

What Is the Expected Return on a Stock?

IAN W. R. MARTIN and CHRISTIAN WAGNER*

ABSTRACT

We derive a formula for the expected return on a stock in terms of the risk-neutral variance of the market and the stock's excess risk-neutral variance relative to that of the average stock. These quantities can be computed from index and stock option prices; the formula has no free parameters. The theory performs well empirically both in and out of sample. Our results suggest that there is considerably more variation in expected returns, over time and across stocks, than has previously been acknowledged.

IN THIS PAPER, WE DERIVE A NEW formula that expresses the expected return on a stock in terms of the risk-neutral variance of the market, the risk-neutral variance of the individual stock, and the value-weighted average of stocks' risk-neutral variance. Then we show that the formula performs well empirically.

The inputs to the formula—the three measures of risk-neutral variance—are computed directly from option prices. As a result, our approach has some distinctive features that separate it from more conventional approaches to the cross section.

First, as it is based on current market prices rather than, say, accounting information, it can in principle be implemented in real time. Nor does it require that we use any historical information. It thus represents a parsimonious alternative to pooling data on many firm characteristics (as, for instance, in Lewellen (2015)).

*Ian Martin is at the London School of Economics. Christian Wagner is at Copenhagen Business School. We thank Harjoat Bhamra; John Campbell; Patrick Gagliardini; Christian Julliard; Marcin Kacperczyk; Binying Liu; Dong Lou; Stefan Nagel (Editor); Christopher Polk; Tarun Ramadorai; Tyler Shumway; Paul Schneider; Andrea Tamoni; Fabio Trojani; Dimitri Vayanos; Tuomo Vuolteenaho; participants at the Western Finance Association Meetings 2017, the 2017 Annual Meeting of the Society for Economic Dynamics, the 2017 CEPR Spring Symposium, the BI-SHoF Conference in Asset Pricing, the AP2-CFF Conference on Return Predictability, the 2016 IFSID Conference on Derivatives, the 4nations Cup 2016; seminar participants at Arrowstreet, AQR, Banca d'Italia, BlackRock, CEMFI, the European Central Bank, the London School of Economics, MIT (Sloan), NHH Bergen, Norges Bank Investment Management, the University of Maryland, the University of Michigan (Ross), and WU Vienna; and two anonymous referees for their comments. Ian Martin is grateful for support from the Paul Woolley Centre and from the ERC under Starting Grant 639744. Christian Wagner acknowledges support from the FRIC Center for Financial Frictions, grant no. DNR102. We have read the *Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

DOI: 10.1111/jofi.12778

Second, our formula provides conditional forecasts at the level of the individual stock. Rather than asking, say, what the unconditional average expected return is on a portfolio of small value stocks, we can ask, what is the expected return on Apple, today?

Third, the formula makes specific quantitative predictions about the relationship between expected returns and the three measures of risk-neutral variance. It does not require estimation of any parameters. This can be contrasted with factor models, in which both factor loadings and the factors themselves are estimated from the data (with all of the associated concerns about data snooping). There is a closer comparison with the Capital Asset Pricing Model (CAPM), which makes a specific prediction about the relationship between expected returns and betas, but even the CAPM requires that the forward-looking betas that come out of theory be estimated based on historical data.

Our approach does not have this deficiency and, as we will show, it performs better empirically than the CAPM. But, like the CAPM, it requires that we take a stance on the conditionally expected return on the market. We do so by applying the results of Martin (2017), who argues that the risk-neutral variance of the market provides a lower bound on the equity premium. In particular, we exploit Martin's more aggressive claim that, empirically, the lower bound is approximately tight, so that risk-neutral variance directly measures the equity premium. We also present results that avoid dependence on this claim, however, by forecasting expected returns *in excess of the market*. In doing so, we isolate the purely cross-sectional predictions of our framework that are independent from the market-timing issue of forecasting the equity premium. As these predictions exploit the cross section as well as the time series, the associated empirical results are stronger in a statistical sense than those of Martin (2017).

We introduce the theoretical framework in Section I. We then show how to construct the three risk-neutral variance measures, and discuss some of their properties, in Section II.

Our main empirical results are presented in Section III. We test the framework for S&P 100 and S&P 500 stocks at forecast horizons ranging from 1 month to 2 years. Papers in the predictability literature typically aim to uncover variables that are statistically significant in forecasting regressions. We share this goal, of course, but as our model makes predictions about the quantitative relationship between expected returns and risk-neutral variances, we hope also to find that the estimated coefficients on the predictor variables are close to *specific numbers* that come out of the theory. For most specifications, we find that that we do not reject the model, whereas we reject the null hypothesis of no predictability at the 6-month, 1-year, and 2-year horizons.

In Section IV, we examine how our findings relate to stock characteristics. Notably, we run panel regressions of realized returns onto beta, size, book-to-market, and past returns. In our sample, size and book-to-market are statistically significant forecasters of excess returns, though not of returns in excess of the market. When we include our predictive variables based on risk-neutral variance, these characteristics become statistically insignificant, but the

risk-neutral variance variables themselves are significant predictors (of both excess returns and excess-of-market returns). Moreover, they enter with coefficients that are insignificantly different from those predicted by our theory. In a similar vein, we show that returns on portfolios sorted on the characteristics are consistent with the model.

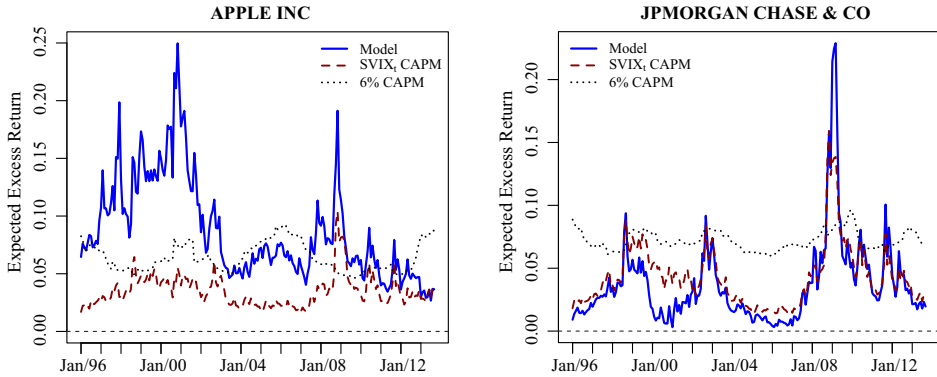
In Section V, we assess the out-of-sample predictive performance of the formula when its coefficients are constrained to equal the values implied by the theory. We compute out-of-sample R^2 coefficients that compare the formula's predictions to those of a range of competitors, as in Goyal and Welch (2008). We start by comparing against competitors that are themselves out-of-sample predictors (in the sense of being based on a priori considerations, without in-sample information). The formula outperforms all such competitors at horizons of 3, 6, 12, and 24 months, both for expected returns and for expected returns in excess of the market.

More ambitiously, we next compare the formula against competitors that have in-sample information. At the 6- and 12-month horizons, the only case in which our model of expected excess returns "loses" is when we allow the competitor predictor to know both the in-sample average realized return across stocks *and* the multivariate in-sample relationship between realized returns and beta, size, book-to-market, and past returns. When we allow the competitor to know only the in-sample average and the univariate relationship between realized returns and any one of the characteristics, our formula outperforms. Even more strikingly, in the case in which we forecast returns in excess of the market, the formula outperforms the competitor armed with knowledge of the in-sample average and of the multivariate relationship.


These empirical successes are particularly notable because the formula makes some dramatic predictions about stock returns. Figure 1 plots the time series of expected excess returns, relative to the riskless asset and relative to the market, for Apple and JPMorgan Chase & Co. over the period January 1996 to October 2014. According to our model, expected returns spiked for both stocks during the depths of the financial crisis of 2008 to 2009. In the case of Apple, this largely reflected a high market-wide equity premium rather than an Apple-specific phenomenon, whereas JP Morgan Chase's expected excess return was high even relative to the market risk premium. The figure also plots expected excess returns computed using the CAPM with 1-year rolling historical betas and the equity premium computed from the SVIX index of Martin (2017), or fixed at 6%, to illustrate the point (which, as we will show, holds more generally) that our model generates more volatility in expected returns, both over time and in the cross section, than does the CAPM.

We conclude in Section VI. The Appendix contains a discussion of the relationship between our volatility measures and implied correlation, and provides details of the bootstrap procedure. Finally, further empirical results and analysis of the finite-sample properties of our block bootstrap procedure are provided in an Internet Appendix.¹

¹ The Internet Appendix may be found in the online version of this article.



Panel A. Expected Excess Returns



Panel B. Expected Returns in Excess of the Market

Figure 1. Expected excess returns and expected returns in excess of the market, annual horizon. This figure illustrates our results. It plots the time series of expected excess returns and expected returns in excess of the market for Apple Inc. and JP Morgan Chase & Co. at an annual horizon (solid line) and, for comparison, the corresponding time series using the CAPM with a constant equity premium of 6% (dotted line) or an equity premium calculated using the SVIX index (dashed line). (Color figure can be viewed at wileyonlinelibrary.com)

Related Literature. A large literature documents the importance of idiosyncratic volatility for future stock returns, though draws mixed conclusions. For instance, Ang et al. (2006) find a negative relation both for total volatility and for idiosyncratic volatility (defined as the residual variance of Fama and French (1993) three-factor regressions on daily returns over the past month). By contrast, Fu (2009) finds a positive relation when idiosyncratic volatility is measured by the conditional variance obtained from fitting an EGARCH model to residuals of Fama and French (1993) regressions on monthly returns.

Our model attributes an important role to average stock variance (measured as the value-weighted sum of individual stock risk-neutral variances), a prediction that we confirm empirically. This result echoes the finding of Herskovic et al. (2016) that idiosyncratic volatility (measured from past returns) exhibits a strong factor structure and that firms' loadings on the common component predict equity returns. Furthermore, our measure of average stock variance may capture a potential factor structure in the cross section of equity options, as documented by Christoffersen, Fournier, and Jacobs (2017) across 29 Dow Jones firms.

Various authors have explored the forecasting power of options-based measures. An et al. (2014) find that increases in implied volatilities of at-the-money call and put options have opposing implications, predicting high and low subsequent stock returns, respectively. Conrad, Dittmar, and Ghysels (2013) study the relationship between risk-neutral moments and realized returns and find a negative, though not statistically significant, relationship between risk-neutral variance and subsequent stock returns. They work with the risk-neutral variance of log returns (following Bakshi, Kapadia, and Madan (2003)), so their volatility indices load particularly strongly on the prices of deep out-of-the-money put options. In contrast to both of these papers, our theoretical results lead us to focus on the risk-neutral variance of index- and stock-level *simple* returns. The resulting volatility indices load equally on the prices of options of all strikes.

Other papers work within the CAPM and attempt to estimate betas more accurately by incorporating forward-looking information from options. French, Groth, and Kolar (1983) estimate beta using a stock's historical return correlation with the market and option-implied volatilities for the stock and the market. Buss and Vilkov (2012) take a similar approach, but estimate correlation from a parametric model that links correlation under the risk-neutral and the objective measure. Chang et al. (2012) make assumptions under which expected correlation can be computed from the ratio of option-implied stock to market skewness; this implies, however, that a firm's implied beta will be positive only if its skewness has the same sign as market skewness, so will typically not provide a meaningful CAPM beta for firms with positive skewness.

In a more closely related, and contemporaneous, paper, Kadan and Tang (2018) adapt an idea of Martin (2017) to derive a lower bound on expected stock returns. To understand the main differences between their approach and ours, recall that Martin starts from an identity that relates the equity premium to a risk-neutral variance term and a (real-world) covariance term. He exploits the identity by arguing that a *negative correlation condition* (NCC) holds for the market return, so that the covariance term is nonpositive in quantitatively reasonable models of financial markets. If so, the risk-neutral variance of the market provides a lower bound on the equity premium.² Kadan and Tang (2018) modify this approach to derive a lower bound for expected stock returns based

² Schneider and Trojani (2019) propose a related approach to forecasting the equity premium based, in part, on variants of the NCC.

on an NCC for individual stocks. But it is trickier to make the argument that the NCC should hold at the individual stock level, so their lower bound applies only for a subset of S&P 500 stocks.

I. Theory

Our starting point is the gross return with maximal expected log return: call it $R_{g,t+1}$, so $\mathbb{E}_t \log R_{g,t+1} \geq \mathbb{E}_t \log R_{i,t+1}$ for any gross return $R_{i,t+1}$. This *growth-optimal return* has the special property, unique among returns, that $1/R_{g,t+1}$ is a stochastic discount factor (SDF). To see this, note that it is attained by choosing portfolio weights $\{g_n\}_{n=1}^N$ on the tradable returns (on stocks, stock options, index options, and the riskless asset) to solve

$$\max_{\{g_n\}_{n=1}^N} \mathbb{E} \log \sum_{n=1}^N g_n R_{n,t+1} \quad \text{such that} \quad \sum_{n=1}^N g_n = 1.$$

The first-order conditions for this problem are

$$\mathbb{E} \left(\frac{R_{i,t+1}}{\sum_{n=1}^N g_n R_{n,t+1}} \right) = \psi \quad \text{for all } i,$$

where ψ is a Lagrange multiplier; we follow Roll (1973) and Long (1990) in assuming that these first-order conditions have an interior solution. Multiplying by g_i and summing over i , we see that $\psi = 1$, and hence that the reciprocal of $R_{g,t+1} \equiv \sum_{n=1}^N g_n R_{n,t+1}$ is an SDF.

We denote by \mathbb{E}_t^* the associated risk-neutral expectation (more precisely, the time $t + 1$ forward-neutral expectation), which is defined via

$$\frac{1}{R_{f,t+1}} \mathbb{E}_t^* X_{t+1} = \mathbb{E}_t \left(\frac{X_{t+1}}{R_{g,t+1}} \right), \tag{1}$$

where $R_{f,t+1}$ is the gross riskless rate from time t to time $t + 1$.

A useful perspective to keep in mind is that of an unconstrained log investor who is marginal for all asset prices, including options, but chooses to invest his or her wealth fully in the market, whose gross return we write as $R_{m,t+1}$. (Martin (2017) and Kremens and Martin (2019) show that this represents a sensible benchmark when forecasting returns on the market and on currencies, respectively.) Such an investor must perceive the market itself as growth optimal, so that if \mathbb{E}_t represents the expectations of the log investor, equation (1) and subsequent equations hold with $R_{g,t+1} = R_{m,t+1}$.

The key property of the growth-optimal portfolio, which follows directly from (1), is that

$$\mathbb{E}_t \frac{R_{i,t+1}}{R_{f,t+1}} - 1 = \text{cov}_t^* \left(\frac{R_{i,t+1}}{R_{f,t+1}}, \frac{R_{g,t+1}}{R_{f,t+1}} \right) \quad \text{for all stocks } i. \tag{2}$$

Thus an asset's risk premium is determined by its risk-neutral covariance with the growth-optimal return.

We start by projecting stock returns onto the growth-optimal portfolio under the risk-neutral measure. That is, for every stock i we decompose

$$\frac{R_{i,t+1}}{R_{f,t+1}} = \alpha_{i,t}^* + \beta_{i,t}^* \frac{R_{g,t+1}}{R_{f,t+1}} + \epsilon_{i,t+1}, \tag{3}$$

where

$$\beta_{i,t}^* = \frac{\text{cov}_t^* \left(\frac{R_{i,t+1}}{R_{f,t+1}}, \frac{R_{g,t+1}}{R_{f,t+1}} \right)}{\text{var}_t^* \frac{R_{g,t+1}}{R_{f,t+1}}} \tag{4}$$

$$\mathbb{E}_t^* \epsilon_{i,t+1} = 0 \tag{5}$$

$$\text{cov}_t^* (\epsilon_{i,t+1}, R_{g,t+1}) = 0. \tag{6}$$

Equations (3) to (5) define $\epsilon_{i,t+1}$, $\beta_{i,t}^*$, and $\alpha_{i,t}^*$, and equation (6) follows as a consequence of equations (3) to (5). Thus the only assumption embodied in equations (3) to (6) is that the appropriate risk-neutral moments exist and are finite, and that $\text{var}_t^* R_{g,t+1}/R_{f,t+1}$ is nonzero. (This last assumption is needed for (4) to be well defined. It rules out the theoretically interesting, but empirically implausible, possibility that the risk-neutral and true probability measures coincide, as in that case the growth-optimal portfolio is riskless.)

It may be helpful to compare this approach to that of Hansen and Richard (1987), who also projected arbitrary returns onto a “distinguished” return—in their case, the minimum-second-moment return, $R_{*,t+1}$, which is proportional to an SDF. This return has the key property that $\mathbb{E}_t(R_{*,t+1}R_{i,t+1}) = \mathbb{E}_t(R_{*,t+1}^2)$ for all tradable returns $R_{i,t+1}$, and hence that

$$\mathbb{E}_t \frac{R_{i,t+1}}{R_{f,t+1}} - 1 = -\frac{R_{f,t+1}}{\mathbb{E}_t R_{*,t+1}} \text{cov}_t \left(\frac{R_{i,t+1}}{R_{f,t+1}}, \frac{R_{*,t+1}}{R_{f,t+1}} \right) \text{ for all stocks } i. \tag{2'}$$

This equation says that risk premia are determined by *true* covariances with a tradable return. It can be rewritten³ as

$$\mathbb{E}_t \frac{R_{i,t+1}}{R_{f,t+1}} - 1 = -(1 + S_t^2) \text{cov}_t \left(\frac{R_{i,t+1}}{R_{f,t+1}}, \frac{R_{*,t+1}}{R_{f,t+1}} \right),$$

where S_t is the maximal conditional Sharpe ratio at time t .

³The results of Hansen and Jagannathan (1991) show that $R_{f,t+1}/\mathbb{E}_t R_{*,t+1} = \mathbb{E}_t(R_{*,t+1}^2)/(\mathbb{E}_t R_{*,t+1})^2 = 1 + S_t^2$. The first equality follows from the key property of $R_{*,t+1}$. The second holds because $R_{*,t+1}$ lies, by definition, at the tangency point of an origin-centered circle to the lower edge of the minimum variance frontier in a mean–standard-deviation diagram.

Equation (2') motivates the decomposition

$$\frac{R_{i,t+1}}{R_{f,t+1}} = \alpha_{i,t} + \beta_{i,t} \frac{R_{*,t+1}}{R_{f,t+1}} + u_{i,t+1}, \quad (3')$$

where

$$\beta_{i,t} = \frac{\text{cov}_t\left(\frac{R_{i,t+1}}{R_{f,t+1}}, \frac{R_{*,t+1}}{R_{f,t+1}}\right)}{\text{var}_t \frac{R_{*,t+1}}{R_{f,t+1}}} \quad (4')$$

$$\mathbb{E}_t u_{i,t+1} = 0 \quad (5')$$

$$\text{cov}_t(u_{i,t+1}, R_{*,t+1}) = 0. \quad (6')$$

We spell this out explicitly to emphasize the analogy between the two approaches. As before, equations (3') to (5') define $u_{i,t+1}$, $\beta_{i,t}$, and $\alpha_{i,t}$, and equation (6') follows as a consequence of equations (3') to (5').⁴ Equations (2') to (6') can be viewed as the theoretical foundation of the factor pricing literature.

But forward-looking real-world covariances are not directly observable, so they must be estimated from time-series data. Such estimates will only approximate the true forward-looking covariances if the econometric environment is sufficiently stable (ergodic, stationary) in a statistical sense. Thus to make these equations empirically useful, one needs to make further assumptions about the factors that must be included to provide a tolerable approximation to the true minimum-second-moment return, about the stochastic properties of $u_{i,t+1}$ across assets and over time, and about the stability of conditional betas over appropriate time horizons.

Broadly speaking, our approach may have a particular advantage at times when beliefs adjust suddenly, whether because of the arrival of information—firm-specific or macroeconomic news, a terrorist attack, natural or unnatural disaster, or something else—or because of a shift in market sentiment or risk aversion. Backward-looking historical covariances will adjust sluggishly at such times, which may be of particular interest to investors, decision makers inside firms, and policy makers who must respond rapidly to changing conditions. By contrast option prices, and hence our formulas, will react almost instantly.

That said, we also need to make assumptions to make our approach implementable in practice. Equations (2) and (4) together imply that

$$\mathbb{E}_t \frac{R_{i,t+1}}{R_{f,t+1}} - 1 = \beta_{i,t}^* \text{var}_t^* \frac{R_{g,t+1}}{R_{f,t+1}}. \quad (7)$$

⁴ By taking risk-neutral expectations of equation (3), we see that $\alpha_{i,t}^* = 1 - \beta_{i,t}^*$. Similarly, by taking real-world expectations of equation (3') and using equation (2') together with the properties of $R_{*,t+1}$ mentioned in footnote 3, we find that $\alpha_{i,t} = 1 - \beta_{i,t}$.

We also have, from equations (3) and (6), that

$$\text{var}_t^* \frac{R_{i,t+1}}{R_{f,t+1}} = \beta_{i,t}^{*2} \text{var}_t^* \frac{R_{g,t+1}}{R_{f,t+1}} + \text{var}_t^* \epsilon_{i,t+1}. \tag{8}$$

What we would *like* to measure is the right-hand side of (7). What we *can* measure is the left-hand side of (8) (as we will show in the next section). To connect the two, we make two assumptions.⁵

First, we approximate $\beta_{i,t}^{*2}$ in equation (8) by linearizing $\beta_{i,t}^{*2} \approx 2\beta_{i,t}^* - k$, where k is a constant. This approximation is reasonable if $\beta_{i,t}^*$ is not too far from one for a typical stock.⁶ In Internet Appendix Section I, we derive the residual that the approximation neglects, and we argue that it is small for most stocks in our sample. We therefore replace (8) with

$$\text{var}_t^* \frac{R_{i,t+1}}{R_{f,t+1}} = (2\beta_{i,t}^* - k) \text{var}_t^* \frac{R_{g,t+1}}{R_{f,t+1}} + \text{var}_t^* \epsilon_{i,t+1}. \tag{9}$$

Using (7) and (9) to eliminate the dependence on $\beta_{i,t}^*$, we have

$$\mathbb{E}_t \frac{R_{i,t+1}}{R_{f,t+1}} - 1 = \frac{1}{2} \text{var}_t^* \frac{R_{i,t+1}}{R_{f,t+1}} + \frac{k}{2} \text{var}_t^* \frac{R_{g,t+1}}{R_{f,t+1}} - \frac{1}{2} \text{var}_t^* \epsilon_{i,t+1}. \tag{10}$$

To make further progress, let $w_{i,t}$ be the market-capitalization weight of stock i in the index. Value-weighting the above equation, we find that

$$\mathbb{E}_t \frac{R_{m,t+1}}{R_{f,t+1}} - 1 = \frac{1}{2} \sum_j w_{j,t} \text{var}_t^* \frac{R_{j,t+1}}{R_{f,t+1}} + \frac{k}{2} \text{var}_t^* \frac{R_{g,t+1}}{R_{f,t+1}} - \frac{1}{2} \sum_j w_{j,t} \text{var}_t^* \epsilon_{j,t+1}. \tag{11}$$

⁵ One might have expected that the risk-neutral covariance that appears in equations (2) and (7) should be observable directly from asset prices without any further assumptions. But Martin (2018) shows that it is in general hard to measure risk-neutral expectations of functions—here, products—of multiple asset returns using asset prices that are observable in practice. (For an exception to this rule, see Kremens and Martin (2019), who exploit quanto contracts to infer risk-neutral covariances between the S&P 500 index and currencies.) In our setting, we could have followed an alternative approach if outperformance options written on $R_{i,t+1} - R_{m,t+1}$ (as a proxy for the theoretical ideal, namely, outperformance options on $R_{i,t+1} - R_{g,t+1}$) were observable and liquid; unfortunately they are not.

⁶ If $k = 1$, this linearization is the tangent to $\beta_{i,t}^{*2}$ at $\beta_{i,t}^* = 1$. Alternatively, if, say, the cross section of betas has mean one and standard deviation σ , then one could set $k = 1 - \sigma^2$ to minimize the mean squared approximation error. As we will see shortly, the precise value of k turns out not to be important. The choice to linearize around $\beta_{i,t}^* = 1$ is not critical, though we think it is natural: if the equal-weighted portfolio of stocks is approximately growth optimal, then $\beta_{i,t}^*$ is close to one on average, while if the market is approximately growth optimal, then $\beta_{i,t}^*$ is close to one on value-weighted average. More generally, we could linearize $\beta_{i,t}^{*2} \approx c\beta_{i,t}^* + d$ for appropriately chosen c and d . For example, the tangent to $\beta_{i,t}^{*2}$ at $\beta_{i,t}^* = \beta_0$, some constant, corresponds to $c = 2\beta_0$ and $d = -\beta_0^2$, or one might want to choose c and d to achieve some other goal (e.g., to minimize the mean squared error for a given distribution of $\beta_{i,t}^*$). If one takes this approach, equations (14) and (15) are unchanged except that $1/2$ is replaced by $1/c$; in particular, the value of d drops out. (See Internet Appendix Section I.) Our empirical results suggest that it is reasonable to linearize around one, that is, to set $c = 2$.

Subtracting (11) from (10), we have

$$\begin{aligned} \mathbb{E}_t \frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} &= \frac{1}{2} \left(\text{var}_t^* \frac{R_{i,t+1}}{R_{f,t+1}} - \sum_j w_{j,t} \text{var}_t^* \frac{R_{j,t+1}}{R_{f,t+1}} \right) \\ &\quad - \frac{1}{2} \left(\text{var}_t^* \epsilon_{i,t+1} - \sum_j w_{j,t} \text{var}_t^* \epsilon_{j,t+1} \right). \end{aligned} \tag{12}$$

Our second assumption is that the final term on the right-hand side of (12), which is zero on value-weighted average, can be captured by a time-invariant stock fixed effect α_i . This fixed-effects formulation, which is econometrically convenient, would follow immediately if, for example, the risk-neutral variances of residuals decompose separably, $\text{var}_t^* \epsilon_{i,t+1} = \phi_i + \psi_t$, and value weights are constant over time.

It will be convenient to define three different measures of risk-neutral variance:

$$\begin{aligned} \text{SVIX}_t^2 &= \text{var}_t^* (R_{m,t+1}/R_{f,t+1}) \\ \text{SVIX}_{i,t}^2 &= \text{var}_t^* (R_{i,t+1}/R_{f,t+1}) \\ \overline{\text{SVIX}}_t^2 &= \sum_i w_{i,t} \text{SVIX}_{i,t}^2. \end{aligned} \tag{13}$$

These measures can be computed directly from option prices, as we show in the next section. The SVIX_t index was introduced by Martin (2017)—the name echoes the related VIX index—but the definitions of stock-level $\text{SVIX}_{i,t}$ and of $\overline{\text{SVIX}}_t$, which measures average stock volatility, are new to this paper. Introducing these definitions into (12), we arrive at our first, purely relative, prediction about the cross section of expected returns in excess of the market (hereafter excess-of-market returns, for short):

$$\frac{\mathbb{E}_t R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha_i + \frac{1}{2} (\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2). \tag{14}$$

We test this prediction by running a panel regression of excess-of-market returns of individual stocks i onto stock fixed effects and excess stock variance $\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2$.

To answer the question posed in the title of the paper, we must also take a view on the expected return on the market itself. To do so, we exploit a result of Martin (2017), who argues that the SVIX index can be used as a forecast of the equity premium and, specifically, that $\mathbb{E}_t R_{m,t+1} - R_{f,t+1} = R_{f,t+1} \text{SVIX}_t^2$. Substituting this into equation (12), we have

$$\frac{\mathbb{E}_t R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \alpha_i + \text{SVIX}_t^2 + \frac{1}{2} (\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2). \tag{15}$$

We test (15) by running a panel regression of realized excess returns on individual stocks i onto stock fixed effects, risk-neutral variance SVIX_t^2 , and excess stock variance $\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}_t^2}$.

As noted above, the fixed effects in (14) and (15) should be zero on value-weighted average. We test this prediction in two ways. First, we consider the weaker prediction $\sum_i w_i \alpha_i = 0$ (where $w_i = \frac{1}{T_i} \sum_t w_{i,t}$ is the average value weight of stock i and T_i the number of time-series observations for stock i). We also test the stronger assumption that $\alpha_i = 0$ for all i , which will hold if risk-neutral residual variance is constant across stocks though not necessarily over time. In this form, we are imposing a tight relationship between a stock's risk-neutral variance and its risk-neutral beta: by (8), stocks with high variances must also have high risk-neutral betas. Making this assumption in (14), for example, we have⁷

$$\frac{\mathbb{E}_t R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \frac{1}{2} \left(\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}_t^2} \right). \tag{16}$$

Correspondingly, if we assume that the fixed effects are constant across i in (15), we end up with a formula for the expected return on a stock that has no free parameters:

$$\frac{\mathbb{E}_t R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \text{SVIX}_t^2 + \frac{1}{2} \left(\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}_t^2} \right). \tag{17}$$

In Section V, we exploit the fact that (16) and (17) require no parameter estimation—only observation of contemporaneous prices—to conduct out-of-sample analysis, and we show that the formulas outperform a range of plausible competitors.

Before we turn to the data, it is worth reiterating our two key assumptions. First, we assume that for stocks in our universe, risk-neutral betas $\beta_{i,t}^*$ are sufficiently close to one to justify our linearization (9). Second, we assume that the risk-neutral variances of residuals—the second term on the right-hand side of equation (12)—can be captured by a fixed-effects formulation.

We emphasize that these assumptions are not appropriate for all assets. Suppose, for example, that asset j is genuinely idiosyncratic—and hence has zero risk premium—but has extremely high, and perhaps wildly time-varying, variance $\text{SVIX}_{j,t}^2$. Then equation (15) cannot possibly hold for asset j . Our assumptions reflect a judgment that such cases are not typical within the universe of stocks that we study (namely, members of the S&P 100 or S&P

⁷ At first sight, equation (16) appears to lead to an inconsistency: if we “set $i = m$,” this seems to imply that $\text{SVIX}_t^2 = \overline{\text{SVIX}_t^2}$, which is not true (as we discuss in Section II below). The correct way to “set $i = m$ ” here is to replace $\text{SVIX}_{i,t}^2$ not with SVIX_t^2 but with its value-weighted sum, SVIX_t^2 . By contrast, it is legitimate to “set $i = m$ ” in linear factor models in which risk premia are expressed in terms of covariances of returns with factors.

500 indices).⁸ This is an empirically testable judgment, and we put it to the test below.

II. Three Measures of Risk-Neutral Variance

The risk-neutral variance terms that appear in our formulas can be calculated from option prices using the approach of Breeden and Litzenberger (1978). Our measure of market risk-neutral variance, $SVIX_t^2$, is determined by the prices of index options:

$$SVIX_t^2 = \frac{2}{R_{f,t+1} S_{m,t}^2} \left[\int_0^{F_{m,t}} \text{put}_{m,t}(K) dK + \int_{F_{m,t}}^{\infty} \text{call}_{m,t}(K) dK \right],$$

where $S_{m,t}$ and $F_{m,t}$ denote the spot and forward (to time $t + 1$) prices of the market, and $\text{put}_{m,t}(K)$ and $\text{call}_{m,t}(K)$ denote the time t prices of European puts and calls on the market that expire at time $t + 1$ with strike K . The length of the period from time t to time $t + 1$ varies according to the horizon of interest. We will therefore forecast 1-month returns using the prices of 1-month options, 3-month returns using the prices of 3-month options, and so on. Throughout the paper, we annualize returns and volatility indices by scaling by horizon length measured in years. The SVIX index (squared) thus represents the price of a portfolio of out-of-the-money puts and calls equally weighted by strike. This definition is closely related to that of the VIX index, the key difference being that VIX weights option prices in inverse-square proportion to their strike.

The corresponding index at the individual stock level is defined in terms of individual stock option prices:

$$SVIX_{i,t}^2 = \frac{2}{R_{f,t+1} S_{i,t}^2} \left[\int_0^{F_{i,t}} \text{put}_{i,t}(K) dK + \int_{F_{i,t}}^{\infty} \text{call}_{i,t}(K) dK \right],$$

where the subscripts i indicate that the reference asset is stock i rather than the market.

Finally, using $SVIX_{i,t}^2$ for all firms available at time t , we calculate the risk-neutral average stock variance index as $\overline{SVIX}_t^2 = \sum_i w_{i,t} SVIX_{i,t}^2$.

We pause to highlight two facts about these volatility indices. First, average stock volatility must exceed market volatility, that is, $\overline{SVIX}_t > SVIX_t$. Given the definitions above, this is an illustration of the slogan that a portfolio of options is more valuable than an option on a portfolio. More formally, it is a consequence of the fact that $\sum_i w_{i,t} \text{var}_t^* R_{i,t+1} > \text{var}_t^* \sum_i w_{i,t} R_{i,t+1}$ or,

⁸There is an analogy with an earlier debate on the testability of the arbitrage pricing theory (APT). Shanken (1982) shows, under the premise of the APT that asset returns are generated by a linear factor model, that it is possible to construct portfolios that violate the APT prediction that assets' expected returns are linear in the factor loadings. Dybvig and Ross (1985) endorse the mathematical content of Shanken's results but dispute their interpretation, arguing that the APT can be applied to certain types of assets (e.g., stocks) but not to arbitrary portfolios of assets.

equivalently, that $\mathbb{E}_t^* \sum_i w_{i,t} R_{i,t+1}^2 > \mathbb{E}_t^* [(\sum_i w_{i,t} R_{i,t+1})^2]$, which follows from Jensen's inequality.

Second, as a rule of thumb, risk-neutral variance is increasing in the time-to-maturity of the underlying options (equivalently, in the length of the period from t to $t + 1$). Formally, assume that the underlying asset does not pay dividends and use put-call parity to write

$$\text{SVIX}_{i,t}^2 = \text{var}_t^*(R_{i,t+1}/R_{f,t+1}) = \frac{2}{R_{f,t+1} S_t^2} \int_0^\infty \underbrace{\text{call}_{i,t}(K)}_{\uparrow \text{ in maturity}} dK - 1.$$

As is well known, if the underlying asset does not pay dividends—a tolerable approximation to reality for the stocks and horizons we consider—a European call and an American call have the same value and hence call prices are increasing in time-to-maturity. Assuming this is not offset by the countervailing effect of increased interest rates $R_{f,t+1}$ over longer horizons, $\text{SVIX}_{i,t}$ should be expected to be monotonic in horizon length. We have found nonmonotonicity to be a useful flag for detecting a small number of extreme outliers in our data, as we discuss further below.

In our empirical work, we start with daily data from OptionMetrics for equity index options on the S&P 100 and on the S&P 500, which provide us with time series of implied volatility surfaces from January 1996 to October 2014. We obtain daily equity index price and return data from CRSP and information on the index constituents from Compustat. We also obtain data on the firms' number of shares outstanding and their book equity to compute their market capitalizations and book-to-market ratios. Using the lists of index constituents, we search the OptionMetrics database for all firms that were included in the S&P 100 or S&P 500 during our sample period, and obtain volatility surface data for these individual firms, where available.

We face the issue that S&P 100 index options and individual stock options are American style rather than European style. The distinction is likely to be relatively minor at the horizons we consider, as the options whose prices we require are out of the money. In any case, the volatility surfaces reported by OptionMetrics address this issue via binomial tree calculations that aim to account for early exercise premia. We take the resulting volatility surfaces as our measures of European implied volatility, following Carr and Wu (2009), among others.

We compute the three measures of risk-neutral variance given in equation (13) for horizons (i.e., option maturities) of 1, 3, 6, 12, and 24 months. We then omit a small number of extreme outliers in our data that violate the monotonicity property of $\text{SVIX}_{i,t}$ across horizons described above.⁹ As summarized in Panel A of Table I, we end up with more than two million firm-day observations for each of the five horizons, covering a total of 869 firms over our sample

⁹ In the daily data, we end up with 2,106,711 firm-day observations after removing 9,648 observations based on nonmonotonicity. In our monthly data for S&P 500 firms, we end up with 102,198 firm-month observations after removing 401 observations based on nonmonotonicity.

Table I
Sample Data

This table summarizes the data used in the empirical analysis. We search the OptionMetrics database for all firms that have been included in the S&P 100 or S&P 500 during the sample period from January 1996 to October 2014 and obtain all available volatility surface data. We use these data to compute firms' risk-neutral variances ($SVIX_{t,t}^2$) for horizons of 1, 3, 6, 12, and 24 months. Panel A summarizes the number of total observations, the number of unique days and unique firms in our sample, as well as the average number of firms for which options data are available per day. For some econometric analysis, we also compile data subsets at a monthly frequency for firms included in the S&P 100 (summarized in Panel B) and the S&P 500 (Panel C).

Panel A: Daily data					
Horizon	30 days	91 days	182 days	365 days	730 days
Observations	2,106,711	2,106,711	2,106,711	2,106,711	2,106,711
Sample days	4,674	4,674	4,674	4,674	4,674
Sample firms	869	869	869	869	869
Average firms/day	451	451	451	451	451
Panel B: Monthly data for S&P 100 firms					
Observations	21,205	20,820	20,247	19,100	16,896
Sample months	224	222	219	213	201
Sample firms	177	176	176	171	167
Average firms/month	95	94	92	90	84
Panel C: Monthly data for S&P 500 firms					
Observations	102,198	100,252	97,340	91,585	80,631
Sample months	224	222	219	213	201
Sample firms	877	869	863	832	770
Average firms/month	456	452	444	430	401

period from January 1996 to October 2014. Across horizons, we have data on 451 firms on average per day, meaning that we cover slightly more than 90% of the firms included in the S&P 500 index. From the daily data, we also compile data subsets at a monthly frequency for firms included in the S&P 100 (Panel B) and the S&P 500 (Panel C).

Figure 2 plots the time series of risk-neutral market variance ($SVIX_t^2$) and average risk-neutral stock variance (\overline{SVIX}_t^2) for the S&P 500; the corresponding time series for the S&P 100 are shown in Figure IA.1 in the Internet Appendix. The dynamics of $SVIX_t^2$ and \overline{SVIX}_t^2 are similar for both indices and across horizons. All of the time series spike dramatically during the financial crisis of 2008. While the average levels of the (annualized) SVIX measures are similar across horizons, their volatility is higher at short than at long horizons. Similarly, the peaks in $SVIX_t^2$ and \overline{SVIX}_t^2 during the crisis and other periods of heightened volatility are most pronounced in short-maturity options.¹⁰

¹⁰ In Appendix A, we show that the ratio of market variance to average stock variance, $SVIX_t^2/\overline{SVIX}_t^2$, can be interpreted as a measure of average risk-neutral correlation between stocks.

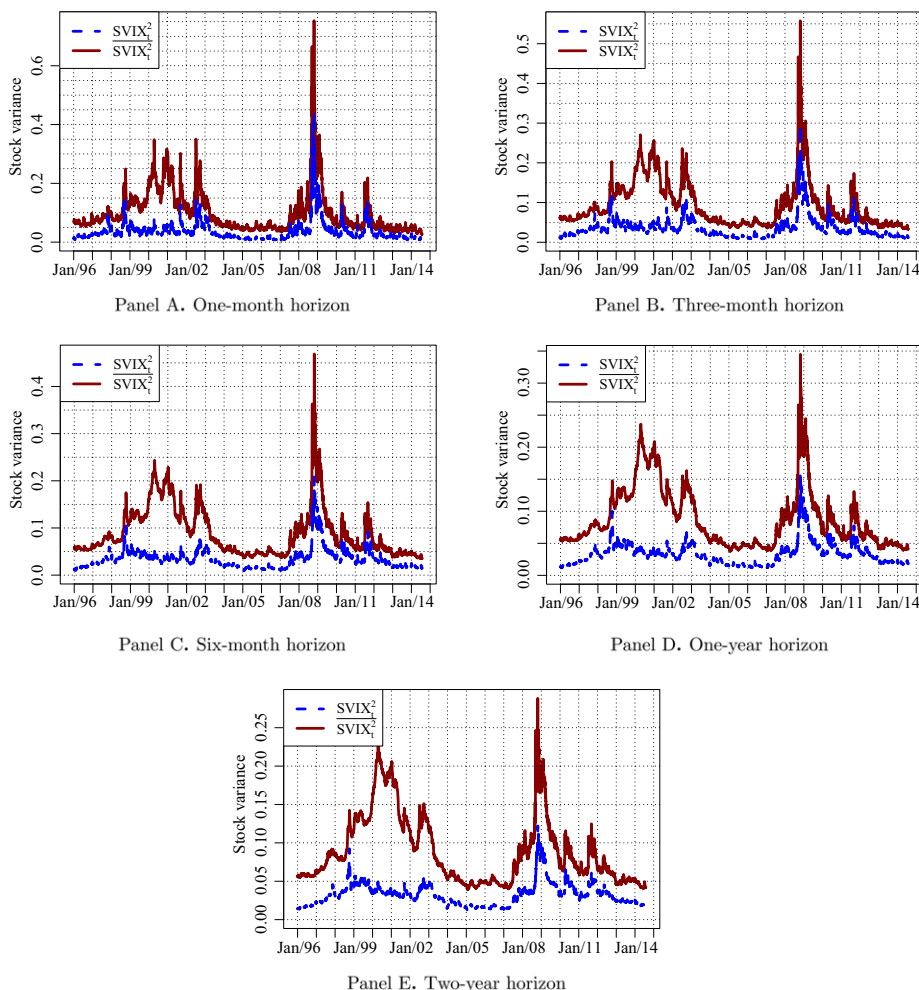


Figure 2. Option-implied equity variance of S&P 500 firms. This figure plots the time series of the risk-neutral variance of the market ($SVIX_t^2$) and of stocks' average risk-neutral variance ($SVIX_t^2$). We compute $SVIX_t^2$ from equity index options on the S&P 500. $SVIX_t^2$ is the value-weighted sum of S&P 500 stocks' risk-neutral variance computed from individual firm equity options. Panels A through D present the variance series implied by equity options with maturities of 1, 3, 6, 12, and 24 months. The data are daily from January 1996 to October 2014. (Color figure can be viewed at wileyonlinelibrary.com)

Figures 3 and 4 show the relationships between risk-neutral stock variances and various firm characteristics, on average and in the time series,

Figure IA.2 in the Internet Appendix plots the time series of $SVIX_t^2 / \overline{SVIX_t^2}$ at 1-month and 1-year horizons for the S&P 100 and S&P 500. Average stock variance was unusually high relative to market variance over the period from 2000 to 2002, indicating that the correlation between stocks was unusually low at that time.




Figure 3. Beta, size, value, momentum, and option-implied equity variance. This figure reports (equally weighted) averages of risk-neutral stock variance ($SVIX_{i,t}^2$, computed from individual firm equity options) of S&P 500 stocks, conditional on firm beta, size, book-to-market, and momentum. At each date t , we assign stocks to decile portfolios based on their characteristics and report the time-series averages of $SVIX_{i,t}^2$ across deciles at a horizon of 1 year (Panels A to D). (Color figure can be viewed at wileyonlinelibrary.com)

respectively. To construct the figures, we sort S&P 500 stocks into portfolios based on their CAPM beta, size, book-to-market ratio, or momentum, and compute the (equally weighted) average $SVIX_{i,t}^2$ for each portfolio at the 12-month horizon.¹¹ We find that $SVIX_{i,t}^2$ is positively related to CAPM beta and inversely related to firm size, both on average and throughout our sample period. In contrast, there is a U-shaped relationship between $SVIX_{i,t}^2$ and book-to-market that reflects an interesting time-series relationship between the two. Growth and value stocks had similar levels of volatility during periods of low index volatility, but value stocks were more volatile than growth stocks during the recent financial crisis and less volatile from 2000 to 2002. We also find a

¹¹ We measure momentum by the return over the past 12 months, skipping the most recent month's return (see, e.g., Jegadeesh and Titman, 1993). Our estimation of conditional CAPM betas based on past returns follows Frazzini and Pedersen (2014): we estimate volatilities by 1-year rolling standard deviations of daily returns and correlations from 5-year rolling windows of overlapping 3-day returns.

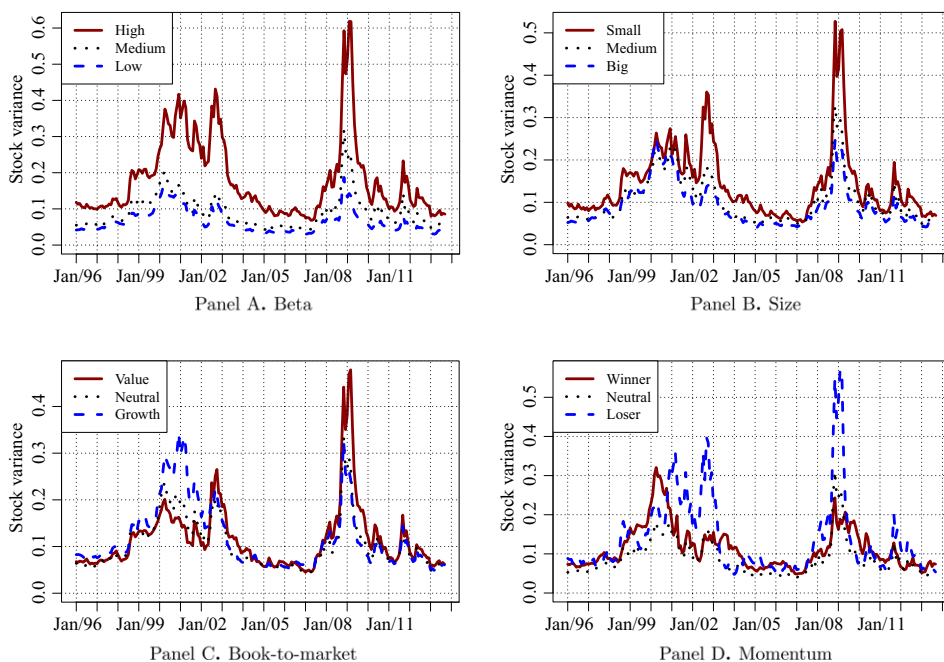


Figure 4. Beta, size, value, momentum, and option-implied equity variance. This figure plots the time series of risk-neutral stock variance ($SVIX_{i,t}^2$) of S&P 500 stocks, conditional on firm beta, size, book-to-market, and momentum. The horizon is 1 year. At each date t , we classify firms as small, medium, or big when their market capitalization is in the bottom, middle, or top tertile of the time t distribution across all firms in our sample, and compute the (equally weighted) average $SVIX_{i,t}^2$. Similarly, we classify firms by their other characteristics at time t . (Color figure can be viewed at wileyonlinelibrary.com)

nonmonotonic relationship between momentum and $SVIX_{i,t}^2$. Interestingly, loser stocks exhibited particularly high $SVIX_{i,t}^2$ from late 2008 until the momentum crash in early 2009.¹²

III. Testing the Model

In this section, we use $SVIX_t^2$, $SVIX_{i,t}^2$, and \overline{SVIX}_t^2 to test the predictions of our model using full-sample information. But before turning to formal tests, we conduct a preliminary exploratory exercise. Specifically, we ask whether, on time-series average, stocks' excess-of-market returns line up with their excess stock variances in the manner predicted by equation (16). To do so, we temporarily restrict attention to firms that were included in the S&P 500 throughout our sample period. For each such firm, we compute time-averaged excess-of-market

¹² We find similar results at the 1-month horizon. See Figures IA.3 and IA.4 in the Internet Appendix. Figure IA.5 plots (equally weighted) average $SVIX_i$ at the 12-month horizon for portfolios double-sorted on size and value.

returns and risk-neutral excess stock variance, $SVIX_i^2 - \overline{SVIX}^2$. Equation (16) implies that for each percentage point difference in $SVIX_i^2 - \overline{SVIX}^2$, we should see half a percentage point difference in excess returns.

The results of this exercise are shown in Figure 5, which is analogous to the security market line of the CAPM. The return horizon matches the maturity of the options used to compute the SVIX indices. We regress average excess-of-market returns on $0.5 \times (SVIX_i^2 - \overline{SVIX}^2)$. Our theory predicts zero intercept and a slope coefficient of one; we find intercepts close to zero and slope coefficients of 0.60, 0.79, 1.00, 1.10, and 1.01 at forecasting horizons of 1, 3, 6, 12, and 24 months, with R^2 ranging from 0.09 to 0.18. Using the same subset of firms, the figures also show decile portfolios sorted by $SVIX_{i,t}$ (indicated by diamonds) and 3×3 portfolios sorted by size and book-to-market (indicated by triangles).

We repeat this exercise for portfolios sorted on firms' risk-neutral variance $SVIX_{i,t}$, using all available firms (lifting the requirement of full-sample-period coverage). Figure 6 shows that average portfolio returns in excess of the market are broadly increasing in portfolios' average volatility relative to aggregate stock volatility, and that $SVIX_i^2 - \overline{SVIX}^2$ captures a sizeable fraction of the cross-sectional variation in returns.

To test the model formally, we start by estimating the pooled panel regression

$$\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha + \gamma \left(SVIX_{i,t}^2 - \overline{SVIX}^2 \right) + \epsilon_{i,t+1}. \quad (18)$$

Based on the formula (16), we would ideally hope to find that $\alpha = 0$ and $\gamma = 1/2$. At a given point in time t , our sample includes all firms that are time t constituents of the index. We compute $R_{m,t+1}$ as the return on the value-weighted portfolio of all index constituent firms included in our sample at time t .

We run the regression using monthly data for the S&P 100 and S&P 500 indices, at return horizons (and hence also option maturities) of 1, 3, 6, 12, and 24 months. Throughout the paper, we calculate standard errors and p -values using a block bootstrap procedure that accounts for time-series and cross-sectional dependencies in the data. Appendix B provides further details about the bootstrap procedure and presents Monte Carlo simulation evidence on the reliability of the procedure in finite samples.

The regression results are shown in Table II. The headline result is that when we conduct a Wald test of the joint hypothesis that $\alpha = 0$ and $\gamma = 0.5$, we do not reject our model at any horizon, with p -values ranging from 0.44 to 0.84 for S&P 100 firms (Panel A) and from 0.49 to 0.63 for S&P 500 firms (Panel B). By contrast, we can reject the hypothesis that $\gamma = 0$ with some confidence in most cases (with p -values of 0.079, 0.020, 0.015, and 0.007 for S&P 100 firms at 3-, 6-, 12-, and 24-month horizons, and p -values of 0.072, 0.068, and 0.077 for S&P 500 firms at 6-, 12-, and 24-month horizons).

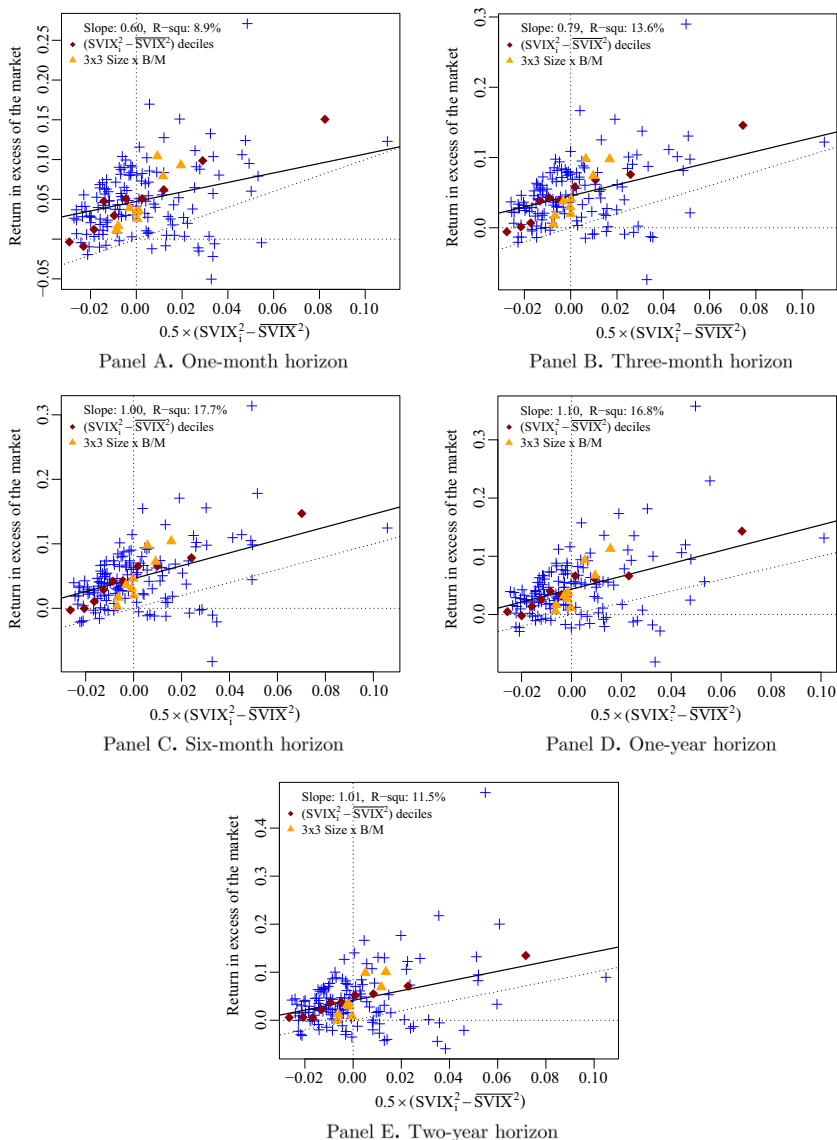


Figure 5. Average equity returns in excess of the market. This figure presents results on the relation between a firm’s equity returns in excess of the market and its risk-neutral variance measured relative to average risk-neutral stock variance. For firms that were constituents of the S&P 500 index throughout our sample period, we compute time-series averages of their returns in excess of the market and their stock volatility relative to stocks’ average volatility ($SVIX_t^2 - \overline{SVIX}^2$). We multiply the stock variance estimate by 0.5 and plot the pairwise combinations (blue crosses) for horizons of 1, 3, 6, 12, and 24 months (Panels A to E). The black line represents the regression fit to the individual firm observations with slope coefficient and R^2 reported in the plot legend. Our theory implies that the slope coefficient of this regression should be one and that the intercept should be zero. The red diamonds represent decile portfolios of firms sorted by $SVIX_{i,t}^2$. Similarly, the triangles in orange represent portfolios of stocks formed according to firms’ size and book-to-market. (Color figure can be viewed at wileyonlinelibrary.com)




Figure 6. Portfolios sorted by stock risk-neutral variance. This figure reports results on the relationship between equity portfolio returns in excess of the market and risk-neutral stock variance measured relative to stocks’ average risk-neutral variance. The horizon is 1 year. At the end of each month, we group all available firms into 10, 25, 50, or 100 portfolios (Panels A to D) based on their individual variance relative to average variance, $SVIX_{i,t}^2 - \overline{SVIX}_t^2$. For each portfolio, we compute the time-series average return in excess of the market and plot the pairwise combinations with the corresponding stock variance estimate multiplied by 0.5. Our theory implies that the slope coefficient of this regression should be one. The black line represents the regression fit to the portfolio observations with slope coefficient and R^2 reported in the plot legend. The sample period is January 1996 to October 2014. (Color figure can be viewed at wileyonlinelibrary.com)

We test prediction (14) by running a panel regression with firm fixed effects,

$$\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha_i + \gamma \left(SVIX_{i,t}^2 - \overline{SVIX}_t^2 \right) + \epsilon_{i,t+1}, \tag{19}$$

and testing the hypothesis that $\gamma = 1/2$ and $\sum_i w_i \alpha_i = 0$.

The results are in Table III. Now γ is significantly different from zero even at the shorter horizons, and in most cases is not significantly different from 0.5. We also find, however, that the value-weighted sum of firm fixed effects is

Table II

Expected Returns in Excess of the Market: Pooled Panel Regressions

This table presents results from regressing stock returns in excess of the market onto stock-level risk-neutral variance measured relative to stocks' average risk-neutral variance ($SVIX_{i,t}^2 - \overline{SVIX}_t^2$) for S&P 100 firms (Panel A) and S&P 500 firms (Panel B). The data are monthly from January 1996 to October 2014. The column labels indicate the return horizons ranging from 1 month to 2 years. The return horizons match the maturities of the options used to compute $SVIX_{i,t}^2$ and \overline{SVIX}_t^2 .

We report estimates of the pooled panel regression specified in equation (18), $\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} =$

$\alpha + \gamma(SVIX_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}$. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. In each panel, we report the regressions' adjusted R^2 and p -values of Wald tests that test whether the regression coefficients take the values predicted by our theory (joint test of zero intercept and $\gamma = 0.5$), whether $\gamma = 0.5$, and whether $\gamma = 0$. The rows labelled "theory adj- R^2 (%)" report the adjusted R^2 obtained when the coefficients are fixed at the values predicted by our theory.

Horizon	30 days	91 days	182 days	365 days	730 days
Panel A: S&P 100 firms					
α	0.008 (0.015)	0.008 (0.014)	0.005 (0.015)	0.007 (0.016)	0.010 (0.016)
γ	0.541 (0.345)	0.551 (0.313)	0.761 (0.328)	0.819 (0.337)	0.723 (0.270)
Adjusted R^2 (%)	0.473	1.185	3.527	6.070	6.665
$H_0 : \alpha = 0, \gamma = 0.5$	0.841	0.832	0.609	0.437	0.439
$H_0 : \gamma = 0.5$	0.906	0.871	0.427	0.344	0.409
$H_0 : \gamma = 0$	0.118	0.079	0.020	0.015	0.007
Theory adj- R^2 (%)	0.463	1.151	3.054	5.005	5.712
Panel B: S&P 500 firms					
α	0.016 (0.015)	0.016 (0.015)	0.013 (0.016)	0.014 (0.019)	0.019 (0.019)
γ	0.301 (0.285)	0.414 (0.273)	0.551 (0.306)	0.553 (0.302)	0.354 (0.200)
Adjusted R^2 (%)	0.135	0.617	1.755	2.892	1.901
$H_0 : \alpha = 0, \gamma = 0.5$	0.489	0.560	0.630	0.600	0.596
$H_0 : \gamma = 0.5$	0.486	0.752	0.869	0.862	0.467
$H_0 : \gamma = 0$	0.291	0.129	0.072	0.068	0.077
Theory adj- R^2 (%)	0.068	0.547	1.648	2.667	1.235

statistically different from zero, though we note that the estimates are fairly small in economic terms (and, consistent with the pooled panel results, we will see below that the model performs well when we drop firm fixed effects entirely, as we do in our out-of-sample analysis).¹³

¹³ Moreover, the fixed effects are not statistically significant if we use portfolios sorted on $SVIX_{i,t}^2$ as test assets. See Tables IA.I, IA.II, IA.III, and IA.IV in the Internet Appendix.

Table III
Expected Returns in Excess of the Market: Panel Regressions with Fixed Effects

This table presents results from regressing stock returns in excess of the market onto stock-level risk-neutral variance measured relative to stocks' average risk-neutral variance ($SVIX_{i,t}^2 - \overline{SVIX}_t^2$) for S&P 100 firms (Panel A) and for S&P 500 firms (Panel B). The data are monthly from January 1996 to October 2014. The column labels indicate the return horizons ranging from 1 month to 2 years. The return horizons match the maturities of the options used to compute $SVIX_{i,t}^2$ and \overline{SVIX}_t^2 . We report estimates of the panel regression with firm fixed effects specified in equation (19), $\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha_i + \gamma(SVIX_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}$, where $\sum_i w_i \alpha_i$ reports the time-series average of the value-weighted sum of firm fixed effects. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. In each panel, we report the regressions' adjusted R^2 and p -values of Wald tests that test whether the regression coefficients take the values predicted by our theory (joint test of zero intercept and $\gamma = 0.5$), whether $\gamma = 0.5$, and whether $\gamma = 0$.

Horizon	30 days	91 days	182 days	365 days	730 days
Panel A: S&P 100 firms					
$\sum_i w_i \alpha_i$	0.026 (0.010)	0.024 (0.009)	0.023 (0.009)	0.022 (0.009)	0.020 (0.009)
γ	0.780 (0.385)	0.833 (0.360)	1.120 (0.348)	1.156 (0.313)	1.018 (0.286)
Adjusted R^2 (%)	1.097	4.013	9.896	16.866	24.071
$H_0 : \sum_i w_i \alpha_i = 0, \gamma = 0.5$	0.026	0.012	0.006	0.002	0.013
$H_0 : \gamma = 0.5$	0.468	0.355	0.074	0.036	0.070
$H_0 : \gamma = 0$	0.043	0.021	0.001	0.000	0.000
Panel B: S&P 500 firms					
$\sum_i w_i \alpha_i$	0.036 (0.008)	0.034 (0.007)	0.033 (0.008)	0.033 (0.008)	0.033 (0.008)
γ	0.560 (0.313)	0.730 (0.313)	0.949 (0.319)	0.917 (0.291)	0.637 (0.199)
Adjusted R^2 (%)	0.398	3.015	7.320	12.637	17.479
$H_0 : \sum_i w_i \alpha_i = 0, \gamma = 0.5$	0.000	0.000	0.000	0.000	0.000
$H_0 : \gamma = 0.5$	0.848	0.461	0.160	0.152	0.491
$H_0 : \gamma = 0$	0.073	0.019	0.003	0.002	0.001

Turning to excess returns (as opposed to excess-of-market returns), we test the prediction of equation (17) by running the regression

$$\frac{R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \alpha + \beta SVIX_t^2 + \gamma (SVIX_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}, \quad (20)$$

and the prediction of equation (15) by running a regression with stock fixed effects,

$$\frac{R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \alpha_i + \beta SVIX_t^2 + \gamma (SVIX_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}. \quad (21)$$

Our model predicts that $\alpha = 0$, $\beta = 1$, and $\gamma = 1/2$ in equation (20), and that $\beta = 1$, $\gamma = 1/2$, and $\sum_i w_i \alpha_i = 0$ in equation (21).

The pooled panel regression results are shown in Table IV. For S&P 100 firms (Panel A), the headline result is again that we do not reject our model at any horizon: p -values of the joint hypothesis test that $\alpha = 0$, $\beta = 1$, and $\gamma = 0.5$ range from 0.55 to 0.69. By contrast, we can reject the joint hypothesis that $\beta = 0$ and $\gamma = 0$ with moderate confidence for 6-, 12-, and 24-month returns (with p -values of 0.064, 0.045, and 0.012, respectively). Notice that as the estimated coefficient γ exploits cross-sectional information, it is estimated more precisely than is β . Our results are therefore consistently stronger, in a statistical sense, than those of Martin (2017).

The corresponding results for S&P 500 firms are reported in Panel B. We do not reject the joint hypothesis that $\alpha = 0$, $\beta = 1$, and $\gamma = 0.5$ at horizons of 1, 3, 6, and 12 months (with p -values between 0.169 and 0.267). We do, however, reject the model at the 24-month horizon: the estimated β is even higher than the theory predicts. We can cautiously reject the joint null that $\beta = 0$ and $\gamma = 0$ at horizons of 6, 12, and 24 months (with p -values of 0.071, 0.092, and 0.036).

The coefficient estimates remain fairly stable, and we draw similar conclusions, when we allow for firm fixed effects in Table V. For S&P 100 firms (Panel A), a Wald test of the joint null hypothesis that $\sum_i w_i \alpha_i = 0$, $\beta = 1$, and $\gamma = 0.5$ does not reject the model (with p -values between 0.11 and 0.36), and we can strongly reject the joint null that $\beta = \gamma = 0$ for horizons of 6, 12, and 24 months (with p -values below 0.01). The β estimates are little changed compared to the pooled panel regressions, while the γ estimates are somewhat higher. The statistical results are more clear-cut for S&P 500 firms when we include firm fixed effects (Panel B). We do not reject the joint null hypothesis implied by our model at horizons up to and including 12 months, and can strongly reject the null that $\beta = \gamma = 0$ at horizons of 6, 12, and 24 months (with p -values of 0.019, 0.008, and 0.002).

We also run these regressions on subsamples of the data. Figure 7 plots the estimated coefficients β and γ using successive yearly and three-yearly subsamples, and shows that our results are not driven by any one subperiod. The figure also helps emphasize the point that the cross-sectional coefficient γ , which exploits the information in the entire cross section of stocks, is estimated more precisely than the “market” coefficient β , which relies on a single time series.

IV. Risk Premia and Stock Characteristics

The results of the previous section show that the model performs well in forecasting stock returns. Nonetheless, we would like to know whether there is return-relevant information in other firm characteristics—notably, CAPM beta, (log) size, book-to-market, and past returns—that is not captured by our predictor variables (see, e.g., Fama and French (1993), Carhart (1997), Lewellen (2015)).

Table IV
Expected Excess Returns: Pooled Panel Regressions

This table presents results from regressing excess returns of S&P 100 firms (Panel A) and S&P 500 firms (Panel B) onto the risk-neutral variance of the market ($SVIX_t^2$) and stock-level risk-neutral variance measured relative to stocks' average risk-neutral variance ($SVIX_{i,t}^2 - \overline{SVIX_t^2}$). The data are monthly from January 1996 to October 2014. The column labels indicate the return horizons ranging from 1 month to 2 years. The return horizons match the maturities of the options used to compute $SVIX_t^2$, $SVIX_{i,t}^2$, and $\overline{SVIX_t^2}$. We report estimates of the pooled panel regression specified in equation (20), $\frac{R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \alpha + \beta SVIX_t^2 + \gamma (SVIX_{i,t}^2 - \overline{SVIX_t^2}) + \varepsilon_{i,t+1}$. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. In each panel, we report the regressions' adjusted R^2 and p -values of Wald tests that test whether the regression coefficients take the values predicted by our theory (zero intercept, $\beta = 1$, and $\gamma = 0.5$), whether $\beta = 0$ and $\gamma = 0$, whether $\gamma = 0.5$, and whether $\gamma = 0$. The rows labelled "theory adj- R^2 (%)" report the adjusted R^2 obtained when the coefficients are fixed at the values predicted by our theory.

Horizon	30 days	91 days	182 days	365 days	730 days
Panel A: S&P 100 firms					
α	0.073 (0.064)	0.035 (0.074)	-0.009 (0.054)	0.001 (0.067)	-0.006 (0.068)
β	-0.001 (2.032)	1.070 (2.263)	2.244 (1.465)	1.956 (1.404)	1.990 (1.517)
γ	0.469 (0.346)	0.489 (0.332)	0.729 (0.340)	0.834 (0.343)	0.736 (0.267)
Adjusted R^2 (%)	0.274	0.942	3.809	6.387	7.396
$H_0 : \alpha = 0, \beta = 1, \gamma = 0.5$	0.550	0.687	0.660	0.566	0.608
$H_0 : \beta = \gamma = 0$	0.356	0.335	0.064	0.045	0.012
$H_0 : \gamma = 0.5$	0.929	0.974	0.500	0.330	0.376
$H_0 : \gamma = 0$	0.175	0.140	0.032	0.015	0.006
Theory adj- R^2 (%)	0.099	0.625	2.509	3.896	4.830
Panel B: S&P 500 firms					
α	0.057 (0.074)	0.019 (0.079)	-0.038 (0.059)	-0.021 (0.071)	-0.054 (0.076)
β	0.743 (2.311)	1.882 (2.410)	3.483 (1.569)	3.032 (1.608)	3.933 (1.792)
γ	0.214 (0.296)	0.305 (0.287)	0.463 (0.320)	0.512 (0.318)	0.324 (0.200)
Adjusted R^2 (%)	0.096	0.767	3.218	4.423	5.989
$H_0 : \alpha = 0, \beta = 1, \gamma = 0.5$	0.267	0.242	0.169	0.184	0.015
$H_0 : \beta = \gamma = 0$	0.770	0.553	0.071	0.092	0.036
$H_0 : \gamma = 0.5$	0.333	0.497	0.908	0.971	0.377
$H_0 : \gamma = 0$	0.470	0.287	0.148	0.108	0.105
Theory adj- R^2 (%)	-0.107	0.227	1.491	1.979	1.660

Table V

Expected Excess Returns: Panel Regressions with Fixed Effects

This table presents results from regressing excess returns of S&P 100 firms (Panel A) and S&P 500 firms (Panel B) onto the risk-neutral variance of the market ($SVIX_t^2$) and stock-level risk-neutral variance measured relative to stocks' average risk-neutral variance ($SVIX_{i,t}^2 - \overline{SVIX}_t^2$). The data are monthly from January 1996 to October 2014. The column labels indicate the return horizons ranging from 1 month to 2 years. The return horizons match the maturities of the options used to compute $SVIX_t^2$, $SVIX_{i,t}^2$, and \overline{SVIX}_t^2 . We report estimates of the panel regression with firm fixed effects specified in equation (21), $\frac{R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \alpha_i + \beta SVIX_t^2 + \gamma(SVIX_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}$, where $\sum_i w_i \alpha_i$ reports the time-series average of the value-weighted sum of firm fixed effects. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. In each panel, we report the regressions' adjusted R^2 and p -values of Wald tests that test whether the regression coefficients take the values predicted by our theory (zero intercept, $\beta = 1$, and $\gamma = 0.5$), whether $\beta = 0$ and $\gamma = 0$, whether $\gamma = 0.5$, and whether $\gamma = 0$.

Horizon	30 days	91 days	182 days	365 days	730 days
Panel A: S&P 100 firms					
$\sum w_i \alpha_i$	0.089 (0.062)	0.051 (0.071)	0.010 (0.051)	0.018 (0.064)	0.003 (0.066)
β	-0.085 (2.041)	0.947 (2.277)	2.091 (1.423)	1.793 (1.325)	1.876 (1.391)
γ	0.734 (0.392)	0.801 (0.387)	1.126 (0.370)	1.225 (0.314)	1.083 (0.273)
Adjusted R^2 (%)	1.211	4.771	11.861	20.003	27.455
$H_0 : \sum_i w_i \alpha_i = 0, \beta = 1, \gamma = 0.5$	0.233	0.363	0.274	0.111	0.184
$H_0 : \beta = \gamma = 0$	0.128	0.103	0.008	0.000	0.000
$H_0 : \gamma = 0.5$	0.551	0.436	0.091	0.021	0.033
$H_0 : \gamma = 0$	0.061	0.038	0.002	0.000	0.000
Panel B: S&P 500 firms					
$\sum w_i \alpha_i$	0.080 (0.072)	0.042 (0.075)	-0.008 (0.055)	0.012 (0.070)	-0.026 (0.079)
β	0.603 (2.298)	1.694 (2.392)	3.161 (1.475)	2.612 (1.493)	3.478 (1.681)
γ	0.491 (0.325)	0.634 (0.331)	0.892 (0.336)	0.938 (0.308)	0.665 (0.205)
Adjusted R^2 (%)	0.650	4.048	10.356	17.129	24.266
$H_0 : \sum_i w_i \alpha_i = 0, \beta = 1, \gamma = 0.5$	0.231	0.224	0.164	0.133	0.060
$H_0 : \beta = \gamma = 0$	0.265	0.119	0.019	0.008	0.002
$H_0 : \gamma = 0.5$	0.978	0.686	0.243	0.155	0.420
$H_0 : \gamma = 0$	0.131	0.056	0.008	0.002	0.001

As a preliminary check, Figure 8 shows that average realized excess returns line up fairly well with our cross-sectional excess return predictor, $0.5(SVIX_{i,t}^2 - \overline{SVIX}_t^2)$, for characteristic-sorted portfolios. The return predictor for a portfolio is calculated by averaging over its constituent stocks. Unless otherwise noted, we work with S&P 500 stocks and at an annual horizon throughout this section.

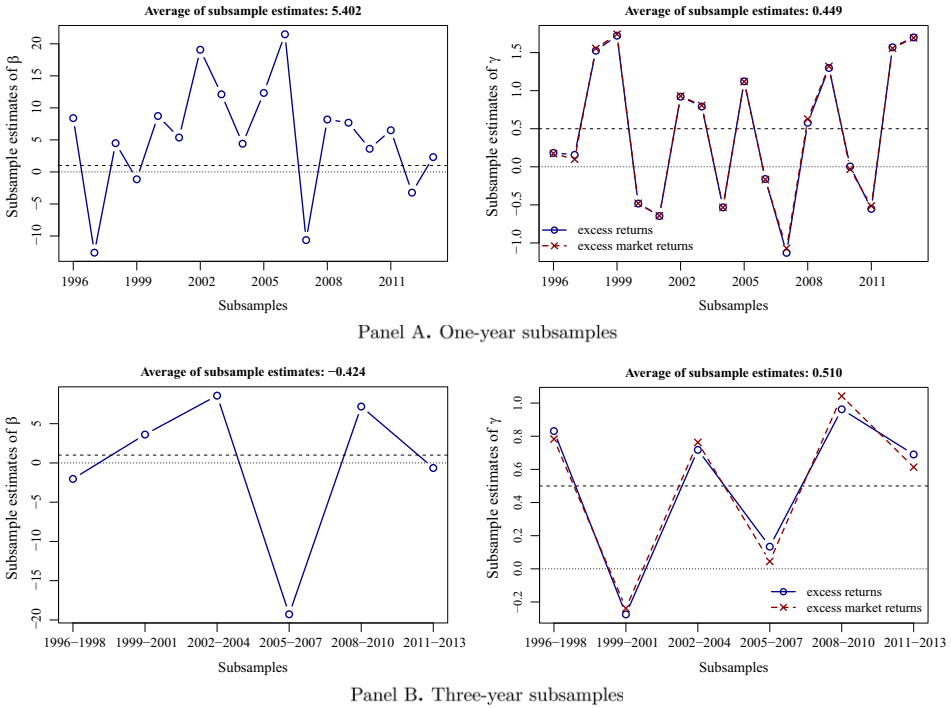


Figure 7. Regression estimates in subsamples. This figure summarizes results from regressing realized excess returns of S&P 500 firms onto the risk-neutral variance of the market, $SVIX_t^2$, and risk-neutral stock variance measured relative to average risk-neutral stock variance, $SVIX_{i,t}^2 - \overline{SVIX}_t^2$. The data are monthly from January 1996 to October 2014. We present results for yearly subsamples in Panel A and for 3-year subsamples in Panel B. The return horizon is 1 year, matching the maturity of the options used to compute $SVIX_t^2$, $SVIX_{i,t}^2$, and \overline{SVIX}_t^2 . We report estimates for the pooled panel regressions (18) and (20). The dashed line in each panel indicates the coefficient value predicted by our theory, that is, $\beta = 1$ and $\gamma = 0.5$. (Color figure can be viewed at wileyonlinelibrary.com)

We test formally whether our framework is able to explain differences in risk premia associated with the various characteristics in two ways: we run regressions of individual stock excess returns onto our predictor variables and the characteristics, and we rerun the regressions of the previous section using portfolios double-sorted on characteristics and on $SVIX_{i,t}^2$ as test assets.

Consider, first, the regressions on characteristics and our predictors. Table VI reports the results for returns in excess of the market. The first column shows the estimated coefficients in a regression of realized excess-of-market returns onto characteristics. We do not find a statistically significant relationship between the characteristics and realized returns in excess of the market (consistent with the findings of Nagel (2005), who documents limited cross-sectional variation in returns on S&P 500 stocks sorted on book-to-market, for example), and we cannot reject the joint hypothesis that the coefficients on all

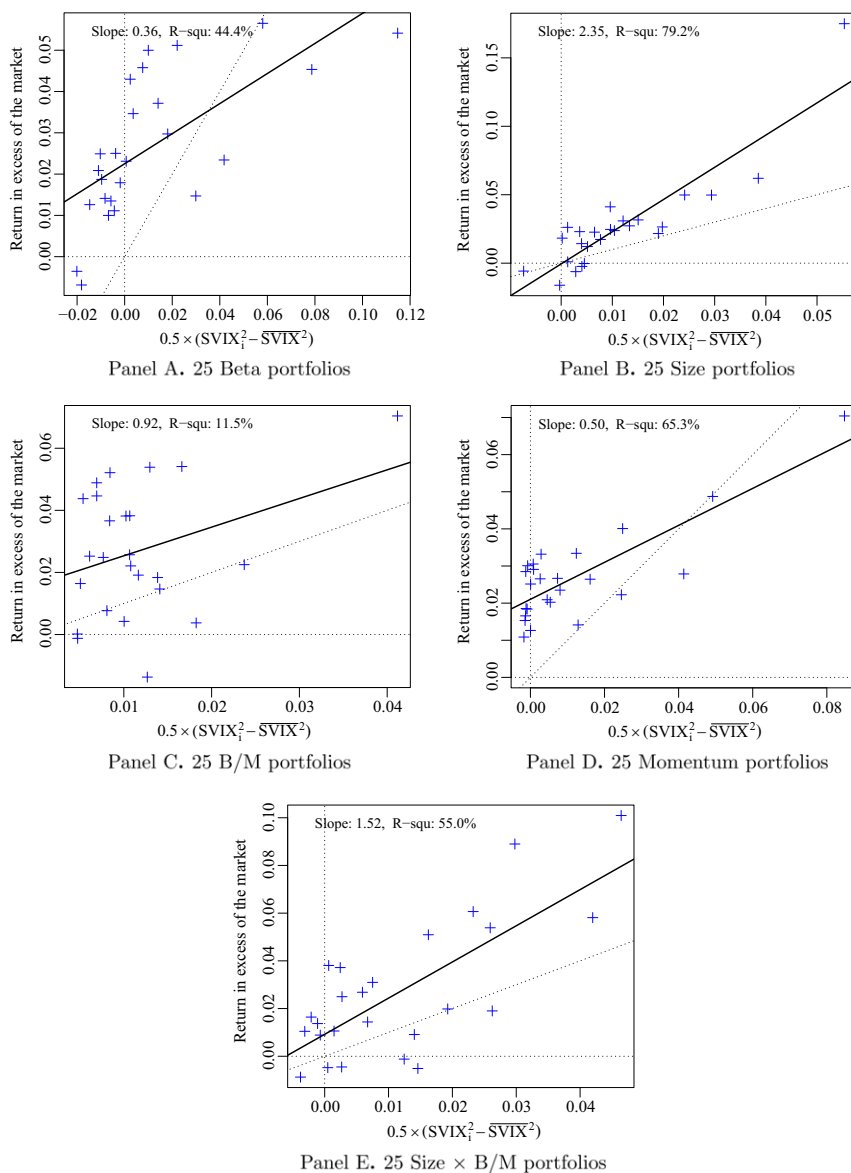


Figure 8. Portfolios sorted by beta, size, book-to-market, and momentum. This figure reports results on the relationship between equity portfolio returns in excess of the market and risk-neutral stock variance measured relative to average firm-level risk-neutral variance. At the end of each month, we form 25 portfolios based on firms’ beta, size, book-to-market, or momentum (Panels A to D) and on a 5×5 conditional double-sort on size and book-to-market (Panel E). For each portfolio, we compute the time-series average return in excess of the market and plot the pairwise combinations with the corresponding stock variance estimate multiplied by 0.5. The black line represents the regression fit to the portfolio observations with slope coefficient and R^2 reported in the plot legend. Our theory implies that the slope coefficient of this regression should be one. The sample period is January 1996 to October 2014. (Color figure can be viewed at wileyonlinelibrary.com)

Table VI
The Relationship between Realized, Expected, and Unexpected
Excess-of-Market Returns and Characteristics

This table presents results from regressing realized, expected, or unexpected equity returns in excess of the market ($y_{i,t+1}$) onto CAPM beta, log size, book-to-market, past return, and risk-neutral stock variance measured relative to stocks' average risk-neutral variance, $SVIX_{i,t}^2 - \overline{SVIX}_t^2$: $y_{i,t+1} = a + b_1 \text{Beta}_{i,t} + b_2 \log(\text{Size}_{i,t}) + b_3 \text{B/M}_{i,t} + b_4 \text{Ret}_{i,t}^{(12,1)} + c(SVIX_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}$. The data are monthly and cover S&P 500 firms from January 1996 to October 2014. The first two columns present results for realized returns, the middle two columns for expected returns, and the last two columns for unexpected returns. In columns labelled "theory," we set the parameter values of our model forecast to the values implied by equation (16); in columns labelled "estimated," we use parameter estimates of a pooled panel regression (i.e., we use the estimates obtained from the regression specified in equation (18) and reported in Panel B of Table II). The return horizon is 1 year. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. The last four rows report the regression's adjusted R^2 and the p -values of Wald tests on joint parameter significance, testing (i) whether all b_i estimates are zero, (ii) whether all b_i estimates are zero and $c = 0.5$, and (iii) whether all nonconstant coefficients are jointly zero.

	Realized returns		Expected returns		Unexpected returns	
			Estimated	Theory	Estimated	Theory
const	0.429 (0.371)	0.277 (0.377)	0.131 (0.073)	0.107 (0.027)	0.298 (0.365)	0.321 (0.359)
Beta _{<i>i,t</i>}	0.016 (0.075)	-0.131 (0.062)	0.113 (0.066)	0.105 (0.016)	-0.097 (0.046)	-0.088 (0.078)
log(Size _{<i>i,t</i>})	-0.018 (0.014)	-0.006 (0.015)	-0.009 (0.006)	-0.009 (0.002)	-0.009 (0.015)	-0.010 (0.013)
B/M _{<i>i,t</i>}	0.032 (0.025)	0.031 (0.027)	0.001 (0.006)	0.001 (0.005)	0.032 (0.026)	0.032 (0.026)
Ret _{<i>i,t</i>} ^(12,1)	-0.051 (0.041)	-0.029 (0.041)	-0.017 (0.018)	-0.015 (0.010)	-0.034 (0.039)	-0.035 (0.040)
SVIX _{<i>i,t</i>} ² - \overline{SVIX}_t^2		0.705 (0.308)				
Adjusted R^2 (%)	1.031	3.969	37.766	37.766	1.051	0.974
$H_0 : b_i = 0$	0.347	0.153	0.435	0.000	0.157	0.619
$H_0 : b_i = 0, c = 0.5$		0.234				
$H_0 : b_i = 0, c = 0$		0.018				

characteristics are zero. In the second column, we add our predictor $SVIX_{i,t}^2 - \overline{SVIX}_t^2$. We find that its estimate is statistically significant individually, and we do not reject the joint hypothesis that it enters with a coefficient of 0.5 while the coefficients on all characteristics are zero; adjusted R^2 increases from 1.0% to 4.0% when we add our predictor variable.

Table VII reports the corresponding results for excess returns. In the absence of our predictor variables, we find that size and book-to-market characteristics are individually statistically significant, and we can reject the joint hypothesis that the coefficients on all characteristics are zero. But once we add $SVIX_t^2$ and

Table VII
The Relationship between Realized, Expected, and Unexpected Returns and Characteristics

This table presents results from regressing realized, expected, and unexpected equity excess returns ($y_{i,t+1}$) onto CAPM beta, log size, book-to-market, past return, risk-neutral market variance ($SVIX_t$), and risk-neutral stock variance measured relative to stocks' average risk-neutral variance ($SVIX_{i,t}^2 - \overline{SVIX}_t^2$): $y_{i,t+1} = a + b_1\text{Beta}_{i,t} + b_2\log(\text{Size}_{i,t}) + b_3\text{B}/\text{M}_{i,t} + b_4\text{Ret}_{i,t}^{(12,1)} + c_0\text{SVIX}_t^2 + c_1(\text{SVIX}_{i,t}^2 - \overline{SVIX}_t^2) + \epsilon_{i,t+1}$. The data are monthly and cover S&P 500 firms from January 1996 to October 2014. The first two columns present results for realized returns, the middle two columns for expected returns, and the last two columns for unexpected returns. In columns labelled "theory," we set the parameter values of our model forecast to the values implied by theory (i.e., we use equation (17)); in columns labelled "estimated," we use parameter estimates of a pooled panel regression (i.e., we use the estimates obtained from the regression specified in equation (20) and reported in Panel B of Table IV). The return horizon is 1 year. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. The last four rows report adjusted R^2 and the p -values of Wald tests of joint parameter significance, testing (i) whether all b_i estimates are zero, (ii) whether all b_i estimates are zero, $c_0 = 1$, and $c_1 = 0.5$, and (iii) whether all nonconstant coefficients are jointly zero.

	Realized returns		Expected returns		Unexpected returns	
			Estimated	Theory	Estimated	Theory
const	0.721 (0.341)	0.452 (0.320)	0.259 (0.133)	0.164 (0.035)	0.462 (0.332)	0.557 (0.331)
Beta _{<i>i,t</i>}	0.038 (0.068)	-0.048 (0.068)	0.082 (0.064)	0.097 (0.018)	-0.044 (0.046)	-0.059 (0.072)
log(Size _{<i>i,t</i>})	-0.030 (0.014)	-0.019 (0.013)	-0.010 (0.007)	-0.009 (0.002)	-0.019 (0.013)	-0.021 (0.013)
B/M _{<i>i,t</i>}	0.071 (0.034)	0.068 (0.038)	0.003 (0.010)	0.001 (0.006)	0.068 (0.038)	0.069 (0.037)
Ret _{<i>i,t</i>} ^(12,1)	-0.049 (0.063)	-0.005 (0.054)	-0.046 (0.042)	-0.026 (0.015)	-0.003 (0.050)	-0.023 (0.058)
SVIX _{<i>t</i>} ²		2.792 (1.472)				
SVIX _{<i>i,t</i>} ² - $\overline{\text{SVIX}}_t^2$		0.511 (0.357)				
Adjusted R^2 (%)	1.924	5.265	17.277	30.482	0.973	1.197
$H_0 : b_i = 0$	0.003	0.201	0.702	0.000	0.187	0.092
$H_0 : b_i = 0, c_0 = 1, c_1 = 0.5$		0.143				
$H_0 : b_i = 0, c_0 = 0, c_1 = 0$		0.001				

$SVIX_{i,t}^2 - \overline{SVIX}_t^2$, we do not reject the joint hypothesis that the coefficients on the characteristics are all zero while those on the volatility measures are equal to their theoretical values of 1 and 0.5. Moreover, adjusted R^2 increases from 1.9% to 5.3% when our predictor variables are added.

The next columns of Tables VI and VII address the relationships between expected excess returns and characteristics, with expected excess returns calculated in two ways, namely, using the coefficients estimated in regressions (18) or (20), and using the theory-implied coefficients given in equations (17)

or (16). (We do so for interest: our theory makes no predictions about these regressions.) The characteristics capture a sizeable fraction of the variation in theory-implied expected returns in excess of the market ($R^2 = 37.8\%$) and theory-implied expected excess returns ($R^2 = 30.5\%$). In both cases, there is a significantly positive relationship between expected returns and beta and a significantly negative relationship between expected returns and size, but the other characteristics do not exhibit a statistically significant relationship to expected returns. When we calculate expected returns using the estimated coefficients from (18) and (20) rather than the theoretical values, the point estimates of the regression coefficients for the characteristics are similar but are estimated less precisely, so are not significantly different from zero.

The last two columns of the tables show that there is little evidence of a systematic relationship between *unexpected* (i.e., realized minus expected) returns and characteristics.

For our second test, we sort stocks into quintiles based on their beta, size, book-to-market, or momentum, and then within each characteristic portfolio we sort firms into quintile portfolios based on $SVIX_{i,t}^2$. We generate forecasts of portfolio-level expected returns by equally weighting the forecasts of the portfolio's constituent stocks' expected returns, and run regressions corresponding to (18) and (19) using the 5×5 portfolios as test assets. The results are shown in Table VIII. Our model is never rejected. In the specification that is *least* favorable to our theory—the fixed-effects regression with size-sorted portfolios—we find a p -value of 0.07 for the joint hypothesis test; all other p -values are above 0.2, and the estimates of γ are close to 0.5. The corresponding results for excess returns are in Internet Appendix Table IA.V. We find similar results when we conduct the double sort in the opposite direction, first sorting on $SVIX_{i,t}^2$ and then on the other characteristic: see Tables IA.VI and IA.VII.

V. Out-of-Sample Analysis

Formulas (16) and (17) have no free parameters, so it is reasonable to hope that they may be well suited to out-of-sample forecasting. In this section, we show that they are. This fact is particularly striking given the substantial variability of the forecasts both in the time series and in the cross section. The former point is consistent with Martin (2017). The latter is new to this paper, and is illustrated in Figure 9, which plots the evolution of the cross-sectional differences in 1-year expected excess returns generated by our model.

We compare the performance of formulas (16) and (17) to various competitor forecasting benchmarks using an out-of-sample R^2 along the lines of Goyal and Welch (2008). Specifically, we define

$$R_{OS}^2 = 1 - \frac{\sum_i \sum_t FE_{M,it}^2}{\sum_i \sum_t FE_{B,it}^2},$$

Table VIII
Excess-of-Market Returns of Characteristics/SVIX_{i,t}
Double-Sorted Portfolios

This table presents results from regressing excess-of-market returns of portfolios double-sorted on characteristics and risk-neutral variance onto portfolio-level risk-neutral variance ($\overline{\text{SVIX}}_{i,t}^2 - \overline{\text{SVIX}}_t^2$, where $\text{SVIX}_{i,t}^2$ is calculated by equally weighting the portfolio's constituent stocks' risk-neutral variances). The data are monthly from January 1996 to October 2014 and the 1-year horizon of the portfolio returns matches the maturity of the options used to compute $\text{SVIX}_{i,t}^2$ and $\overline{\text{SVIX}}_t^2$. Panel A reports estimates of the pooled panel regression specified in equation (18), $\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha + \gamma(\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2) + \epsilon_{i,t+1}$. Panel B reports estimates of the panel regression with portfolio fixed effects specified in equation (19), $\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha_i + \gamma(\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2) + \epsilon_{i,t+1}$, where $\sum_i w_i \alpha_i$ reports the time-series average of the value-weighted sum of portfolio fixed effects. Values in parentheses are standard errors obtained from the block bootstrap procedure described in Appendix Section B. In each panel, we report the regressions' adjusted R^2 and p -values of Wald tests that test whether the regression coefficients take the values predicted by our theory (zero intercept and $\gamma = 0.5$) and whether $\gamma = 0$. For the pooled panel regressions, the row labelled "theory adj- R^2 (%)" reports the adjusted R^2 obtained when the coefficients are fixed at the values predicted by our theory.

	Beta	Size	B/M	Mom
Panel A: Pooled panel regressions				
α	0.015 (0.019)	0.013 (0.020)	0.015 (0.020)	0.014 (0.019)
γ	0.495 (0.311)	0.572 (0.323)	0.502 (0.327)	0.559 (0.319)
Adjusted R^2 (%)	8.391	9.908	8.098	10.245
$H_0 : \alpha = 0, \gamma = 0.5$	0.635	0.593	0.635	0.613
$H_0 : \gamma = 0.5$	0.987	0.823	0.996	0.890
$H_0 : \gamma = 0$	0.112	0.076	0.125	0.088
Theory adj- R^2 (%)	7.598	8.995	7.232	8.555
Panel B: Panel regressions with portfolio fixed effects				
$\sum_i w_i \alpha_i$	0.015 (0.017)	0.008 (0.005)	0.014 (0.016)	0.019 (0.017)
γ	0.794 (0.490)	0.941 (0.529)	0.711 (0.507)	0.864 (0.491)
Adjusted R^2 (%)	13.010	16.419	12.679	15.020
$H_0 : \sum_i w_i \alpha_i = 0, \gamma = 0.5$	0.439	0.070	0.479	0.212
$H_0 : \gamma = 0.5$	0.549	0.405	0.677	0.459
$H_0 : \gamma = 0$	0.106	0.075	0.161	0.079

where $FE_{M,it}$ and $FE_{B,it}$ denote the forecast errors for stock i at time t based on our model and on a benchmark prediction, respectively. Our model outperforms a given benchmark if the corresponding R_{OS}^2 is positive.

What are the natural competitor benchmarks? One possibility is to give up on trying to make differential predictions across stocks, and simply to use a

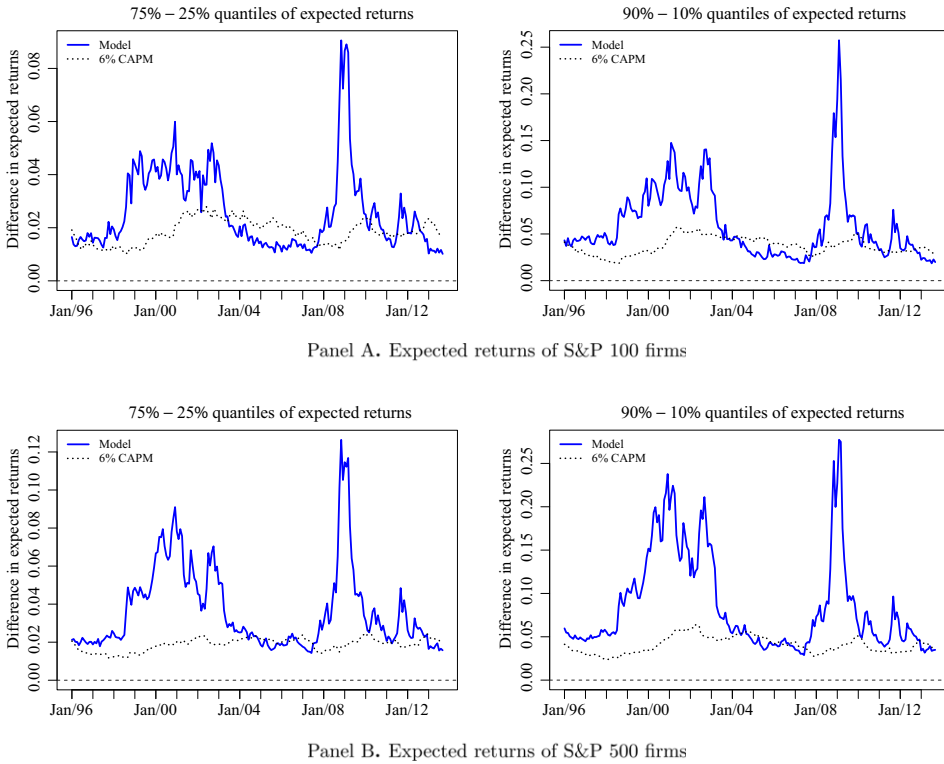


Figure 9. Cross-sectional variation in expected returns. This figure plots time series of cross-sectional differences in 1-year expected excess returns generated by our model and by CAPM forecasts. The CAPM forecasts use conditional betas (estimated from historical returns) and a constant 6% per annum equity premium. The plots show the difference in the 75% and 25% quantiles of expected returns (on the left) and the difference in the 90% and 10% quantiles of expected returns (on the right) for S&P 100 stocks (Panel A) and S&P 500 stocks (Panel B). The data are monthly and cover S&P 500 stocks from January 1996 to October 2014. (Color figure can be viewed at wileyonlinelibrary.com)

forecast of the expected return on the market as a forecast for each individual stock. We consider various ways of doing so. We use the market's historical average excess return as an equity premium forecast, following Goyal and Welch (2008) and Campbell and Thompson (2008), and we use the S&P 500 ($\overline{S\&P\ 500}_t$) and the CRSP value-weighted index (\overline{CRSP}_t) as proxies for the market. We also use the risk-neutral variance of the market, $SVIX_t^2$, to proxy for the equity premium, as suggested by Martin (2017). Lastly, we consider a constant excess return forecast of 6% per annum, corresponding to long-run estimates of the equity premium used in previous research.

More ambitious competitor models would seek to provide differential forecasts of individual firm stock returns, as we do. Again, we consider several alternatives. One natural thought is to use the historical average of firms' stock excess returns ($\overline{RX}_{i,t}$). Another is to estimate firms' conditional CAPM betas

Table IX
Out-of-Sample Forecast Accuracy

This table presents results on the out-of-sample accuracy of our model relative to benchmark predictions. To compare the forecast accuracy of the model to that of the benchmarks, we compute an out-of-sample R^2 , defined as $R_{OS}^2 = 1 - \frac{\sum_i \sum_t FE_M^2}{\sum_i \sum_t FE_B^2}$, where FE_M and FE_B denoted the forecast errors from our model and a benchmark prediction, respectively. Panel A evaluates forecasts of expected equity excess returns, as given in equation (17), and Panel B evaluates forecasts of expected equity returns in excess of the market return, as given in equation (16). The data are monthly and cover S&P 500 stocks from January 1996 to October 2014. The column labels indicate the return horizons ranging from 1 month to 2 years. The return horizons match the maturities of the options used to compute $SVIX_t^2$, $SVIX_{i,t}^2$, and \overline{SVIX}_t^2 . For Panel A, the benchmark forecasts are the risk-neutral market variance ($SVIX_t^2$), the time- t historical average excess returns of the S&P 500 ($\overline{S\&P\ 500}_t$) and the CRSP value-weighted index ($CRSP_t$), a constant prediction of 6% per annum, the stock's risk-neutral variance ($SVIX_{i,t}^2$), the time t historical average of the firms' stock excess returns ($\overline{RX}_{i,t}$), and conditional CAPM implied predictions, where we estimate the CAPM betas from historical return data. For Panel B, we use $SVIX_{i,t}^2$, a random walk (i.e., zero return forecast), and the conditional CAPM as benchmarks.

Horizon	30 days	91 days	182 days	365 days	730 days
Panel A: Expected excess returns					
$SVIX_t^2$	0.09	0.57	1.77	3.08	2.77
$\overline{S\&P\ 500}_t$	0.09	0.79	2.56	3.82	4.46
$CRSP_t$	-0.09	0.24	1.43	1.70	0.88
6% p.a.	-0.01	0.46	1.84	2.54	2.06
$SVIX_{i,t}^2$	0.95	1.87	1.55	2.17	7.64
$\overline{RX}_{i,t}$	1.40	4.97	11.79	27.10	56.67
$\widehat{\beta}_{i,t} \times \overline{S\&P\ 500}_t$	0.09	0.79	2.54	3.76	4.72
$\widehat{\beta}_{i,t} \times CRSP_t$	-0.06	0.28	1.46	1.68	1.61
$\widehat{\beta}_{i,t} \times SVIX_t^2$	0.04	0.46	1.58	2.87	2.91
$\widehat{\beta}_{i,t} \times 6\% \text{ p.a.}$	0.00	0.47	1.84	2.48	2.58
Panel B: Expected returns in excess of the market					
Random walk	0.16	0.76	1.92	3.07	1.99
$(\widehat{\beta}_{i,t} - 1) \times \overline{S\&P\ 500}_t$	0.18	0.80	1.98	3.10	2.17
$(\widehat{\beta}_{i,t} - 1) \times CRSP_t$	0.21	0.89	2.14	3.35	2.83
$(\widehat{\beta}_{i,t} - 1) \times SVIX_t^2$	0.11	0.62	1.68	2.80	2.01
$(\widehat{\beta}_{i,t} - 1) \times 6\% \text{ p.a.}$	0.19	0.83	2.04	3.19	2.49

from historical return data and combine the beta estimates with the aforementioned market premium predictions. We also consider firm-level risk-neutral variance ($SVIX_{i,t}^2$) as a competitor forecasting variable, motivated by Kadan and Tang (2018), who show that under certain conditions $SVIX_{i,t}^2$ provides a lower bound on stock i 's risk premium.

The results for expected excess returns are shown in Panel A of Table IX. Our formula (17) outperforms all of the above competitors at the 3-, 6-, 12-, and 24-month horizons, and its relative performance (as measured by R_{OS}^2) almost invariably increases with forecast horizon, at least up to the 1-year

Table X
Model Out-of-Sample Forecasts versus In-Sample
Benchmark Predictions

This table presents results on the out-of-sample accuracy of our model relative to benchmark predictions that also include in-sample information on returns and/or firm characteristics. To compare the forecast accuracy of the model to that of the benchmarks, we compute an out-of-sample R^2 , defined as $R_{OS}^2 = 1 - \frac{\sum_i \sum_t FE_M^2}{\sum_i \sum_t FE_B^2}$, where FE_M and FE_B denoted the forecast errors from our model and a benchmark prediction, respectively. Panel A evaluates forecasts of expected equity excess returns, as given in equation (17), and Panel B evaluates forecasts of expected equity returns in excess of the market return, as given in equation (16). The data are monthly and cover S&P 500 stocks from January 1996 to October 2014. The column labels indicate the return horizons ranging from 1 month to 2 years. The return horizons match the maturities of the options used to compute $SVIX_t^2$, $SVIX_{i,t}^2$, and \overline{SVIX}_t^2 . For Panel A, the benchmark forecasts are the in-sample average market excess return, a conditional CAPM forecast that uses the in-sample average market excess return as an estimate of the equity premium, the in-sample average return across all stocks, and the fitted values of predictive in-sample regressions of stock returns in excess of the market on CAPM betas, log market capitalization, book-to-market ratios, stock momentum, and all four firm characteristics. For Panel B, we use analogous predictions based on returns in excess of the market.

Horizon	30 days	91 days	182 days	365 days	730 days
Panel A: Expected excess returns					
In-sample avg mkt	-0.05	0.31	1.52	1.90	1.42
In-sample avg all stocks	-0.09	0.17	1.26	1.42	0.56
$\widehat{\beta}_{i,t} \times$ in-sample avg mkt	-0.03	0.34	1.54	1.87	2.04
Beta $_{i,t}$	-0.09	0.16	1.22	1.30	0.56
log(Size $_{i,t}$)	-0.19	-0.17	0.62	0.21	-1.34
B/M $_{i,t}$	-0.18	-0.03	0.89	0.77	0.00
Ret $_{i,t}^{(12,1)}$	-0.10	0.15	1.09	1.05	-0.76
All	-0.25	-0.30	0.26	-0.53	-2.71
Panel B: Expected returns in excess of the market					
In-sample avg all stocks	0.11	0.58	1.60	2.48	0.95
$(\widehat{\beta}_{i,t} - 1) \times$ in-sample avg mkt	0.20	0.86	2.11	3.29	2.63
Beta $_{i,t}$	0.11	0.58	1.60	2.45	0.95
log(Size $_{i,t}$)	0.05	0.39	1.27	1.90	0.12
B/M $_{i,t}$	0.07	0.50	1.47	2.31	0.88
Ret $_{i,t}^{(12,1)}$	0.10	0.56	1.47	2.05	0.03
All	0.03	0.34	1.11	1.46	-0.64

horizon.¹⁴ At the 1-year horizon, R_{OS}^2 ranges from 1.68% to 3.82% depending on the competitor benchmark, with the exception of the historical average stock return, $\overline{RX}_{i,t}$, which the formula outperforms by a much wider margin, with an R_{OS}^2 above 27%. This dramatic outperformance reflects an advantage of our approach: it does not rely on historical data. This is particularly important

¹⁴ The R_{OS}^2 results are based on expected excess returns defined as $\mathbb{E}_t R_{i,t+1} - R_{f,t+1}$, that is, we multiply the left and the right sides of equations (16) and (17) by $R_{f,t+1}$.

for stocks with short return histories that may not be representative of future returns. For example, at the peak of the dot-com bubble young tech firms had extremely high historical average returns over their short histories. In such cases, employing the historical average as a predictor may lead to large forecast errors for subsequent returns.

The results for expected returns in excess of the market, which focus on the purely cross-sectional dimension of stock returns, are shown in Panel B and are, if anything, even stronger. We adjust the conditional CAPM predictions appropriately (by multiplying the equity premium by beta minus one), and we add a “random walk” forecast of zero. Formula (16) outperforms all of the competitors at every horizon, with the outperformance increasing in forecast horizon up to 1 year. At the 1-year horizon, R_{OS}^2 is around 3% relative to each of the benchmarks.

More surprisingly, our model is competitive with—and at horizons of 6 months or more, typically outperforms—a range of predictions based on *in-sample* information. The first three lines of Table X, Panel A, compare the performance of the excess-return formula (17) to the in-sample average equity premium and to the in-sample average excess return on a stock (each of which makes the same forecast for every stock’s return), and also to estimated beta multiplied by the in-sample equity premium (which differentiates across stocks). In each case, R_{OS}^2 is increasing with forecast horizon up to 1 year and is positive at horizons of 3, 6, 12, and 24 months.

The next five lines compare the model forecasts to in-sample predictions based on firm characteristics. We calculate in-sample predictions as the fitted values from pooled univariate regressions (with intercepts) of excess returns onto conditional betas, log size, book-to-market ratios, or past returns, and from a pooled multivariate regression onto all four characteristics. The formula outperforms each of the individual characteristics at horizons of 6 and 12 months, and is competitive with the multivariate model.

The corresponding results for returns in excess of the market are shown in Panel B of Table X. Formula (16) outperforms the univariate characteristics-based competitors at all horizons; remarkably, it even beats the in-sample *multivariate* model at horizons from 1 month to 1 year.

VI. Conclusion

We conclude by highlighting some distinctive features of our approach to the cross section of expected stock returns.

First, our theory identifies, *ex ante*, specific variables that should forecast stock returns. A comparison can be made with the CAPM, which identifies market betas as the relevant quantities. But market betas must be estimated if this prediction is to be tested. At times, when markets are turbulent, historical betas may not accurately reflect the idealized forward-looking betas called for by the CAPM, or by factor models more generally; and if the goal is to forecast returns over, say, a 1-year horizon, one cannot respond to this critique by taking refuge in the last 5 minutes of high-frequency data. In contrast, our predictive

variables, which are based on option prices, are observable in real time and inherently forward-looking.

Second, the theory makes quantitative predictions about the signs and sizes of the coefficients on these predictive variables in forecasting regressions. By contrast, the factor model approach to the cross section has both the advantage and the disadvantage of imposing almost no structure, and therefore says *ex ante* little about the anticipated signs, and nothing about the sizes, of coefficient estimates. (The CAPM does predict that the slope of the security market line should equal the market risk premium, but it is silent on the size of the market risk premium.)

Our approach performs well in and out of sample, particularly over 6-, 12-, and 24-month horizons. The model does a good job of accounting for realized returns on portfolios sorted on characteristics known to be problematic for previous generations of asset pricing models. When we run stock-level panel regressions of realized returns onto characteristics and our volatility predictor variables, our volatility variables drive out the characteristics and are themselves statistically significant, and we do not reject the hypothesis that the associated coefficients take the values predicted by our theory.

As the coefficients in the formula for the expected return on a stock are theoretically motivated, we need only observe the market prices of certain options to implement the formula. No estimation is required, so we avoid the critique of Goyal and Welch (2008). We show, moreover, that the formula outperforms a range of competitor predictors out of sample—even competitors with knowledge of the in-sample relationship between expected returns and characteristics.

Our real-time measure of the expected return on a stock has many potential applications in asset pricing and corporate finance. For example, we are currently exploring the reaction of expected stock returns to macroeconomic and firm-specific news announcements. As expected (or “required”) rates of return are a key determinant of investment decisions, our results also have important implications for macroeconomics more generally—notably because our approach generates considerably more variation in expected returns, both over time and across stocks, than does, say, the CAPM. This points toward a quantitatively and qualitatively new view of risk premia.

Appendix A: A Measure of Correlation

In this section, we show that the ratio $SVIX_t^2 / \overline{SVIX}_t^2$ can be interpreted as an approximate measure of average risk-neutral correlation between stocks. Note first that

$$\begin{aligned} & \text{var}_t^* R_{m,t+1} - \sum_i w_{i,t}^2 \text{var}_t^* R_{i,t+1} \\ &= \sum_{i \neq j} w_{i,t} w_{j,t} \text{corr}_t^*(R_{i,t+1}, R_{j,t+1}) \sqrt{\text{var}_t^* R_{i,t+1} \text{var}_t^* R_{j,t+1}}, \end{aligned}$$

so we can define a measure of average correlation, ρ_t , as

$$\rho_t = \frac{\text{var}_t^* R_{m,t+1} - \sum_i w_{i,t}^2 \text{var}_t^* R_{i,t+1}}{\sum_{i \neq j} w_{i,t} w_{j,t} \sqrt{\text{var}_t^* R_{i,t+1} \text{var}_t^* R_{j,t+1}}}$$

Now, we have

$$\begin{aligned} \rho_t &\approx \frac{\text{var}_t^* R_{m,t+1}}{\sum_i w_{i,t}^2 \text{var}_t^* R_{i,t+1} + \sum_{i \neq j} w_{i,t} w_{j,t} \sqrt{\text{var}_t^* R_{i,t+1} \text{var}_t^* R_{j,t+1}}} \\ &= \frac{\text{var}_t^* R_{m,t+1}}{\left(\sum_i w_{i,t} \sqrt{\text{var}_t^* R_{i,t+1}}\right)^2}. \end{aligned}$$

This last expression features the square of average stock volatility, rather than average stock variance, in the denominator, but we can approximate $\left(\sum_i w_{i,t} \sqrt{\text{var}_t^* R_{i,t+1}}\right)^2 \approx \sum_i w_{i,t} \text{var}_t^* R_{i,t+1}$. (The approximation neglects a Jensen’s inequality term: the left-hand side is strictly smaller than the right-hand side.) This leads us to the correlation measure

$$\rho_t \approx \frac{\text{var}_t^* R_{m,t+1}}{\sum_i w_{i,t} \text{var}_t^* R_{i,t+1}} = \frac{SVIX_t^2}{\overline{SVIX}_t^2}. \tag{A1}$$

Appendix B: Bootstrap Procedure

Our empirical analysis uses a large set of panel data in which residuals may be correlated across firms and over time. Petersen (2009) provides an extensive discussion of how such cross-sectional and time-series dependencies in panel data may bias standard errors in OLS regressions and suggests using two-way clustered standard errors. In further analysis, he finds that standard errors obtained from a bootstrap procedure based on firm clusters are identical to the two-way-clustered standard errors in his panel data. We choose to work with bootstrap standard errors because this is the more conservative approach in our setup for two reasons. First, our monthly data generate overlapping observations at return horizons exceeding 1 month. Second, our data are characterized by high but less than perfect coverage of the cross section of index constituent firms, due to limited availability of option data.

To alleviate biases in standard errors that arise from applying asymptotic theory to finite samples, we use a nonparametric bootstrap procedure based

on resampling. More specifically, because our data are characterized by time-series dependence, we use an overlapping block resampling scheme (originally proposed by Kuensch, 1989) to handle serial correlation and heteroskedasticity. The block bootstrap procedure also takes cross-sectional dependencies into account. Using a large number of bootstrap samples, we estimate the bootstrap covariance matrix and estimate Wald statistics, as we describe in more detail in Appendix Section B1 below. In Section B2 of the Appendix, we provide simulation evidence on the finite-sample properties of the block bootstrap procedure. Detailed results are in the Internet Appendix.

B.1. Implementation of the Block Bootstrap Procedure

We first describe details of the block bootstrap procedure that we apply for pooled panel regressions of returns in excess of the market. Next, we discuss adjustments to the procedure in regressions of excess returns (instead of excess-of-market returns) and adjustments to the procedure when using firm fixed-effects regressions (instead of pooled panel regressions). We then discuss adjustments for portfolio regressions (compared to regressions at the individual firm level).

Pooled Panel Regressions of Returns in Excess of the Market: We use a block bootstrap approach to generate $b = 1, \dots, B$ bootstrap samples by resampling from the actual panel data, as suggested by Kuensch (1989). From the actual data, we need dates, firm identifiers, firms' stock returns in excess of the risk-free rate, firms' risk-neutral variances ($\text{SVIX}_{i,t}^2$), and firms' market capitalizations.

- (1) We generate $B = 1,000$ bootstrap samples of panel data, where the number of time periods in each sample matches the number of time periods in the actual data. More specifically, we generate a bootstrap sample b as follows:
 - (a) Start the resampling procedure by randomly drawing a block of time-length T , which corresponds to the return prediction horizon and the maturity of the options used to compute the SVIX quantities.¹⁵ From the block drawn, randomly select a subset of firms.¹⁶

¹⁵ In time-series bootstraps, it is possible to implement automated procedures that determine the block length based on the properties of the time series (e.g., Politis and White (2004) and Patton, Politis, and White (2009)). These procedures are not implementable in our panel data setup as different firm time series may suggest different block lengths but we need to choose a single block length across all firms to account for the cross-sectional dependencies in the data over time. For instance, for $T = 12$ months, we find that applying such a procedure for different firm time series of $\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}_t^2}$ would suggest block lengths between approximately 8 and 24 months. We repeat our bootstrap procedure with these block lengths of 8 and 24 months, instead of 12 months, and find that our conclusions remain unchanged. We therefore set the block length equal to return horizon T to account for overlapping observations and follow the suggestion of Lahiri (1999) of keeping the block length fixed to allow for overlaps in the blocks.

¹⁶ The idea is to account for the empirical reality that options data may not be available for all firms. For the large number of bootstrap samples $B = 1,000$ that we use, the results of randomly

- (b) Draw further (overlapping) blocks, with replacement, until the bootstrap sample has the same number of time periods as the actual data.
- (c) For every point in time in the bootstrap sample b , determine the firms' market weights and compute the value-weighted average of individual stocks' risk-neutral variance, that is, $\overline{\text{SVIX}}_t^2 = \sum_i w_{i,t} \text{SVIX}_{i,t}^2$, the market return as the return on the value-weighted portfolio, and the stocks' returns in excess of the market.
- (2) For each bootstrap sample, run the pooled panel regression of returns in excess of the market onto risk-neutral excess stock variance,

$$\frac{R_{i,t+1} - R_{m,t+1}}{R_{f,t+1}} = \alpha + \gamma \left(\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2 \right) + \epsilon_{i,t+1},$$

and collect the $B = 1,000$ bootstrap estimates of α and γ .

- (3) Using the $B = 1,000$ bootstrap estimates of α and γ , compute the bootstrap covariance matrix of α and γ . Using this bootstrap covariance matrix, we compute Wald statistics for hypothesis tests.¹⁷ Building on the asymptotic refinement achieved from bootstrapping the covariance matrix, we use the Wald tests' asymptotic distribution to compute the p -values. We explore the finite-sample properties of this bootstrap procedure in Section B2 of the Appendix; our simulation evidence suggests that the approach works well.

Pooled Panel Regressions of Excess Returns: The bootstrap procedure for pooled panel regressions of excess returns is essentially the same as the one described for returns in excess of the market above. The only modifications are:

- In Step 1, we also include the risk-neutral market variance (SVIX_t^2) in the resampling procedure.
- In Step 2, we run the regression of excess returns on SVIX_t^2 and $\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2$, and collect the $B = 1,000$ bootstrap estimates of α , β , and γ .
- In Step 3, we compute the bootstrap covariance matrix for α , β , and γ and use it to compute standard errors and to conduct hypothesis tests.

selecting a subset of firms or including all firms that are available in a drawn block leads to identical results. Conceptually, our approach is similar to the bootstrap using firm clusters described by Petersen (2009) in his footnote 12.

¹⁷ We prefer to compute the Wald statistic based on the bootstrap covariance matrix rather than to bootstrap the Wald statistic because our approach explicitly takes cross-sectional dependencies as well as overlapping observations and other time-dependencies into account. Qualitatively, our results are very similar when we bootstrap Wald statistics that are computed using a double-clustered covariance matrix as suggested by Petersen (2009). The quantitative bootstrap results of the Wald tests can be quite different when using a nonclustered covariance matrix, but we would still not reject the model. As a further check, we also verified that the p -values of bootstrapped likelihood ratio test statistics are identical to those of the bootstrapped Wald statistics computed from nonclustered covariance matrices.

Regressions with Firm Fixed Effects: For the bootstraps of the firm fixed-effects regressions, we adjust the procedure for the pooled panel regressions described above as follows:

- In Step 2, we run the regression with firm fixed effects α_i (instead of the intercept α) and
 - we compute the value-weighted sum of firm fixed effects at every date in every bootstrap sample, that is, $\alpha_t = \sum_i w_{i,t} \alpha_i$
 - in each bootstrap sample, we compute $\bar{\alpha}$ as the time-series average of α_t
 - we collect the $B = 1,000$ estimates of $\bar{\alpha}$ (instead of intercept α)
- In Step 3, we compute the bootstrap covariance matrix with $\bar{\alpha}$ (instead of intercept α) and use it to compute standard errors and to conduct hypothesis tests.

Portfolio Regressions: The bootstraps for pooled panel and fixed-effects regressions using excess returns and excess-of-market returns of portfolios follow the corresponding firm-level procedures described above. The only difference is that in Step 1(a) we use all portfolios rather than resampling in the cross section, as we have a balanced panel of portfolio data.

B.2. Finite-Sample Properties of the Block Bootstrap Procedure

To provide evidence for the reliability of our bootstrap procedure in finite samples, we conduct a simulation study. We simulate S samples on which we impose the null hypothesis and within each sample we repeat the bootstrap procedure from Section B1 above with B iterations. We then compare the empirical quantiles of the Wald statistic in the simulated data to the quantiles of the χ^2 distribution, that is, the Wald statistic's asymptotic distribution. These results suggest that our procedure, using the bootstrap covariance matrix to compute the Wald statistic and then using the asymptotic distribution to infer its p -value, is reasonable. Next, we compare the rejection frequency for the null hypothesis in the simulated data (on which we imposed the null hypothesis) to the nominal size of the test. These results provide further support for our empirical approach.

Given the enormous computational demand of this exercise with an additional $S \times B$ bootstrap samples to be generated and evaluated, we focus on the pooled panel regressions of S&P 100 firms' returns in excess of the market. We simulate data under the null hypothesis by imposing $\alpha = 0$ and $\gamma = 0.5$ and drawing blocks of innovations from the regression residuals (from the specification in Panel A of Table II). The block resampling scheme follows the approach described above in Section B1 and again serves to account for cross-sectional and time-series dependencies. We start by setting the number of simulations to $S = 200$ and the number of bootstrap iterations to $B = 99$, following the choice of Piatti and Trojani (2014) in a similar double-bootstrap exercise. We show that the results are similar when we increase the number of simulations to $S = 400$ and the number of bootstrap iterations to $B = 198$. The subsequent

discussion is based on the results for the 1-year horizon. We then show that our conclusions are very similar for other horizons.

Empirical and Asymptotic Quantiles of the Wald Statistic: Panel A in Figure IA.6 compares the empirical quantiles of the Wald statistic in the simulated data to the quantiles of the Wald statistic's asymptotic χ^2 distribution. With vertical lines marking the 90%, 95%, and 99% quantiles, the plot shows that the empirical quantiles are virtually identical to the quantiles of the χ^2 distribution beyond the 95% quantile; only in the very far tails of the distribution do the critical values from the empirical distribution exceed those from the χ^2 distribution. These results suggest that our approach of using the bootstrap covariance matrix to compute the Wald statistic, and then using the asymptotic distribution to infer the p -value of the Wald statistic, should work well.

Nominal Size and Empirical Rejection Frequencies: Panel B in Figure IA.6 compares the empirical rejection frequencies of our bootstrap approach when applied to simulated data (on which we impose the null hypothesis) to the corresponding nominal size of the test. That is, we compute the fraction of samples in which the bootstrap procedure leads to a rejection of the hypothesis when using the nominal size given on the x -axis. Similar to Panel A, the dotted and dashed lines plot the 90%/10%, 95%/5%, and 99%/1% quantiles to mark the economically interesting regions, where we care about rejections. We find that empirical rejection frequencies are well aligned with nominal size, particularly within the economically interesting regions, and that differences in empirical rejection frequencies and nominal size are too small to lead to incorrect inference in our empirical analysis. To illustrate this, the large symbol in the plot indicates the p -value of the Wald statistic that we obtain from our empirical test of the model in the data; this p -value is 0.437 as reported in Panel A of Table II. These results suggest that our empirical approach performs well.

Figure IA.7 shows that the empirical quantiles of the Wald statistic in the simulated data also line up well with quantiles of the Wald statistic's asymptotic χ^2 distribution at horizons of 3 and 6 months. At the shortest (longest) horizon of one (24) month(s), the empirical quantiles appear somewhat too low (high) compared to the asymptotic quantiles. Nonetheless, the comparison of empirical rejection frequencies in the simulated data to the nominal sizes used in the tests in Figure IA.8 suggests that our approach performs well at all horizons. All results are very similar when increasing the number of simulations and bootstrap iterations to $S = 400$ and $B = 198$ as we show in Figure IA.9. Overall, the alignment of empirical rejection frequencies in the simulated data with nominal sizes used in the tests improves slightly when increasing S to 400 and B to 198.

REFERENCES

- An, Byeong-Je, Andrew Ang, Turan G. Bali, and Nusret Cakici, 2014, The joint cross section of stocks and options, *Journal of Finance* 69, 2279–2337.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.

- Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.
- Breedon, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.
- Buss, Adrian, and Grigory Vilkov, 2012, Measuring equity risk with option-implied correlations, *Review of Financial Studies* 25, 3113–3140.
- Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies* 21, 1509–1531.
- Carhart, Mark, 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.
- Chang, Bo-Young, Peter Christoffersen, Kris Jacobs, and Gregory Vainberg, 2012, Option-implied measures of equity risk, *Review of Finance* 16, 385–428.
- Christoffersen, Peter, Mathieu Fournier, and Kris Jacobs, 2017, The factor structure in equity options, *Review of Financial Studies* 31, 595–637.
- Conrad, Jennifer, Robert F. Dittmar, and Eric Ghysels, 2013, Ex ante skewness and expected stock returns, *Journal of Finance* 68, 85–124.
- Dybvig, Philip H., and Stephen A. Ross, 1985, Yes, the APT is testable, *Journal of Finance* 40, 1173–1188.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.
- French, Dan W., John C. Groth, and James W. Kolari, 1983, Current investor expectations and better betas, *The Journal of Portfolio Management* 10, 12–17.
- Fu, Fangjian, 2009, Idiosyncratic risk and the cross-section of expected stock returns, *Journal of Financial Economics* 91, 24–37.
- Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.
- Hansen, Lars P., and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.
- Hansen, Lars P., and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.
- 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, 249–283.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65–91.
- Kadan, Ohad, and Xiaoxiao Tang, 2018, A bound on expected stock returns, Working paper, Washington University in St. Louis.
- Kremens, Lukas, and Ian Martin, 2019, The quanto theory of exchange rates, *American Economic Review* 109, 810–843.
- Kuensch, Hans R., 1989, The jackknife and the bootstrap for general stationary observations, *Annals of Statistics* 17, 1217–1241.
- Lahiri, Soumendra N., 1999, Theoretical comparisons of block bootstrap methods, *Annals of Statistics* 27, 386–404.
- Lewellen, Jonathan, 2015, The cross-section of expected stock returns, *Critical Finance Review* 4, 1–44.
- Long, John B., 1990, The numeraire portfolio, *Journal of Financial Economics* 26, 29–69.
- Martin, Ian, 2017, What is the expected return on the market? *Quarterly Journal of Economics* 132, 367–433.
- Martin, Ian, 2018, Options and the gamma knife, *Journal of Portfolio Management* 44, 47–55.
- Nagel, Stefan, 2005, Short sales, institutional investors and the cross-section of stock returns, *Journal of Financial Economics* 78, 277–309.

- Patton, Andrew, Dimitris N. Politis, and Halbert White, 2009, Correction to “Automatic block-length selection for dependent bootstrap,” *Econometric Reviews* 28, 372–375.
- Petersen, Mitchell A., 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of Financial Studies* 22, 435–480.
- Piatti, Ilaria, and Fabio Trojani, 2014, Dividend growth predictability and the price-dividend ratio. forthcoming.
- Politis, Dimitris N., and Halbert White, 2004, Automatic block-length selection for dependent bootstrap, *Econometric Reviews* 23, 53–70.
- Roll, Richard, 1973, Evidence on the “growth-optimum” model, *Journal of Finance* 28, 551–566.
- Schneider, Paul, and Fabio Trojani, 2019, (Almost) model-free recovery, *Journal of Finance* 74, 323–370.
- Shanken, Jay, 1982, The arbitrage pricing theory: Is it testable? *Journal of Finance* 37, 1129–1140.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher’s website:

Appendix S1: Internet Appendix.
Replication Code.