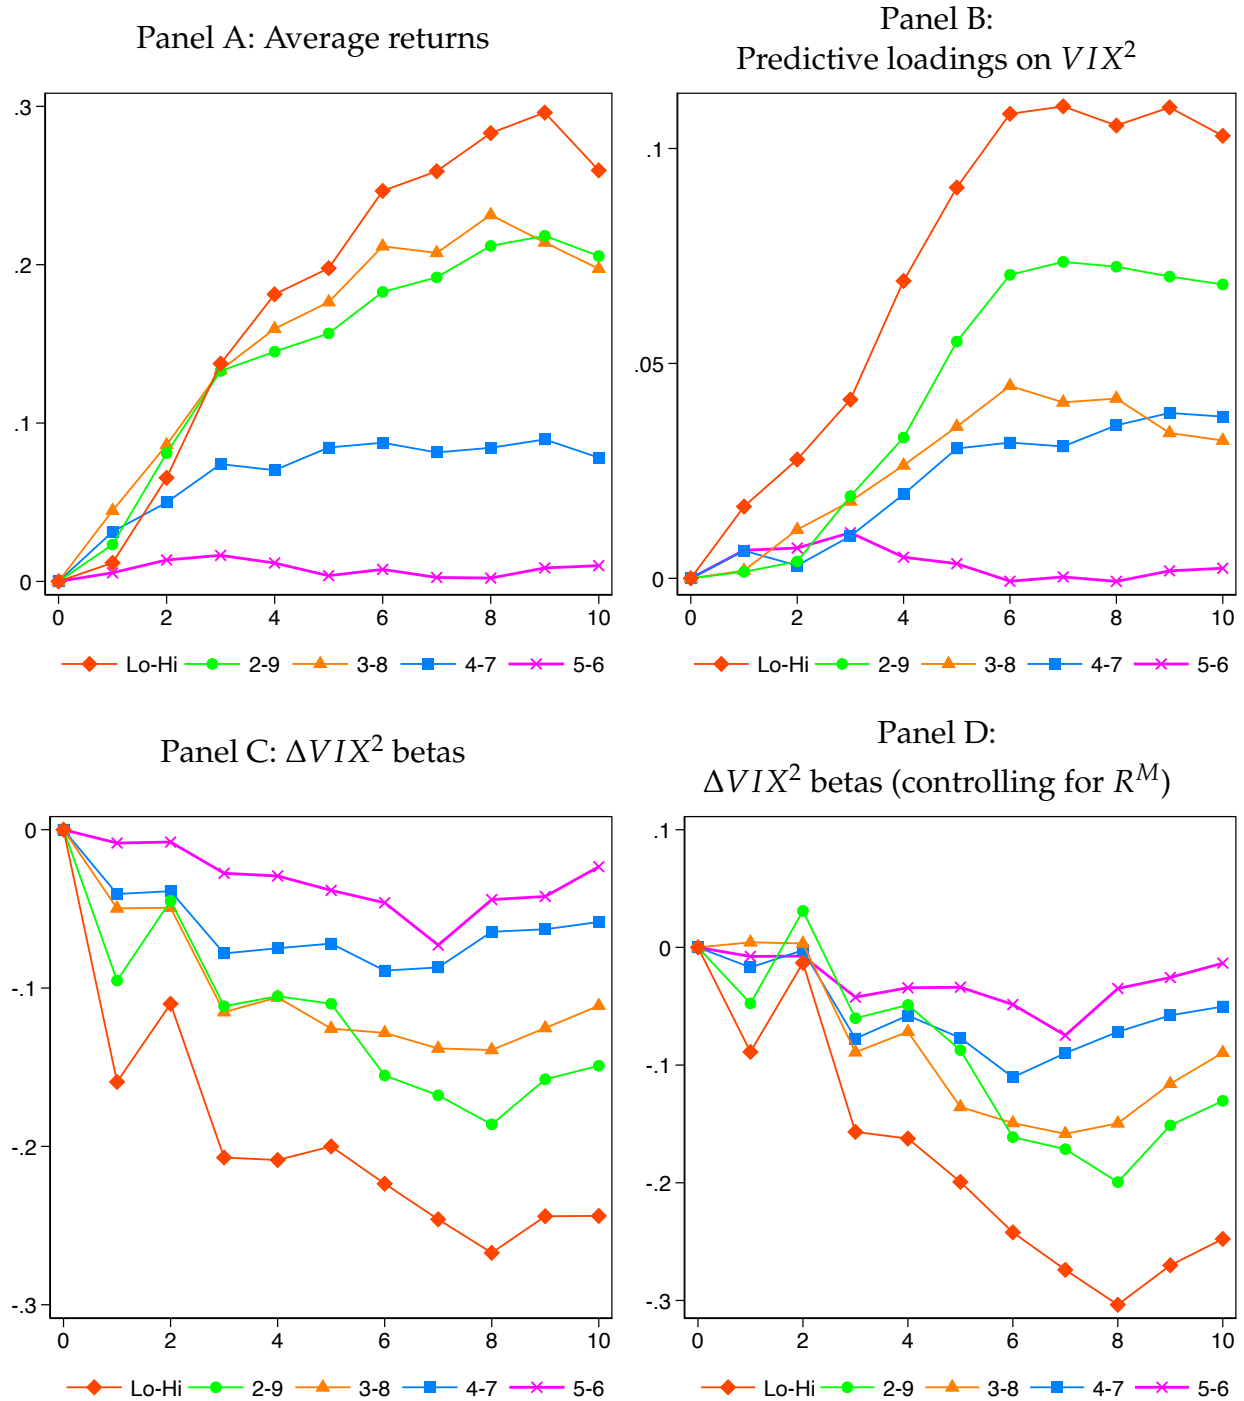
Figure 6: Volatility risk persistence of reversal returns

The figure shows the betas of reversal strategies on changes in VIX^2 on the first day after portfolio formation. The betas are estimated by running the regressions

$$R^p_{t,t+h} \;=\; \alpha_{p,h} + \beta^p_{VIX,h}\Delta VIX^2_{t,t+1} + \epsilon^p_{t,t+h},$$

where $R^p_{t,t+h}$ is the cumulative excess return on portfolio p from the portfolio formation date t to $t + h$ and $\Delta VIX^2_{t,t+1}$ 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 7: 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.

Panel A: Market — Panel B: ΔVIX^2 — Panel C: Market and ΔVIX^2
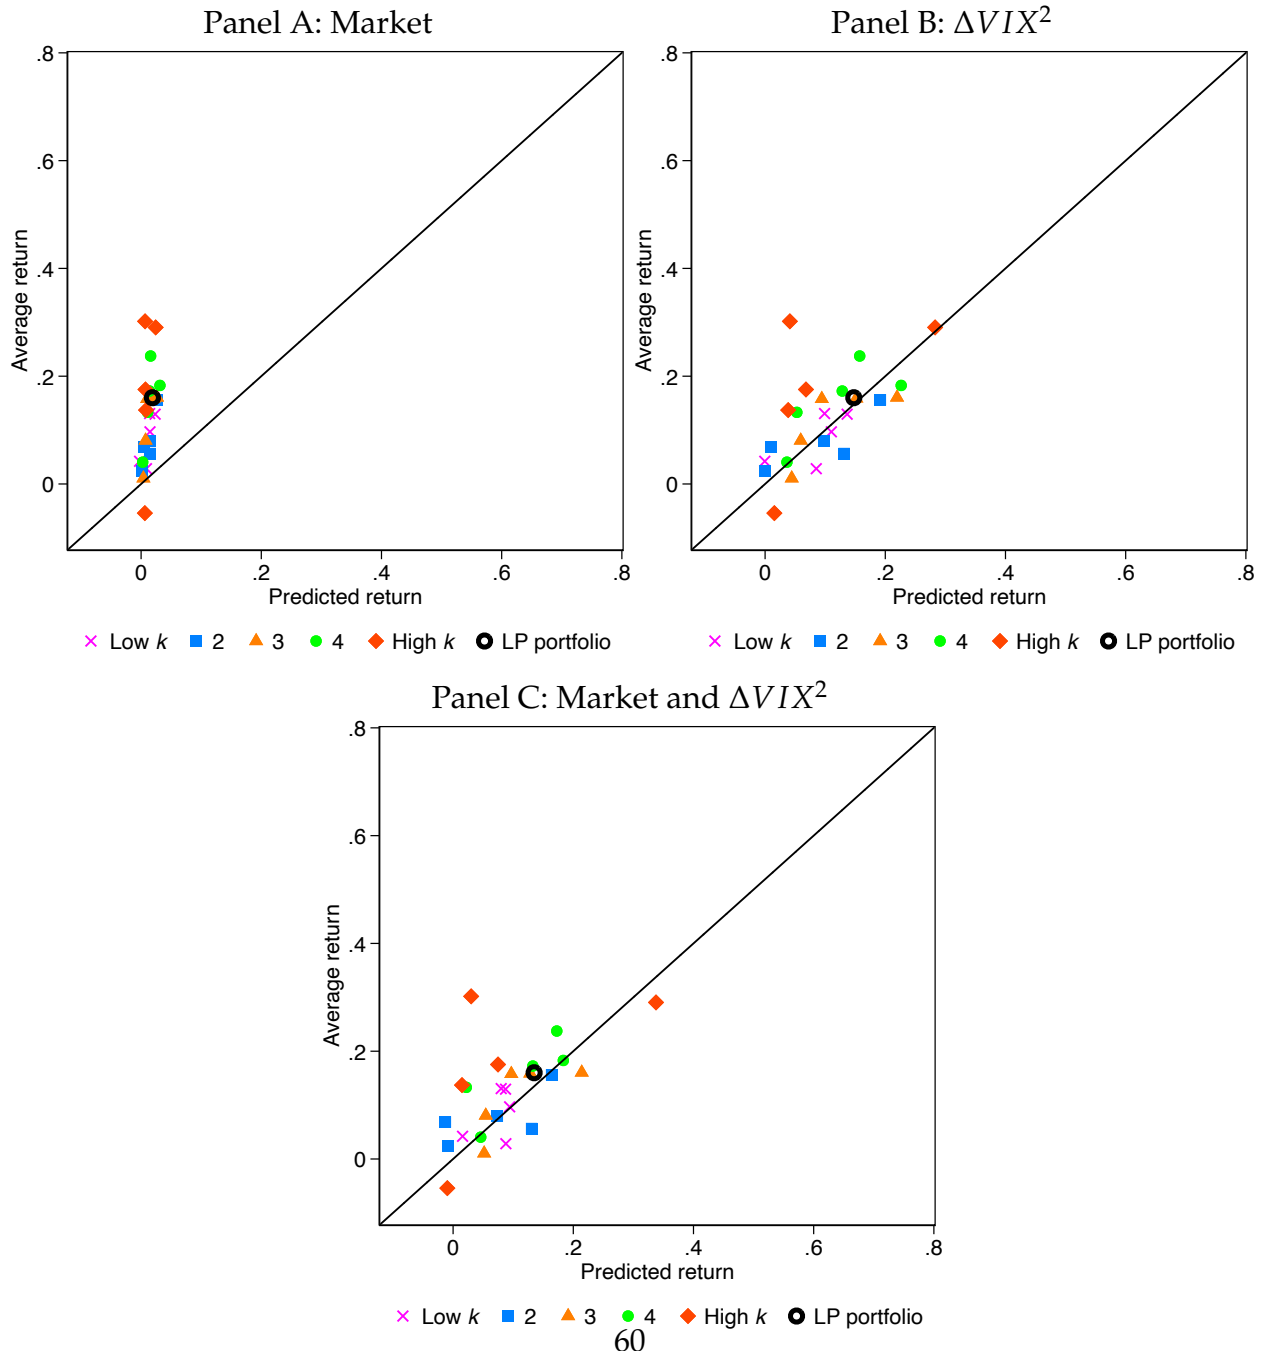
Internet Appendix for

"Liquidity and Volatility"

IA.1 Public news about fundamentals

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

$$p_{i,1} = \bar{v}_i + v_i + u_i, \tag{IA.1}$$

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

$$p_{i,t} = \bar{v}_i + E_t^Q[u_i] + \frac{\phi E_t^Q\left[\sigma_{v,i}^2\right]}{\sigma_{x,i}^2} x_i. \tag{IA.2}$$

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:

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

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

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

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

1

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

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

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 ($N \to \infty$),

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

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:

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

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: ($\sum_{i=1}^{N} x_i \left(E_\tau^Q[u_i] - E_0^Q[u_i]\right) \to 0$). In contrast, the liquidity provider has a negative exposure 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.2 Proofs

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

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

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:

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

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 (9). $\qquad \square$

Proof of Lemma 1. The result is obtained by re-arranging (9) as applied to date 0, solving for x_i, and substituting $\Delta p_{i,1} = p_{i,0} - \bar{v}_i$. $\qquad \square$

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

Proof of Lemma 3. The result is obtained by multiplying (11) by $-x_i$. $\qquad \square$

Proof of Proposition 2. The first equality is the definition of a beta. The second equality substitutes in Lemma 3 as follows:

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

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

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

$$\tag{IA.12}$$

$$= -\frac{\phi k x_i^2}{\sigma_{x,i}^2}. \tag{IA.13}$$

This gives (13). The second to last equality uses the factor structure of volatility as given by (6). The last equality uses the fact that $\varsigma_{v,i}^2$ is idiosyncratic and hence uncorrelated with σ_m^2. Finally, substitute (10) into (13) to get to get (14). $\qquad\square$

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:

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

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

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

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

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

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

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

$$p_{i,t} \;=\; E_t^Q \left[p_{i,1} \right] \;=\; \bar{v}_i + E_t^Q \left[u_i \right] + E^Q \left[v_i | \mathcal{F}_t \right]. \tag{IA.18}$$

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

$$E^Q \left[v_i | \mathcal{F}_t \right] \;=\; E^Q \left[E \left[v_i | \sigma_{v,i}^2, \mathcal{F}_t \right] \Big| \mathcal{F}_t \right] \;=\; E^Q \left[\frac{\phi \sigma_{v,i}^2}{\sigma_{x,i}^2} x_i \Big| \mathcal{F}_t \right] \;=\; \frac{\phi E_t^Q \left[\sigma_{v,i}^2 \right]}{\sigma_{x,i}^2} x_i. \tag{IA.19}$$

The last equality uses the fact that v_i is idiosyncratic and hence independent of $E_t^Q \left[u_i \right]$ conditional on $\sigma_{v,i}^2$. Plugging (IA.19) into (IA.18) gives (IA.2).

ii. Re-arranging (IA.2) to solve for $-x_i$ gives (IA.3).

4

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

$$\beta_{u_i,\sigma_m} = \frac{Cov\left(E_\tau^Q\left[u_i\right] - E_0^Q\left[u_i\right], E_\tau^Q\left[\sigma_m^2\right] - E_0^Q\left[\sigma_m^2\right]\right)}{Var\left(E_\tau^Q\left[\sigma_m^2\right] - E_0^Q\left[\sigma_m^2\right]\right)}. \qquad \text{(IA.20)}$$

Taking the covariance of the right side of (IA.4) with $E_\tau^Q\left[\sigma_m^2\right] - E_0^Q\left[\sigma_m^2\right]$ and dividing by $Var\left(E_\tau^Q\left[\sigma_m^2\right] - E_0^Q\left[\sigma_m^2\right]\right)$ 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.

$\square$

Proof of Proposition 4.

i. Part (i) follows from re-arranging (17) as applied to $t = 0$.

ii. Part (ii) follows from differencing (17) between dates τ and 0 and multiplying by $-x_i$.

$\square$

Proof of Proposition 5. The results follow from taking the covariance of the right side of (19) with $E_t^Q\left[\sigma_m^2\right] - E_0^Q\left[\sigma_m^2\right]$ and dividing by $Var\left(E_t^Q\left[\sigma_m^2\right] - E_0^Q\left[\sigma_m^2\right]\right)$ for $t \in \{\tau, 1\}$. In the cases $\gamma_{i,t} = \gamma E_t^Q\left[\sigma_m^2\right]$ or $\gamma_{i,t} = \gamma E_t^Q\left[\sigma_{v,i}^2\right]$ the term $\beta_{\gamma_i,\sigma_m}$ simplifies to γ. $\square$

Proof of Proposition 6. This result follows by taking the expected value of (19) as applied to $t \in \{\tau, 1\}$ and simplifying using Proposition 5. $\square$

IA.3 Additional empirical results

IA.3.1 Idiosyncratic volatility and market volatility

Table IA.1 assesses the relationship between idiosyncratic and market volatility formally. 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

5

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—in this case VIX squared—is an equally powerful predictor of both. This again confirms Prediction 1.

IA.3.2   Volatility co-movement portfolio sorts

Table IA.2 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 Prediction 1, the ex-post co-movement coefficients are all positive, indicating that higher market volatility is associated with higher idiosyncratic volatility.

Table IA.3 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 shows that the corresponding CAPM alphas are very similar to the raw returns and statistically significant.

Table IA.4 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.4 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.5 shows the results of Fama-Macbeth regressions with an option-implied price of risk. These results are analogous to Table 8 for our main portfolio sorts. From Table IA.5, 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.

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***	0.969***				
	(0.012)	(0.013)				
Mkt. vol.			0.961***			
			(0.006)			
VIX^2				0.913***	1.429***	
				(0.016)	(0.024)	
Mkt. var.						1.418***
						(0.008)
Constant	−1.881***	10.555***	13.934***	−0.338***	3.463***	4.550***
	(0.258)	(0.269)	(0.118)	(0.120)	(0.181)	(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: 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

	Lo–Hi	2–9	3–8	4–7	5–6
Low k	52.01	52.22	48.50	47.92	45.88
2	57.50	56.73	54.90	54.32	53.72
3	55.88	56.47	54.91	53.79	54.01
4	47.55	49.10	47.96	46.95	46.72
High k	30.74	34.21	33.94	34.75	32.90

Panel B: Idiosyncratic volatility

	Lo–Hi	2–9	3–8	4–7	5–6
Low k	1.89	2.03	2.20	2.33	2.43
2	1.54	1.62	1.69	1.74	1.77
3	1.70	1.78	1.85	1.89	1.90
4	2.08	2.18	2.25	2.28	2.30
High k	2.95	3.07	3.19	3.25	3.30

Panel C: Average turnover

	Lo–Hi	2–9	3–8	4–7	5–6
Low k	1.30	1.37	1.43	1.47	1.50
2	0.96	0.99	1.03	1.05	1.07
3	1.04	1.08	1.11	1.12	1.13
4	1.34	1.39	1.42	1.43	1.44
High k	2.00	2.09	2.13	2.18	2.19

Panel D: Sorting-day return

	Lo–Hi	2–9	3–8	4–7	5–6
Low k	−6.00	−3.18	−2.03	−1.14	−0.37
2	−4.87	−2.67	−1.71	−0.97	−0.32
3	−5.27	−2.89	−1.85	−1.05	−0.34
4	−6.23	−3.41	−2.17	−1.22	−0.39
High k	−8.57	−4.55	−2.88	−1.62	−0.53

Panel E: Pre-sorting k

	Lo–Hi	2–9	3–8	4–7	5–6
Low k	−0.39	−0.41	−0.44	−0.48	−0.49
2	0.26	0.26	0.26	0.26	0.26
3	0.68	0.68	0.68	0.68	0.68
4	1.20	1.20	1.20	1.20	1.20
High k	2.39	2.41	2.44	2.45	2.46

Panel F: Post-sorting k

	Lo–Hi	2–9	3–8	4–7	5–6
Low k	0.49	0.45	0.45	0.45	0.46
2	0.79	0.74	0.73	0.71	0.72
3	1.17	1.07	1.04	1.05	1.04
4	1.68	1.50	1.45	1.47	1.45
High k	2.51	2.06	2.04	2.08	2.06

Table IA.3: 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 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

	Lo–Hi	2–9	3–8	4–7	5–6
Low *k*	0.13	0.13	0.10	0.03	0.04
2	0.16	0.08	0.06	0.07	0.02
3	0.16	0.16	0.16	0.08	0.01
4	0.18	0.24	0.17	0.13	0.04
High *k*	0.29	0.30	0.18	0.14	−0.05

Panel B: *t*-statistics

	Lo–Hi	2–9	3–8	4–7	5–6
Low *k*	2.46	3.17	2.57	0.71	1.18
2	3.62	2.53	1.97	2.66	1.01
3	3.68	4.52	5.33	2.79	0.40
4	3.29	5.43	4.62	3.63	1.21
High *k*	3.00	5.01	3.45	2.91	−1.11

Panel C: Standard deviations

	Lo–Hi	2–9	3–8	4–7	5–6
Low *k*	3.34	2.88	2.55	2.69	2.45
2	2.66	2.14	1.97	1.88	1.70
3	2.91	2.40	2.06	1.96	1.82
4	3.55	2.99	2.58	2.46	2.43
High *k*	5.62	3.99	3.65	3.40	3.42

Panel D: Sharpe ratios

	Lo–Hi	2–9	3–8	4–7	5–6
Low *k*	0.28	0.32	0.27	0.07	0.12
2	0.42	0.26	0.20	0.26	0.10
3	0.39	0.47	0.54	0.29	0.04
4	0.37	0.56	0.47	0.38	0.12
High *k*	0.37	0.54	0.34	0.29	−0.11

Panel E: CAPM alphas

	Lo–Hi	2–9	3–8	4–7	5–6
Low *k*	0.10	0.11	0.08	0.02	0.04
2	0.13	0.06	0.04	0.06	0.02
3	0.14	0.14	0.15	0.07	0.01
4	0.15	0.22	0.16	0.12	0.04
High *k*	0.26	0.29	0.16	0.13	−0.06

Panel F: CAPM alpha *t*-statistics

	Lo–Hi	2–9	3–8	4–7	5–6
Low *k*	2.00	2.73	2.14	0.47	1.21
2	2.99	2.04	1.41	2.50	0.82
3	3.25	3.96	4.94	2.42	0.26
4	2.72	4.99	4.26	3.41	1.09
High *k*	2.67	4.74	3.17	2.70	−1.22

Table IA.4: 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

$$R^p_{t,t+5} = \alpha_p + \beta^p_{VIX}\Delta VIX^2_{t,t+5} + \beta^p_M R^M_{t,t+5} + \epsilon^p_{t,t+5},$$

where $R^p_{t,t+5}$ is the cumulative excess return on reversal strategy portfolio p from the portfolio formation date t to $t+5$, $R^M_{t,t+5}$ is the excess return on the market portfolio, and $\Delta VIX^2_{t,t+5}$ 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^p_M. 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^p_{VIX}						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^p_{VIX}						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^p_M						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.5: 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