A Bound on Expected Stock Returns*

Ohad Kadan† and Xiaoxiao Tang‡

November 2015

Abstract

We present a sufficient condition under which the prices of options with different strike prices written on a particular stock can be used to calculate a lower bound on the expected returns of that stock. The sufficient condition imposes a restriction on a combination of the stock's systematic and idiosyncratic risk. The lower bound is forward-looking and can be calculated on a high-frequency basis for stocks with liquid option trading. We estimate the lower bound empirically for constituents of the S&P 500 index and study its cross-sectional properties. We find that the bound increases with beta and decreases with size. The bound also provides an economically meaningful signal on future realized stock returns.

1 Introduction

A traditional point of view in the asset pricing literature is that option prices can teach us a great deal about the volatility of underlying assets. The idea is that option prices along with some assumptions on the price process (such as in Black and Scholes (1973)) can be used to derive a forward looking estimate of the return volatility of the underlying asset. More generally, beginning with the work of Breeden and Litzenberger (1978), it has been well understood that option prices can be used to extract forward-looking risk neutral probabilities, allowing one to price a variety of assets. What remains a puzzle to this day is whether option prices can be used to elicit any useful information on forward-looking expected returns, derived from the physical (true) distribution of the returns of the underlying asset.[1] It is for this reason that the Recovery theorem, proposed by Ross (2015), ignited much controversy, as it suggests that given a regularity condition on the pricing kernel, option prices can in fact be used to recover the physical probability distribution of asset returns,

*We thank Kerry Back, Phil Dybvig, Ian Martin, Luca Pezzo, seminar participants at University of Adelaide, Rice, University of Sydney, UNSW, Southwest Institute of Finance and Economics, and Washington University in St. Louis as well as participants at the IDC Summer Conference for helpful comments and suggestions.

†Olin Business School, Washington University in St. Louis, kadan@wustl.edu.

‡Olin Business School, Washington University in St. Louis, xtang22@wustl.edu.

[1] The conventional approach to estimating pricing kernels from option prices relies on using the risk-neutral probability and then supplementing it with a physical probability derived from historical returns. See Jackwerth (2000) and Ait-Sahalia and Lo (2000).

and, in particular their implied expected returns. Ross's result, however, has been criticized by Borovička, Hansen, and Scheinkman (2014), who argue that it restricts the dynamics of the stochastic discount factor in an unrealistic manner, and thereby, the recovered probability distribution differs from the true one.

Martin (2015) takes a different approach to address this problem. He derives the following simple and yet useful relation for the expected return of any asset j,

$$\mathrm{E}_t R_{j,t+1} - R_{f,t} = \frac{Var_t^*\left(R_{j,t+1}\right)}{R_{f,t}} - Cov_t\left(M_{t+1} R_{j,t+1}, R_{j,t+1}\right). \qquad (*)$$

Namely, the expected excess returns of asset j as of time t, equals the risk-neutral variance of the asset's returns discounted at the risk free rate, less the covariance between the product of the stochastic discount factor and the asset return, and the asset return.[2] Martin then restricts attention to the case in which asset j is the market, and he explores whether $Cov_t\left(M_{t+1} R_{m,t+1}, R_{m,t+1}\right)$ is weakly negative, where $R_{m,t+1}$ is the return on the market. He terms this the *Negative Correlation Condition* (NCC). The idea is that if $Cov_t\left(M_{t+1} R_{m,t+1}, R_{m,t+1}\right) \leq 0$, then (*) implies that the deflated risk-neutral variance can be used as a lower bound on the expected market premium,

$$\mathrm{E}_t R_{m,t+1} - R_{f,t} \geq \frac{Var_t^*\left(R_{m,t+1}\right)}{R_{f,t}}. \qquad (**)$$

Moreover, Martin (2015) shows that $Var_t^*\left(R_{m,t+1}\right)$ can be relatively easily calculated from the prices of options written on the S&P 500 index with different strike prices.[3] Thus, to the extent that the NCC holds, Martin's argument presents us with a simple procedure to calculate a high-frequency lower bound on expected market returns. Martin then goes on to show that the NCC indeed holds theoretically under very mild conditions in a variety of asset pricing settings. Essentially, what is needed to justify the NCC is that relative risk aversion in the economy be no less than 1. Martin also shows that the NCC holds empirically given typical factor structures for the stochastic discount factor.

Martin's approach is appealing. His assumptions are very mild, leading to a credible forward-looking estimate of a lower bound on the market premium. Furthermore, his analysis suggests that the lower bound he obtains is tight.

In this paper we follow Martin's approach noting that (*) holds true for any asset. The question then is whether the NCC holds for individual assets as well. If it does, then one can use (**) to calculate a lower bound on expected stock returns cross-sectionally using options written on individual stocks. Intuitively, since the market is a weighted average of individual assets, and since the NCC holds for the market as a whole, one would expect the NCC to hold for a significant portion of individual assets. However, Martin's approach does not provide us with a procedure to test whether a particular stock satisfies the NCC. In our main theoretical contribution we provide

[2] As in Martin (2015), a star is used to denote moments under the risk-neutral measure. Thus, Var_t^* refers to a risk-neutral variance whereas Cov_t refers to a physical covariance.

[3] Martin (2013) shows that the risk neutral variance is closely related to an index of simple variance swaps (SVIX).

such a tool by deriving a sufficient condition for the NCC to hold for individual stocks. Our condition, which relies on a first-order approximation, establishes that the NCC holds for stocks with a relatively low combination of systematic and idiosyncratic risk compared to the relative risk-aversion in the market.

Our sufficient condition is easily checked empirically provided that one makes some plausible assumption on the level of relative risk aversion in the economy. We show that given conventional estimates of relative risk aversion, the NCC holds for 50%-90% of S&P 500 constituents, where the range depends on the level of conservatism one would like to impose on the assumed risk aversion. It follows that for a large cross-section of stocks we can modify (**) to obtain the following lower bound on the expected returns of individual stocks

$$\mathrm{E}_t R_{j,t+1} - R_{f,t} \geq \frac{Var_t^*\left(R_{j,t+1}\right)}{R_{f,t}}. \tag{***}$$

This lower bound can be estimated empirically from options with different strike prices using the procedure developed in Martin (2015).[4] We need, however, to restrict our empirical analysis to stocks that have highly liquid traded options so that the lower bound can be calculated in a sufficiently precise manner. To this end, we confine our examination to 652 constituents of the S&P 500 index during the years 2005-2014, which pass our screens for the availability of option price data. This approach yields 575,197 stock/day combinations for which we can estimate the lower bound. We then use this extensive panel to study the properties of these lower bounds empirically.

The average lower bound during our sample period ranges from 8.4% to 19.5% depending on the level of conservatism one exerts when applying the NCC. The average lower bound drops significantly to 6.7%-15.2% when we exclude the crisis period of 2008-2009. Indeed, during the crisis period we observe a dramatic spike in expected stock returns consistent with the findings of Martin (2015) on the market premium. Cross-sectionally we find that the lower bound decreases with firm size, in line with the classic Fama and French (1992) findings, which are based on realized rather than forward-looking estimates. Strikingly, and in contrast to the Fama-French findings, the lower bound is strongly correlated with the CAPM beta. Thus, while beta does not seem to explain cross-sectional variations in realized stock returns, it appears to be a strong determinant of investors' expectations as reflected in option prices. We do not find any relation between the firm's book-to-market ratio and the lower bound. Additionally, momentum appears to have a moderate effect on the lower bound, and to the extent it exists, it works in the opposite direction from the one explored in Jegadeesh and Titman (1993). Namely, being a recent winner induces investors to lower rather than increase their expectations regarding a given stock. These findings are consistent with prior findings using other approaches to estimating forward looking expected returns. In particular, Berk and van Binsbergen

[4] Martin's (2015) approach to estimating the lower bound requires the use of European style options. However, options written on individual stocks in the U.S. are of the American style. Our estimates are mostly based on out-of-the money options for which the difference in price between American and European options is not big. In Section 5 we estimate the error resulting from using American instead of European options, and we find it to typically be rather small.

(2015) use capital flows to/from mutual funds to test asset pricing models and find that the CAPM is the closest model to the model investors use to make capital allocation decisions. Similarly, Brav, Lehavy, and Michaely (2005) calculate forward looking expected returns using analyst expectations and find them to be positively correlated with firm beta and negatively correlated with size.

In our last analysis we ask whether the forward looking lower bounds we obtain have any predictive value. That is, are the lower bounds on expected stock returns correlated with future realized returns? Clearly, our sample period of 10 years is short for a longitudinal asset pricing test, and our estimates are likely quite noisy. Nevertheless, our lower-bound estimates appear to provide a large and economically meaningful (if noisy) signal on future stock returns. The signal appears to be stronger for stocks in which our lower bound is more binding and over horizons of 6-12 months. For example, stocks in the lowest decile of the lower bound on expected returns have average realized returns of 8.6% in the next twelve-months period, while stocks in the highest decile have average realized returns of 36.0%. The difference of 27.5% per year is very large economically.

The paper proceeds as follows. In Section 2 we explain how we derive the lower bounds for individual assets. In Section 3 we explain how the lower bounds can be estimated from option prices. Section 4 discusses our cross-sectional empirical analysis of the lower bounds. In Section 5 we discuss and evaluate the precision of the approximations we are using. We conclude in Section 6.

2 Derivation of the Lower Bound

2.1 General Approach

Our setup follows that of Martin (2015) with the difference being that Martin focuses on the market as a whole, while we focus on the cross-section of individual assets. Consider a standard dynamic asset pricing model with uncertainty at time t about the realization of asset returns at time $T > t$. Assume no arbitrage so that a stochastic discount factor, M_T, exists satisfying $\mathrm{E}_t\left(M_T R_{j,T}\right) = 1$ for the gross returns $R_{j,T}$ from time t to T of any asset j. Denote $R_{f,t} = 1/\mathrm{E}_t\left(M_T\right)$ the gross risk free rate between times t and T. Then, it is standard that

$$\frac{Var_t^*\left(R_{j,T}\right)}{R_{f,t}} = \mathrm{E}_t(M_T R_{j,T}^2) - R_{f,t}, \tag{1}$$

for all assets j, where $Var_t^*(\cdot)$ denotes variance taken under the risk-neutral measure. As Martin (2015) shows, the expected excess return of any asset j is given by

$$\begin{aligned}
\mathrm{E}_t R_{j,T} - R_{f,t} &= \mathrm{E}_t R_{j,T} - R_{f,t} + \mathrm{E}_t(M_T R_{j,T}^2) - \mathrm{E}_t(M_T R_{j,T}^2) \\
&= \frac{Var_t^*(R_{j,T})}{R_{f,t}} - \mathrm{E}_t(M_T R_{j,T}^2) + \mathrm{E}_t R_{j,T} \quad \text{(using (1))} \\
&= \frac{Var_t^*(R_{j,T})}{R_{f,t}} - \left[\mathrm{E}_t(M_T R_{j,T}^2) - \mathrm{E}_t R_{j,T} \mathrm{E}_t\left(M_T R_{j,T}\right)\right] \quad \text{(using that } \mathrm{E}_t\left(M_T R_{j,T}\right) = 1\text{)} \\
&= \frac{Var_t^*(R_{j,T})}{R_{f,t}} - Cov_t\left(M_T R_{j,T}, R_{j,T}\right),
\end{aligned} \tag{2}$$

where $Cov_t(\cdot,\cdot)$ denotes the covariance operator as of time t. Note that in (2) both the expectation and covariance are taken under the physical probability measure whereas the variance is taken under the risk-neutral measure.

Definition 1 *(Martin (2015)). We say that the Negative Correlation Condition (NCC) holds for asset j if $Cov_t\left(M_T R_{j,T}, R_{j,T}\right) \leq 0$.*

Note that (2) implies that if the NCC holds for asset j, then

$$\mathrm{E}_t R_{j,T} - R_{f,t} \geq \frac{Var_t^*(R_{j,T})}{R_{f,t}}. \tag{3}$$

Thus, whenever the NCC holds for asset j, the risk neutral return variance scaled by the risk free return serves as a lower bound for the expected excess return of asset j. On the other hand, if the NCC fails, then the inequality in (3) is reversed, and we obtain an upper bound on expected asset returns.

Martin (2015) restricts attention to the case in which j is the market portfolio. He shows that in this case, the NCC holds in a variety of standard asset pricing models subject to relative risk aversion being at least 1, allowing him to obtain a lower bound on the expected market premium. Our goal here is to show that the NCC holds for many individual assets satisfying a somewhat stricter (but still quite standard) condition. Thus, we can use (3) to obtain cross-sectional bounds on expected stock returns.

2.2 Martin's Argument for the Validity of the NCC

To motivate our analysis we begin by reviewing Martin's (2015) basic argument for why the NCC is satisfied when j is the market portfolio. We then show that Martin's argument does not directly apply for individual assets, leading to our own analysis in the next section.

Martin shows that under a mild condition on risk aversion, when the asset being considered is the market the NCC holds for standard work-horse models in asset pricing including a one-period investment problem, a dynamic consumption/investment model with separable utility, and a dynamic consumption/investment model with recursive utility. For brevity we only review his one period analysis and refer the reader to his paper for the other models.

Assume the existence of a representative agent maximizing expected utility. The utility function $u(\cdot)$ is assumed to be twice differentiable with $u' > 0$ and $u'' < 0$. At time t the agent needs to construct an optimal portfolio out of N assets whose random returns will be realized at time $T > t$ and are denoted by $R_{i,T}$, $i = 1, 2, ...N$. The agent consumes only once at time T. The optimal portfolio is the market portfolio and we denote its return by $R_{m,T}$.

The agent's problem is thus to choose portfolio weights $\{w_i\}$ to solve

$$\max_{\{w_i\}} \mathrm{E}_t u(\sum w_i R_{i,T})$$
$$s.t. \sum w_i = 1.$$

At optimum, we have the first-order condition:[5]

$$\mathrm{E}_t[u'(R_{m,T})R_{i,T}] = \lambda, \forall i, \tag{4}$$

where λ is a positive Lagrange multiplier and the market portfolio is

$$R_{m,T} \equiv \sum_{i=1}^{N} w_i R_{i,T}. \tag{5}$$

Dividing both sides of (4) by λ, we have that

$$\mathrm{E}_t[\frac{u'(R_{m,T})}{\lambda} R_{i,T}] = 1, \forall i.$$

Therefore, $\frac{u'(R_{m,T})}{\lambda}$ is a stochastic discount factor in this economy,[6] and we conclude that M_T is proportional to $u'(R_{m,T})$. Thus, for the NCC to hold in this setup for some specific asset j we need

$$Cov_t\left(u'\left(R_{m,T}\right) R_{j,T}, R_{j,T}\right) \leq 0. \tag{6}$$

In particular, if j is the market portfolio we need

$$Cov_t\left(u'\left(R_{m,T}\right) R_{m,T}, R_{m,T}\right) \leq 0.$$

Denote

$$\gamma\left(w\right) \equiv -w \frac{u''\left(w\right)}{u'\left(w\right)},$$

the coefficient of relative risk aversion at wealth level w. Note that we do not need to assume in this setup that $\gamma\left(w\right)$ is constant or monotone in w. If $\gamma\left(R_{m,T}\right) \geq 1$ then $u'\left(R_{m,T}\right) R_{m,T}$ is a decreasing function of $R_{m,T}$ and therefore $Cov_t\left(u'\left(R_{m,T}\right) R_{m,T}, R_{m,T}\right) \leq 0$.[7] We thus have

[5] To justify this one needs to introduce assumptions on the validity of differentiation of an expectation. We skip these standard technical details for brevity.

[6] It is easy to see that $\lambda = E[R_{m,T} u'(R_{m,T})]$.

[7] This argument relies on the fact that if g is a decreasing function then $Cov\left(g\left(x\right), x\right) \leq 0$.

Proposition 2 *(Martin (2015)). Suppose* $\gamma(R_{m,T}) \geq 1$*, then the NCC holds for the market portfolio. In this case,* $\frac{Var^*(R_{m,T})}{R_{f,t}}$ *serves as a lower bound on the expected market premium between time* t *and* T*.*

Thus, as long as relative risk aversion in the economy is at least 1, $\frac{var^*(R_{m,T})}{R_{f,t}}$ is a lower bound on the expected market premium. As shown by Martin, this lower bound can be calculated from option prices on a high frequency basis.

Martin's argument is a powerful statement on the market portfolio, but it cannot be directly applied to individual assets. To see this note first that the sign of (6) depends on the covariance between $R_{j,T}$ and $u'(R_{m,T})$ and on the covariance between $R_{j,T}$ and itself, i.e., the variance of $R_{j,T}$. The former is related to the systematic risk of the asset and is likely to be negative since u' is decreasing. The latter is related to the idiosyncratic risk of the asset and is always positive. Thus, we have two conflicting effects driving the sign of the covariance in (6). The curvature of u', which is driven by risk aversion, will determine which one of the two effects dominates. Thus, intuitively, determining whether the NCC holds for a particular asset should depend on the asset's systematic risk, idiosyncratic risk, and on the prevailing level of risk aversion.

It is also instructive to note that the monotonicity argument used by Martin to sign the covariance does not extend to the case of individual assets. Indeed, for any given asset j, (6) can be written as

$$Cov_t\left(R_{j,T}u'(\sum_{i=1}^{N} w_i R_{i,T}), R_{j,T}\right) \leq 0. \tag{7}$$

But now, asking that $u'(\sum_{i=1}^{N} w_i R_{i,T})R_{j,T}$ be monotone decreasing in $R_{j,T}$ is no longer sufficient to guarantee that the covariance is negative. In fact, taking a first order derivative with respect to $R_{j,T}$ will often result in a positive value. However, this is not conducive to the conclusion that the covariance is positive.[8] Rather, the sign of the covariance depends on the entire correlation structure between $R_{j,T}$ and all other assets, which is being ignored if one only takes partial derivatives. In the next section we present an alternative approach to signing the covariance in (6).

2.3 The NCC for Individual Assets

To overcome the difficulty of calculating the covariance in (7) we use a first-order approximation leading to sufficient conditions under which the NCC holds for individual assets. As in Martin (2015), our approach is valid for both static and dynamic standard asset pricing models. We present the one-period model here and relegate the dynamic models to Appendix I.

[8] As noted above, Martin (2015) uses the fact that if g is a decreasing function then $Cov(g(x), x) \leq 0$. This result, however, does not extend to multivariate functions. For example, even if $g(x, y)$ is decreasing in x, one cannot generally conclude that $Cov(g(x, y), x) \leq 0$. Rather, this depends also on the covariance between x and y.

Our goal is to find a condition under which (7) is satisfied. Define $f : \mathbb{R}^N \to \mathbb{R}$ by

$$f(R_{1,T}, R_{2,T}, ...R_{N,T}) = R_{j,T} u'(\sum_{i=1}^{N} w_i R_{i,T}). \tag{8}$$

The first order multivariate Taylor expansion of f around $(\mathrm{E}_t R_{m,T}, \mathrm{E}_t R_{m,T}, ...\mathrm{E}_t R_{m,T})$ gives

$$\begin{aligned} f(R_{1,T}, R_{2,T}, ...R_{N,T}) \quad &\approx \quad f(\mathrm{E}_t R_{m,T}, \mathrm{E}_t R_{m,T}, ...\mathrm{E}_t R_{m,T}) \\ &\quad + \sum_{i=1}^{N} f_i(\mathrm{E}_t R_{m,T}, \mathrm{E}_t R_{m,T}, ...\mathrm{E}_t R_{m,T})(R_{i,T} - \mathrm{E}_t R_{m,T}), \end{aligned} \tag{9}$$

where f_i is the partial derivative of f with respect to its ith argument.

Partially differentiating (8) gives for $i \neq j$

$$f_i(R_{1,T}, R_{2,T}, ...R_{N,T}) = w_i R_{j,T} u''(\sum_{i=1}^{N} w_i R_{i,T}) \tag{10}$$

and for $i = j$

$$f_j(R_{1,T}, R_{2,T}, ...R_{N,T}) = u'(\sum_{i=1}^{N} w_i R_{i,T}) + w_j R_{j,T} u''(\sum_{i=1}^{N} w_i R_{i,T}). \tag{11}$$

Using (10) and (11), we can rewrite (9) as

$$\begin{aligned} f(R_{1,T}, R_{2,T}, ...R_{N,T}) \quad &\approx \quad f(\mathrm{E}_t R_{m,T}, \mathrm{E}_t R_{m,T}, ...\mathrm{E}_t R_{m,T}) \\ &\quad + u'(\sum_{i=1}^{N} w_i \mathrm{E}_t R_{m,T})(R_{j,T} - \mathrm{E}_t R_{m,T}) \\ &\quad + \sum_{i=1}^{N} w_i \mathrm{E}_t R_{m,T} u''(\sum_{i=1}^{N} w_i \mathrm{E}_t R_{m,T})(R_{i,T} - \mathrm{E}_t R_{m,T}) \\ &= \quad f(\mathrm{E}_t R_{m,T}, \mathrm{E}_t R_{m,T}, ...\mathrm{E}_t R_{m,T}) \\ &\quad + u'(\mathrm{E}_t R_{m,T})(R_{j,T} - \mathrm{E}_t R_{m,T}) + \sum_{i=1}^{N} w_i \mathrm{E}_t R_{m,T} u''(\mathrm{E}_t R_{m,T})(R_{i,T} - \mathrm{E}_t R_{m,T}). \end{aligned}$$

We can now plug this result into the left hand side of (7) obtaining

$$\begin{aligned}
Cov_t\left(R_{j,T}u'(\sum_{i=1}^{N}w_iR_{i,T}),R_{j,T}\right) &= Cov_t\left(f(R_{1,T},R_{2,T},...R_{N,T}),R_{j,T}\right)\\
&\approx Cov_t\left(u'(\mathrm{E}_tR_{m,T})(R_{j,T}-\mathrm{E}_tR_{m,T}),R_{j,T}\right)\\
&\quad +Cov_t\left(\sum_{i=1}^{N}w_i\mathrm{E}_tR_{m,T}u''(\mathrm{E}_tR_{m,T})(R_{i,T}-\mathrm{E}_tR_{m,T}),R_{j,T}\right)\\
&= u'(\mathrm{E}_tR_{m,T})Var_t\left(R_{j,T}\right)\\
&\quad +\mathrm{E}_tR_{m,T}u''(\mathrm{E}_tR_{m,T})\sum_{i=1}^{N}w_iCov_t\left(R_{i,T},R_{j,T}\right)\\
&= u'(\mathrm{E}_tR_{m,T})Var_t\left(R_{j,T}\right)+\mathrm{E}_tR_{m,T}u''(\mathrm{E}_tR_{m,T})Cov_t\left(R_{m,T},R_{j,T}\right),
\end{aligned}$$

where the penultimate equality follows from the fact that $\mathrm{E}R_{m,T}$ is constant as of time t, and the last equality follows from (5). Equivalently,

$$Cov_t\left(u'\left(R_{m,T}\right)R_{j,T},R_{j,T}\right)\approx u'(\mathrm{E}_tR_{m,T})\left[Var_t\left(R_{j,T}\right)+\mathrm{E}_tR_{m,T}\frac{u''(\mathrm{E}_tR_{m,T})}{u'(\mathrm{E}_tR_{m,T})}Cov_t\left(R_{m,T},R_{j,T}\right)\right]. \tag{12}$$

Then, (12) can be written as

$$\begin{aligned}
Cov_t\left(u'\left(R_{m,T}\right)R_{j,T},R_{j,T}\right) &\approx u'(\mathrm{E}_tR_{m,T})\left[Var_t\left(R_{j,T}\right)-\gamma\left(\mathrm{E}_tR_{m,T}\right)Cov_t\left(R_{m,T},R_{j,T}\right)\right]\\
&= u'(\mathrm{E}_tR_{m,T})Var_t\left(R_{j,T}\right)\gamma\left(\mathrm{E}_tR_{m,T}\right)\left[\frac{1}{\gamma\left(\mathrm{E}_tR_{m,T}\right)}-\frac{Cov_t\left(R_{m,T},R_{j,T}\right)}{Var_t\left(R_{j,T}\right)}\right].
\end{aligned}$$

We conclude that, up to a first order approximation, the NCC holds for asset j if

$$\frac{1}{\gamma\left(\mathrm{E}_tR_{m,T}\right)}\leq\frac{Cov_t\left(R_{m,T},R_{j,T}\right)}{Var_t\left(R_{j,T}\right)}. \tag{13}$$

Note that when $Cov_t\left(R_{m,T},R_{j,T}\right)\leq 0$, (13) implies that the NCC must fail (since risk aversion is positive). On the other hand, for "positive beta" assets where $Cov_t\left(R_{m,T},R_{j,T}\right)>0$, (13) can be written as

$$\gamma\left(\mathrm{E}_tR_{m,T}\right)\geq\frac{Var_t\left(R_{j,T}\right)}{Cov_t\left(R_{m,T},R_{j,T}\right)}. \tag{14}$$

Intuitively, this means that the NCC holds for asset j with a positive beta if relative risk aversion is high enough to make the covariance between $R_{j,T}$ and $u'\left(R_{m,T}\right)$ outweigh the covariance between $R_{j,T}$ and itself, i.e., the variance of $R_{j,T}$.As (14) shows, the ratio of the asset return's variance to its covariance with the market return is fundamental for evaluating whether the NCC holds for asset j. It will be convenient to denote this ratio by

$$\delta_{j,t}\equiv\frac{Var_t\left(R_{j,T}\right)}{Cov_t\left(R_{m,T},R_{j,T}\right)}. \tag{15}$$

Thus, (14) is equivalent to $\gamma\left(\mathrm{E}_t R_{m,T}\right) \geq \delta_{j,t}$. The next proposition summarizes the discussion thus far, establishing a sufficient condition for the NCC and for the validity of the lower bound.

Proposition 3 *Up to a first order approximation, the NCC holds for asset j whenever $Cov_t\left(R_{m,T}, R_{j,T}\right) > 0$ and $\gamma\left(\mathrm{E}_t R_{m,T}\right) \geq \delta_{j,t}$. For such assets, $\frac{Var^*(R_{j,T})}{R_{f,t}}$ serves as a lower bound on the asset's expected excess return between time t and T. The NCC fails when $Cov_t\left(R_{m,T}, R_{j,T}\right) \leq 0$.*

Note that when j is the market, $\delta_{j,t} = \delta_{m,t} = \frac{Var_t\left(R_{m,T}\right)}{Cov_t\left(R_{m,T}, R_{m,T}\right)} = 1$. Thus, in this case, Proposition 3 agrees with Martin's Proposition 2. More generally, Proposition 3 provides us with a sufficient condition to check whether any given asset j satisfies the NCC, and so, whether $\frac{Var^*(R_j)}{R_f}$ is a valid lower bound on its expected excess returns. We next study this condition and explain its economic meaning.

2.4 Combined Risk: The Economic Meaning of δ_j

The condition in Proposition 3 says that for the NCC to hold for asset j we need relative risk aversion in the economy to be greater than $\delta_{j,t}$, this is in contrast to Martin's result which asked that risk aversion be uniformly greater than 1 for the lower bound on the market premium to hold. To obtain some intuition for our condition, how it compares to Martin's condition, and the type of assets for which it is likely to hold we now provide two different economic interpretations of $\delta_{j,t}$.

2.4.1 Relation to CAPM Beta and Idiosyncratic Risk

While $\delta_{j,t}$ is time varying, for estimation we will follow the standard approach (e.g., for estimating betas) and assume it is constant for some (short) period of time and denote it by δ_j. Then, δ_j could be easily estimated from time-series data of asset returns and the returns on the market. We denote by $\hat{\delta}_j$ the estimator of δ_j,

$$\hat{\delta}_j = \frac{Var\left(R_{j,t}\right)}{Cov\left(R_{j,t}, R_{m,t}\right)}. \tag{16}$$

To gain further intuition consider the following standard one-factor model

$$R_{j,t} = \alpha_j + \beta_j R_{m,t} + \varepsilon_{j,t}, \tag{17}$$

in which asset j's returns are being regressed over the market returns and $Cov\left(\varepsilon_{j,t}, R_{m,t}\right) = 0$. Considering the OLS estimation of this regression model, the slope is given by

$$\hat{\beta}_j = \frac{Cov\left(R_{j,t}, R_{m,t}\right)}{Var\left(R_{m,t}\right)},$$

and the R-squared is given by

$$\begin{aligned}\rho_j^2 &= \frac{\hat{\beta}_j^2 Var(R_{m,t})}{Var(R_{j,t})} \\ &= \frac{\hat{\beta}_j Cov(R_{j,t}, R_{m,t})}{Var(R_{j,t})} \\ &= \frac{\hat{\beta}_j}{\hat{\delta}_j}.\end{aligned}$$

We conclude that

$$\hat{\delta}_j = \frac{\hat{\beta}_j}{\rho_j^2}. \tag{18}$$

In words, δ_j is estimated as the ratio of the asset's beta and the R-squared from a regression of the asset returns on the market returns. Since $\frac{1}{\rho_j^2}$ is a measure of the asset's idiosyncratic risk, we obtain that δ_j is large for assets with either large systematic risk (beta) and/or large idiosyncratic risk. We thus call δ_j the *combined risk* of asset j, as it accounts for both idiosyncratic and systematic risk of the asset.

One extreme case is when an asset has idiosyncratic risk of zero, i.e., $\rho_j^2 = 1$. In this case $\hat{\delta}_j = \hat{\beta}_j$. Another extreme case would be if idiosyncratic risk is very large and so ρ_j^2 would be (approximately) zero. In this case $\hat{\delta}_j$ diverges to infinity. The typical cases would be somewhere in between, in which $0 < \rho_j^2 < 1$ and hence the asset's combined risk is strictly higher than $\hat{\beta}_j$ as it is inflated by a factor of $\frac{1}{\rho_j^2}$.

Corollary 4 *For any asset j, $\hat{\delta}_j \geq \hat{\beta}_j$ with equality occurring only for assets with zero idiosyncratic risk.*

Since the weighted average of asset betas in the market equals 1 (the market beta), it follows from Corollary 4 that the weighted average of all combined risks will be strictly larger than 1.

According to Proposition 3, for the NCC to hold for asset j we need relative risk aversion to be greater than the combined risk (δ_j). Corollary 4 shows that the NCC for individual assets is typically more demanding than Martin's NCC for the market. Indeed, his condition only required that risk aversion be greater than 1, while our condition requires that risk aversion be greater than δ_j, which is, on average, strictly greater than 1.

2.4.2 Relation to the Reverse CAPM Regression and Estimation Bias

More insight into the meaning of δ_j can be obtained by considering the reverse CAPM regression in which market returns are regressed on stock returns

$$R_{m,t} = \omega_j + \nu_j R_{j,t} + \eta_{j,t}. \tag{19}$$

Then, from (16),

$$\hat{\delta}_j = \frac{1}{\hat{v}_j}. \tag{20}$$

Thus, another interpretation of δ_j is as the reciprocal of the slope coefficient from a reverse CAPM regression. This interpretation is important for practical estimation since it has implications for the bias of the estimator.[9] To see this note that

$$\mathrm{E}\left(\hat{\delta}_j\right) = \mathrm{E}\left(\frac{1}{\hat{v}_j}\right) \geq \frac{1}{\mathrm{E}\left(\hat{v}_j\right)} = \frac{1}{v_j} = \delta_j,$$

where the first equality follows from (20), the inequality follows from Jensen's inequality, the second equality follows from the fact that OLS coefficients are unbiased, and the last equality follows by taking expectations on both sides of (17) and (19) and then comparing coefficients. We have obtained the following corollary.

Corollary 5 *For any asset j, $\mathrm{E}\left(\hat{\delta}_j\right) \geq \delta_j$. That is, the estimator for δ_j is upward biassed.*

This result is important since it implies that by using $\hat{\delta}_j$ we are actually being conservative. Namely, we require risk aversion to be larger than $\hat{\delta}_j$ while it actually can be somewhat lower.

We next turn to studying the type of stocks for which the NCC holds in the data. These are the stocks for which our lower bound will be valid.

2.5 Testing for the NCC

By Proposition 3, for the NCC to hold we need that $\gamma \geq \delta_j$, where $\delta_j = \frac{Var(R_j)}{Cov(R_j,R_m)}$ is asset $j's$ combined risk and γ is the relative risk aversion in the economy. While δ_j can be easily estimated from data, γ is not directly observable. The finance literature has provided a wide range of reasonable values for relative risk aversion. Bliss and Panigirtzoglou (2004, Table 7) gather estimates of relative risk aversion from several prior studies. Their table shows estimates anywhere between 0 and 55. Recent studies in asset pricing typically consider relative risk aversion levels between 1 and 10 as being "reasonable." For example, Mehra and Prescott (1985) argue that relative risk aversion should be lower than 10, and Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2011) use a relative risk aversion coefficient of either 7.5 or 10 for their calibrations. Such levels are supported by recent estimates such as in Vissing-Jorgensen and Attanasio (2003), who estimate relative risk aversion between 5 to 10 under realistic assumptions for Epstein-Zin Euler equations. Similarly, Bliss and Panigirtzoglou (2004), when considering power utility, estimate relative risk aversion between 3 to 10 as implied from option prices.

To see whether the condition $\gamma \geq \delta_j$ is empirically reasonable, and so whether the NCC may hold for typical stocks, we begin by calculating δ_j from historical monthly stock returns (obtained from CRSP) for common stocks (share code 10 and 11) for

[9] We thank Kerry Back for highlighting this point to us.

the time period January 1995 to December 2014. We estimate δ_j as the variance of the stock monthly return divided by covariance of its return with the CRSP value weighted returns. We restrict attention to S&P 500 constituents starting from the year 2005 ($N = 692$). We consider this sample since these stocks have relatively active option trading and thus relevant data in the OptionMetrics database. As a result, these stocks can serve as a baseline sample for our lower-bound calculations in Section 4. Panel A of Table 1 reports summary statistics for δ_j calculated over the entire sample period. The mean δ_j is 6.5 and the median is 5.2. Thus, if one believes that risk aversion in the economy is at least 5, then the NCC holds for about 50% of S&P constituents. If one holds the more liberal view that relative risk aversion is at least 10, then the NCC holds for 90% of these stocks. The table also shows that δ_j is negatively correlated with size. Thus, the NCC is more likely to hold for large stocks. Panel B reports parallel results with δ_j being calculated using a 60-months rolling window rather then over the entire sample period. The median here is just 4.3 suggesting that risk aversion larger than this level would imply that the NCC holds for at least 50% of S&P 500 constituents. As before, risk aversion of 10 covers about 90% of the S&P 500 universe.

This discussion demonstrates that checking whether the NCC holds for a particular stock ultimately depends on one's views regarding relative risk aversion in the economy. To fix ideas we divide all S&P 500 constituents at each month τ starting from January 2005 into four groups denoted by *conservative, moderate, liberal, and very liberal* based on their $\hat{\delta}_{j,\tau}$ estimated over a rolling window of 60 months from $\tau - 60$ to $\tau - 1$ as follows

$$\text{stock } j \text{ in month } \tau \text{ is} \begin{cases} \textit{conservative} & \text{if} \quad 1 \leq \hat{\delta}_{j,\tau} < 4 \\ \textit{moderate} & \text{if} \quad 4 \leq \hat{\delta}_{j,\tau} < 7 \\ \textit{liberal} & \text{if} \quad 7 \leq \hat{\delta}_{j,\tau} \leq 10 \\ \textit{very liberal} & \text{if} \quad \hat{\delta}_{j,\tau} > 10 \end{cases}. \tag{21}$$

Our main analysis will be centered around the conservative, moderate, and liberal stocks, for which risk aversion in the range of 1-4, 4-7, or 7-10 respectively guarantees the validity of the lower bound. For very liberal stocks the lower bound is unlikely to be valid, and the sign of the covariance in (7) is likely positive. Thus, instead of a lower bound we will likely be getting an upper bound in this case. In addition as can be seen from Table 1, the very liberal group accounts for just around 10% of S&P 500 constituents.

Table 2 reports summary statistics for $\hat{\delta}_{j,\tau}$ for each of the four groups. The average number of stocks in the conservative group is 189 in a given month. The average and median $\hat{\delta}_{j,\tau}$ in this group is 2.8. The stocks in this group are quite large with a mean market cap of $28.2 billion, average B/M ratio of 45%, and average beta of 1. The average R-squared for regressions against the market in this group is 38%. The moderate group averages 133 stocks per month. The average $\hat{\delta}_{j,\tau}$ in this group is 5.3 and the median is 5.1. The stocks in this group are still quite large with an average market cap of $15.5 billion, average B/M ratio of 51%, average beta of 1.3, and average R-squared of 26%. The liberal group averages 48 stocks per month with a mean $\hat{\delta}_{j,\tau}$ of 8.3 and a median of 8.1. The stocks in this group are quite small with a average market cap of \$12.5 billion, average B/M of 51%, average beta of 1.5, and average R-squared of 18%.

As expected, among these three groups, when combined risk becomes larger both the systematic risk (reflected in beta) and idiosyncratic risk (reflected in the inverse of R-squared) become larger. Simultaneously, firms that belong to the more liberal groups are also smaller in size.

Finally, the very liberal group is the smallest with 46 stocks per month on average. The mean $\hat{\delta}_{j,\tau}$ in this group is 33.8 and the median is 14. Thus, the mean here is strongly influenced by large extreme values as reflected by the very large standard deviation. The average beta in this group is 0.9 and the R-squared is just 7%. Thus, the large combined risk for this group is mostly reflecting a very high level of idiosyncratic risk.

3 Estimating the Lower Bound

To estimate the lower bound on the return of asset j we follow a methodology similar to Martin (2015). The idea is that the risk-neutral variance of a stock return can be readily calculated from option prices with various strike prices. Formally, let $d_{j,t}$ denote the present value of dividends paid between times t and T, and consider a set of call and put European options on asset j with strike prices K ranging from 0 to infinity with the same maturity, T. Denote the prices of these options by $call_{j,t}(K)$ and $put_{j,t}(K)$, respectively. Martin shows that

$$\frac{1}{R_{f,t}} var_t^* R_{j,T} = \frac{2}{S_{j,t}^2} \left[\int_0^{F_{j,t}} put_{j,t}(K)\, dK + \int_{F_{j,t}}^{\infty} call_{j,t}(K)\, dK \right] \tag{22}$$

where

$$F_{j,t} = R_{f,t}(S_{j,t} - d_{j,t}),$$

is the forward price of asset j as of time t for delivery at time T. Note that as long as $d_{j,t}$ is not too large, the integration in (22) is performed using options that are primarily out the money for both the put and the call options.

To obtain a numerical estimate of (22) we consider all put options with strike prices less than or equal to $F_{j,t}$ and call options with strike prices larger than $F_{j,t}$ on asset j. Assume there are N_P such put options and N_C such call options available at a given point in time t for which we would like to estimate (22) with corresponding strike prices $K_1^P < ... < K_{N_P}^P < K_1^C < ... < K_{N_P}^C$. Denote the prices of these options by $put_1, ..., put_{N_p}$ and $call_1, ..., call_{N_C}$ respectively. Then, our numerical estimate for the bracketed integrals in (22) is

$$\sum_{i=1}^{N_P-1} put_i \left(K_{i+1}^P - K_i^P\right) + \sum_{i=1}^{N_C-1} call_{i+1} \left(K_{i+1}^C - K_i^C\right) + \left(K_1^C - K_{N_P}^P\right) \min\left(put_{N_P}, call_1\right). \tag{23}$$

Figure 1 illustrates this numerical integration. Note that our numerical approach estimates the integral from below by using the minimum price at each interval, which is consistent with our goal to obtain a lower bound. This approach is somewhat different from Martin's, who uses the price at the mid-point of each interval and then relies on the convexity of the option prices to justify the validity of the lower bound. Our more conservative approach is mandated by the option data for individual stocks, which often does not offer equally spaced strike prices, and thus does not allow us to rely on a similar convexity argument. Note also that dropping the left and right tails from the integral in (22) (i.e., starting the integration at K_1^P and ending at $K_{N_C}^C$) again has the effect of lowering the integral, in line with obtaining a valid lower bound.

A limitation unique to our setting of dealing with individual assets is the fact that options for individual stocks in U.S. markets are all of the American style, while the correct implementation of (22) requires European style options. Since the prices of the American options include an early exercise premium, using American option instead of European option tends to overestimate the lower bound. It is known, however, that the value of the early exercise premium in out-of-the-money options is small (see Barone-Adesi and Whaley, 1987). In Section 5 we provide estimates for the positive bias introduced by using American instead of European options and show that this bias is indeed typically small.

We next turn to discuss how we obtained our lower bound estimates in more detail.

4 Empirical Analysis

4.1 Data and Main Calculations

To calculate (22) we obtain option price data from OptionMetrics. An observation in this database consists of the closing bid and ask prices for an option on a given date for a given stock. Typically, a stock will have multiple options traded on it with different strike prices and maturities.

It is well known that options for individual stocks may be quite illiquid. Thus, we limit our attention to a subset of stocks which are known to be highly traded and liquid. In particular, we focus our attention on options written on stocks which were included in the S&P 500 index starting from 2005. Our sample period spans January 2005 to August 2014. We choose to begin our sample period in 2005 since before that year option trading on individual stocks (as documented on OptionMetrics) has been quite thin. Our focus is on including in our sample stocks that have a rich and liquid set of options, which will allow us to calculate the integrals in (22) as precisely as possible. We have no intention to be fully comprehensive in our sample and are inclined to drop stocks for which option trading is too thin rather than including them in the analysis.

Our aim is to calculate the lower bound (22) for each S&P 500 constituent and each day during our sample period. This goal, however, is too ambitious since even among S&P 500 constituents we find many stocks that have rather thinly traded options or that the number of options is too small to estimate the integrals. Thus, we use the following criteria to further screen illiquid options (similar or related filters have been used in Figlewski, 2008 and Benzoni, Collin-Dufresne, and Goldstein, 2011):

- We drop options with maturity horizon of less than 30 days, as these options are often thinly traded.
- We drop observations that have missing values in their bid and/or ask prices.
- We do not estimate the integrals in (22) for a given day and stock if the total volume of options with the same maturity is less than 20 contracts. We do not feel comfortable trusting the prices of such thinly traded options.[10]
- We do not calculate the integrals in (22) for a given day and stock if the lowest strike to closing price ratio is greater than 0.8 or the highest strike to closing price ratio is less than 1.2. In these cases the left and/or right tails are very large and so we are not likely to obtain a good estimate.
- We do not calculate the integrals for a given day and stock if the number of distinct options for a given maturity is less than 20, as the grid would become too coarse to obtain a viable estimate.
- We do not calculate the integrals for a given day and stock if the maximum distance between two adjacent strike prices is greater than the maximum of 20% of the closing stock price and 10. Again, this would imply a very coarse grid.
- We drop all options with a non-standard settlement.

After applying all these screens we are still left with 54,485,878 observations which yield 575,197 stock/day combinations for which we can estimate (22) to obtain the lower bounds on expected excess returns of 652 distinct stocks. Thus, while our screens our quite aggressive to ensure that our estimates are meaningful, we are still left with a healthy cross-section of stocks allowing us to perform a variety of asset pricing tests.

For each stock/day combination in this sample we estimate the lower bound (22) using (23) for each available maturity separately and then annualize it. For these calculations we take $S_{j,t}$, as the closing price of the underlying security, and the put and call prices used for the integration as the average of the closing bid and ask prices of these options. We then estimate the lower bound for a given day as the average of the lower bounds across different maturities for that day.[11]

[10] We have repeated all the analysis with a screen of at least 50 contracts traded per day. This results in about 10% drop in the number of daily observations but no material effect on any of our results.

[11] An alternative approach would be to take the maximum lower bound across different maturities (as this is the most binding of the lower bounds). We have tried this approach as well and it yields similar results.

Besides the OptionMetrics data we also draw data from CRSP and Compustat to calculate firm characteristics such as beta, size, book-to-market ratio and momentum. These characteristics are used in our cross-sectional Fama-MacBeth regressions.

4.2 Example: Microsoft, P&G, and Apple

Before moving to a full-scale empirical analysis we begin by illustrating the lower bound for three household names: Microsoft, P&G, and Apple. Microsoft and P&G typically belong to the conservative group with an average combined risk of 2.99 and 3.57, respectively. Apple typically belongs to the moderate group with average combined risk of 5.11.

Figure 2 presents the time series of the lower bound for these three stocks between January 2005 and August 2014. Consider Microsoft first. The average lower bound is 6%. A clear pattern is that the lower bound spiked very significantly during the years of the financial crisis. Excluding the years 2008-2009 the average lower bound on expected excess returns for Microsoft is 4.5%. Considering P&G, the average lower bound is 3% and it also exhibits a sharp increase during the crisis period. Excluding this period the average lower bound is 2.3%. For Apple, average lower bound is much larger at 14%. It decreases to 12% once the crisis period is excluded.

Apart from the crisis period, the expected returns of both Microsoft and P&G appear to be quite stable over time. By contrast, Apple's expected excess returns appear to be trending down somewhat from around 18% in 2005 (the iPod and Steve Job's era) to about 10% (the iPhone 6 and Tim Cook's era), with the crisis period interrupting this trend.

4.3 Summary Statistics for the Lower Bounds

We now turn to analyzing the lower bound for all stocks in our sample. Figure 3 and Table 3 provide summary statistics for different sub-samples. Consider first the stocks in the conservative group (for which $1 \leq \delta_{j,\tau} < 4$). The average lower bound for these stocks is 8.4% and the median is 6.1%. Figure 3 also clearly shows that the expected return for these stocks spiked significantly during the crisis years of 2008-2009, in line with the pattern observed for the three household names discussed above. Excluding these two years shows that during "normal times" the average lower bound on these stocks is 6.7% and the median is 5.5%. Note that the distribution of lower bounds is right skewed, which is driven by some very large occasional values. Turning to the moderate group ($4 \leq \delta_{j,\tau} < 7$), the average lower bound is 16.0% and the median is 11.0% for the whole sample period. Excluding 2008-2009 we have an average lower bound of 12.0% and a median of 9.6%. Evidently, the lower bounds of the expected excess returns in this group are higher than in the conservative group. Considering the liberal group ($7 \leq \delta_{j,\tau} \leq 10$) we obtain an even higher lower bound, which spikes occasionally not only during the financial crisis but also in later years. This noise may be attributed to the fact that this group is very small with only 29.4 stocks on average. The average lower bound in this group is 19.3% and the median is 12.9%. Excluding the crisis period we have a mean of 15.2% and a median of 11.1%.

Thus, as the combined risk ($\delta_{j,\tau}$) becomes larger, the lower bound we obtain rises as well. This can stem from two different sources. First, as $\delta_{j,\tau}$ becomes larger the lower bound becomes tighter and (if the risk aversion is sufficiently low) may even switch to becoming an upper bound. Second, a higher combined risk corresponds to higher systematic and/or idiosyncratic risk. And, as shown in Panel B of Table 2, higher $\delta_{j,\tau}$ corresponds to higher beta and smaller size. To the extent that these two risks are compensated for in expected returns, we should expect a higher lower bound.

Considering the very liberal group ($\delta_{j,\tau} > 10$) we again see a rather volatile pattern of lower bounds, which may be attributed to the small number of stocks in this group (21.8 on average). Recall that for this group it is quite likely that what we are calculating is an upper bound rather than a lower bound on expected excess returns. In line with this view, the average bound we obtain is 19.5% and the median is 13.8%. Excluding 2008-2009 we still have a quite high mean of 14.2% and median of 10.7%.

It is interesting to compare these results to the average over the entire sample (Panel E of Figure 3) and for the S&P 500 (Panel F of Figure 3). The latter is simply our version of Martin's (2015) calculations obtained using options on the S&P 500 index (compared this to Figure 3 in Martin (2015)) . The former is the average of the lower bounds for all stocks in our sample considering all the four groups lumped together. The average lower bound for the market premium we obtain is 4.6% and the median is 3.7%. These numbers are comparable to those obtained by Martin (2015) over a larger sample period. The average lower bound across all stocks in our sample is 12.6% and the median is 8.3%. The fact that the average lower bound is higher than the lower bound for the market is not a coincidence. The reason for this a that our lower bounds are convex functions of returns, and thus the lower bound for a portfolio is strictly lower than the weighted average of lower bounds. To see this point formally note that by the convexity of the variance operator

$$Var^{*}\left(\sum_{j=1}^{N} w_j R_j\right) < \sum_{j=1}^{N} w_j Var^{*}\left(R_j\right),$$

implying that the lower bound for the market is strictly lower than the weighted average of lower bounds for individual stocks. Thus, it would be inappropriate to estimate Martin's lower bound for the market premium by averaging our lower bounds for individual stocks.

4.4 Cross-Sectional Analysis of the Bound

Having documented the time-series summary statistics of the lower bounds we now turn to study how they vary in the cross-section of stocks. We hypothesize that firm characteristics that have been documented to affect average realized returns will be reflected also in the forward-looking lower bounds on expected returns we study. The usual suspects are of course beta, size, and book-to-market (Fama and French (1992)) as well as momentum (Jegadeesh and Titman (1993)).

Our goal is to perform a standard Fama-MacBeth (1973) cross-sectional analysis on a monthly basis in which the estimated bound on expected excess returns serves as a left-hand side variable. Thus, we perform a standard analysis replacing average historical returns with forward looking lower bounds of expected returns. To this end, we first calculate for each stock in our sample an average monthly lower bound by averaging the daily estimates of the lower bound for each month during our sample period. We then run cross-sectional regressions of the monthly lower bound against beta, size, book-to-market ratio, and momentum.[12] We perform this analysis for the entire sample as well as separately for each of the four groups defined in (21).

Table 4 reports the time-series averages of the cross-sectional coefficient estimates along with Newey-West adjusted standard errors. The results are striking. To begin, consider the entire sample. The first result is that beta has a highly significant and positive coefficient. That is, firms with high beta are associated with a high lower bound on their expected returns in line with the classic CAPM. Thus, when considering forward-looking expected returns it appears that beta is getting its life back after being "announced dead" in Fama and French (1992). Second, and just as important, size obtains its expected signs as in Fama and French (1992). Indeed, the coefficient of size is negative and highly significant. By contrast, the book-to-market ratio is not significant unlike in Fama and French (1992). Finally and somewhat surprisingly, the coefficient on momentum is negative and significant, suggesting that stocks that experienced a run-up in their price in the past 12 month are associated with lower forward looking expectations.

When considering the conservative, moderate, and liberal groups separately we obtain quite similar results. The coefficient of beta is positive and significant and the coefficient of size is negative and significant for all three groups. The coefficient of book-to-market is not significant for any of the groups. Momentum is negative and significant for the moderate and liberal groups but not for the conservative group. Finally, when considering the very liberal group – none of the coefficients is significant, apparently reflecting the large noise in the estimation of the bound for this group and its small size.

Overall, the cross-sectional results suggest that beta, size, and momentum are reflected in market expectations for individual assets. The result for beta is particularly important since it comes up as a major determinant of expected stock returns unlike in the studies using realized returns. This is consistent with previous studies using other approaches for estimating forward looking returns as in Berk and van Binsbergen (2015) and Brav, Lehavy, and Michaely (2005). As for recent returns, they still play a significant role, but rather than expecting continuation, it appears that investors are expecting reversals in the short term.

[12] We estimate beta by regressing the stock realized excess returns on the market excess return over a rolling window of 60 months prior to the month of interest. We estimate firm size as the log of the year-end market-cap for the fiscal year that preceded the month of interest and book-to-market is the log of the ratio of the book-value of equity to market-cap as of the end of the preceding fiscal year. Finally, momentum is estimated as the stock return during the 12 months preceding the month of interest.

4.5 Predictive Value of the Bound

In the next set of analyses we ask if the lower bound on expected excess returns provides a valuable signal on realized future returns. Finding such a predictive value is a challenging task for at least three reasons. First, our sample period of 10 years is quite short for a longitudinal asset pricing study especially given the likely noisiness of our estimates. Thus, our ability to identify predictability, to the extent it exists, is limited. Second, it may be that the market expectations reflected in option prices are systematically wrong or irrational, and so they (or lower bounds thereof) are not providing a valuable signal about future returns. Third, it may well be that the lower bound we obtain is far from being binding. Thus, variations in the lower bound may not tell us much about variations in realized future returns.

Despite these challenges, we do find evidence of predictability as we describe below. This evidence is especially strong over 6-12 months investment horizons and in the moderate and liberal groups. This is perhaps to be expected as the moderate and liberal groups are the ones for which the NCC is likely to bind (relevant risk aversion of 4-10). Therefore, it should be expected that variations in the lower bound for these groups would be more reflective of variations in future returns.

To evaluate whether the lower bound delivers an informative signal about future returns, in each month t we sort all stocks in our sample based on their monthly average lower bound. We then divide the stocks in each month into ten deciles based on their average bound. Decile 1 consists of the stocks with the lowest estimates and Decile 10 with the highest estimates. We then calculate the equal weighted average of stock realized returns for each decile over three different horizons: one month ($t+1$), six months ($t+1$ to $t+6$), and 12 months ($t+1$ to $t+12$). If the lower bound provides an informative signal, then we expect stocks in lower deciles to show lower average realized returns compared to stocks in the higher deciles.

The results for this analysis are documented in Table 5 for the entire sample and for each of the four groups separately. Consider first Panel B, which reports the time-series averages of the six-month returns for each decile as well as the returns of a portfolio which is long in Decile 10 and short it Decile 1. Considering the entire sample (first column) we see a quite monotone pattern. The average six-month returns in the low decile is 3.73% as compared to 15.16% in the top decile. The difference of 11.43% is very large economically and significant at the 10% level. Similar results are obtained for each of the sub-groups. For example, the moderate group (column 3) shows an increase in average realized returns of 5.66% in the low decile to 20.09% in the top decile, a difference of 14.43% over a six months period.[13]

In Panel C we consider a 12-month investment horizon. The results here are even more striking. The entire sample (column 1) shows an almost perfect monotone trend. Indeed, average returns for deciles 1 through 5 are all below 15% while for deciles 6 through 10 they are all above 15%. The average return in the low decile equals 8.57% and it rises to 36.04% in the top decile – a difference of 27.47% over 12 months (significant at the 10% level). A similar pattern exists in each of the individual groups

[13] Note that we use Newey-West standard errors with three lags in this analysis to account for the potential autocorrelations resulting from the overlapping windows of return estimation.

with statistical significance being strongest in the moderate and liberal groups. For example, in the liberal group, the average return in the low decile is 7.7% compared to 39.6% in the top decile. This difference is very large economically.

In Panel A, which reports the one-month returns, we see weaker results. The trends in realized returns are still somewhat monotone and the differences between the two extreme deciles are still always positive. However, they are not statistically significant. For example, for the entire sample the difference between average return in the top and bottom deciles is 1.34% per month.

To summarize, the results in Table 5 provide evidence in support of a predictive signal incorporated in the lower bound on expected stock returns. The evidence is economically large and statistically significant for an investment horizon of 6-12 months. While we do see evidence of predictability in all groups, the evidence is more pronounced within stocks in the moderate and liberal groups for which the lower bounds are more likely to be binding.

5 Precision of Approximations

Our analysis relies on two types of approximations. First, to study whether the NCC holds for a particular asset we rely on a first-order Taylor approximation. Second, when estimating the lower bound we do not account for the fact that the options being used are American. In this section we estimate the error associated with these approximations. Our overall conclusion is that in most cases this error is small or even negligible.

5.1 First Order Approximation for the NCC

A key advantage of our approach is that the form of the utility function need not be known in order to check whether the NCC holds or to estimate the lower bound for a particular asset. Instead, all that is needed is an assumption on the relative risk aversion (which may not be constant) in the economy and its magnitude relative to the combined risk of the asset. However, in order to obtain this "utility irrelevance" we replaced $u'(R_{m,t}) R_{j,t}$ with its first order Taylor approximation in the calculation of $Cov(u'(R_{m,t}) R_{j,t}, R_{j,t})$. Similar approximations are rather common in asset pricing models attempting to linearize or log-linearize non-linear expressions. In order to assess the error associated with this approximation we consider standard utility functions, which allow us to estimate $Cov(u'(R_{m,t}) R_{j,t}, R_{j,t})$ precisely without resorting to an approximation. We can then test how often the first order approximation leads us to an incorrect inference about the NCC. Specifically, we assume that u takes the CRRA form $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$, with γ varying between 2 to 10. We then repeat the analysis in Section 2.5, but instead of using the condition $\gamma \geq \delta_{j,\tau}$ (relying the approximation) we check directly whether $cov(u'(R_{m,t}) R_{j,t}, R_{j,t}) \leq 0$ holds true.

Table 6 reports the type I (false positive) and type II (false negative) errors associated with using the approximation. Consider for example the case $\gamma = 3$. The table shows that using the approximation $\gamma \geq \delta_{j,\tau}$ implies a probability of 0.76%

of concluding that the NCC holds while in fact it does not (type I error). At the same time, the probability of concluding that the NCC fails while it actually holds is about 11.8% (type II error). Thus, in this case, the approximation seems rather conservative in line with Corollary 5. It rarely leads us to accept the NCC when it fails, but it occasionally leads us to conclude that the NCC fails when it actually holds. This conservatism seems somewhat desirable as it prevents us from relying on the lower bound in cases where it is not valid.

As γ grows larger the probability of type I errors increases and that of type II errors declines sharply. For example, when $\gamma = 9$ the probability of type I errors is about 15% and the probability of type II errors is 3.4%. Overall, for levels of risk aversion up to 8 it appears that our approximation is quite conservative, with type I error probabilities less than 10%. For higher levels of risk aversion type I errors become more likely while the likelihood of type II errors diminishes.

5.2 Using American Options to Calculate the Lower Bound

The estimation of the lower bound in (22) relies on the prices of European options. However, all options on individual stocks in the U.S. are of the American style, introducing a potential upward bias in our lower bound estimation due to the early exercise premium (EEP). It is important to note that the options we are using are mostly out-of-the-money, a case in which the EEP is known to be relatively small. Still, in this section we evaluate the magnitude of this potential bias.

A key advantage of (22) is that it makes no assumptions regarding the underlying framework. To estimate the EEP we need to make additional assumptions on the dynamics of the underlying security prices. To this end, we follow the framework in MacMillan (1986) and Barone-Adesi and Whaley (1987), who offer an analytic approximation for the EEP of American options in the Black-Scholes framework. Specifically, we calculate the lower bound in (22) by obtaining option prices from OptionMetrics and subtracting from them the estimated EEP to obtain a synthetic price of a corresponding European option.

The results are reported in Table 7, which is analogous to Table 3, with the only difference being the use of the synthetic European option prices instead of the American options prices. As expected, the lower bounds in Table 7 are all lower than in Table 3, but the differences are typically small. To illustrate, the median lower bound for the entire sample in Panel A of Table 3 is 8.28% as opposed to 7.96% in Panel A of Table 7, a difference of 32 basis points annually. When considering the conservative group the median is 6.07% as opposed to 5.86%, a difference of 21 basis points. Similar differences apply to the three other groups. The differences are even smaller when considering Panel B of Tables 3 and 7, which exclude the crisis period.

When considering the mean of the lower bound distribution, the differences for the conservative groups are also small at 39 basis points. The differences in means become larger for the other three groups. The deviation between the mean and the median in that regard is a result of what looks like outliers in the right tail, specifying lower bounds close to or higher than 100%.

In conclusion, our estimates suggest that the use of American options inflates the lower bound by less than 50 basis points for the conservative group and by about 1-3 percentage points for the other groups. As a robustness test we also repeated the empirical analysis discussed in Sections 4.4 and 4.5 replacing the original lower bounds by the modified lower bounds obtained using the synthetic European options. None of the conclusions is affected.

6 Conclusion

To be completed.

Appendix I

In this appendix we derive approximations and conditions for when an asset satisfies NCC for two dynamic models. First is a standard consumption/investment model with separable utility in which risk aversion may not be constant. Second is a dynamic model with a recursive Epstein-Zin utility.

A Dynamic Model with Separable Utility

Consider a representative investor with time-separable utility $u(\cdot)$, which is increasing and concave and a subjective discount factor $\beta \in (0,1)$. The investor faces $N > 1$ assets with random return $R_{i,t+1}$ between time t and $t+1$. In each period $0, 1, 2, ...$ the investor needs to allocate his initial wealth W_t among consumption C_t and investment in each asset i, $w_{i,t}$.

Assume there is an inter-temporal representative investor with separable utility. Her value function $J(\cdot)$ is defined recursively as a function of her wealth W_t as follows

$$\begin{aligned} J(W_t) &= \max_{C_t, \{w_{i,t}\}} \left[u(C_t) + \beta \mathrm{E}_t J\left((W_t - C_t) \sum_{i=1}^{N} w_{i,t} R_{i,t+1} \right) \right] \\ s.t. \sum_{i=1}^{N} w_{i,t} &= 1. \end{aligned}$$

The first-order condition for $w_{i,t}$ is

$$\beta \mathrm{E}_t J' \left((W_t - C_t) \sum_{i=1}^{N} w_{i,t} R_{i,t+1} \right) (W_t - C_t) R_{i,t+1} = \lambda,$$

where λ is a positive multiplier. It follows that

$$M_{t+1} \equiv \frac{\beta J' \left((W_t - C_t) \sum_{i=1}^{N} w_{i,t} R_{i,t+1} \right) (W_t - C_t)}{\lambda}$$

is a time-t stochastic discount factor (pricing claims between time t and $t+1$). Since $W_{t+1} = (W_t - C_t) \sum w_i R_{i,t+1}$ we have

$$M_{t+1} = \frac{\beta J'(W_{t+1})(W_t - C_t)}{\lambda}.$$

Since the representative investor holds the market the portfolio $(w_{1,t},...,w_{N,t})$ is the market portfolio as of time t, and we denote the return on this portfolio between times t and $t+1$ by

$$R_{m,t+1} = \sum_i w_{i,t} R_{i,t+1}.$$

Thus, $W_{t+1} = (W_t - C_t) R_{m,t+1}$. This shows that M_{t+1} is proportional to $J'((W_t - C_t) R_{m,t+1})$. Thus, the NCC holds for asset j at time t if and only if $Cov_t (J'((W_t - C_t) R_{m,t+1}) R_{j,t+1}, R_{j,t+1}) \leq 0$.

Define $f : \mathbb{R}^N \to \mathbb{R}$ by

$$f(R_{1,t+1}, R_{2,t+1}, ...R_{N,t+1}) = R_{j,t+1} J'((W_t - C_t) \sum_{i=1}^{N} w_{i,t} R_{i,t+1}). \tag{24}$$

An analysis parallel to that in Section 2.3 yields the following first order approximation

$$Cov_t \left(J'((W_t - C_t) R_{m,t+1}) R_{j,t+1}, R_{j,t+1} \right) \approx J'(\mathrm{E}_t W_{t+1}) \left(Var_t (R_{j,t+1}) - \Gamma(\mathrm{E}_t W_{t+1}) Cov_t (R_{m,t+1}, R_{j,t+1}) \right),$$

where $\Gamma(\mathrm{E}_t W_{t+1}) = -\mathrm{E}_t W_{t+1} \frac{J''(\mathrm{E}_t W_{t+1})}{J'(\mathrm{E}_t W_{t+1})}$ is the investor's relative risk aversion with respect to his life-time utility evaluated at $\mathrm{E}_t W_{t+1}$.

Thus, similar to the conclusion in Section 2.3 we have that the NCC holds for asset j at time t if $\delta_{j,t} = \frac{Var_t(R_{j,t+1})}{Cov_t(R_{m,t+1}, R_{j,t+1})}$ is lower than relative risk aversion.

Proposition 6 *Up to a first order approximation, the NCC holds for asset j whenever $\delta_{j,t} \leq \Gamma(\mathrm{E}_t W_{t+1})$. For such assets, $\frac{var^*(R_{j,t+1})}{R_{f,t}}$ serves as a lower bound on the asset's expected excess return between times t and $t+1$.*

Note that we have obtained this result in a traditional consumption/investment framework in which returns are exogenous while consumption and investment are endogenous. A similar result can be obtained in an endowment economy, where consumption is assumed exogenous and prices are determined endogenously.

A Dynamic Consumption/Investment Model with Recursive Utility

Consider an infinitely lived representative investor with recursive value function

$$V_t = J(C_t, \mu(V_{t+1})),$$

where the function J is an aggregator mapping current consumption C_t, and the certainty equivalent of future life time value, $\mu(V_{t+1})$, to current value, V_t. We follow Epstein and Zin (1989) and consider the following functional form

$$J(C, \mu) = \left[(1-\delta) C^{1-\rho} + \delta \mu^{1-\rho} \right]^{\frac{1}{1-\rho}}, \qquad \rho \geq 0$$

and the certainty equivalent function is

$$\mu\left(V_{t+1}\right)=\left[E_t\left(V_{t+1}^{1-\gamma}\right)\right]^{\frac{1}{1-\gamma}}, \qquad \gamma>0.$$

In the above expressions, $\rho=\frac{1}{\psi}$, where ψ is the inter-temporal elasticity of substitution (IES), γ is the relative risk aversion, and δ is a subjective discount factor.

The investor faces $N>1$ assets with random return $R_{i,t+1}$ between time t and $t+1$. In each period the investor needs to allocate his initial wealth W_t among consumption C_t and investment weight in each asset i, $w_{i,t}$ to maximize his life time value subject to the budget constraint

$$\left(W_t-C_t\right)R_{m,t+1}=W_{t+1},$$

where $R_{m,t+1}=\sum_1^N w_{i,t}R_{i,t+1}$. We further assume that market returns are i.i.d.

The stochastic discount factor in this economy takes the form

$$M_{t+1}=\delta^{\theta}\left(\frac{C_{t+1}}{C_t}\right)^{\frac{-\theta}{\psi}}R_{m,t+1}^{\theta-1},$$

where

$$\theta\equiv\frac{1-\gamma}{1-\rho}.$$

Since C_t is known as of time t, we have that M_{t+1} is proportional to $C_{t+1}^{-\theta/\psi}R_{m,t+1}^{\theta-1}$. It follows that $M_{t+1}R_{j,t+1}$ is proportional to

$$C_{t+1}^{-\theta/\psi}R_{m,t+1}^{\theta-1}R_{j,t+1}=\left(\frac{C_{t+1}}{W_{t+1}}\right)^{-\theta/\psi}W_{t+1}^{-\theta/\psi}R_{m,t+1}^{\theta-1}R_{j,t+1}.$$

Since market returns are i.i.d., the consumption to wealth ratio is constant. It follows $M_{t+1}R_{j,t+1}$ is proportional to $W_{t+1}^{-\theta/\psi}R_{m,t+1}^{\theta-1}R_{j,t+1}$. Moreover, from the budget constraint

$$W_{t+1}=(W_t-C_t)R_{m,t+1}=W_t(1-\frac{C_t}{W_t})R_{m,t+1},$$

and since both W_t and $\frac{C_t}{W_t}$ are known by time t we have that $M_{t+1}R_{j,t+1}$ is proportional to

$$\left(\sum_i w_iR_{i,t+1}\right)^{\theta-1-\theta/\psi}R_{j,t+1}=\left(\sum_i w_iR_{i,t+1}\right)^{-\gamma}R_{j,t+1}.$$

Thus, $Cov_t(M_{t+1}R_{j,t+1},R_{j,t+1})$ has the same sign as

$$Cov_t\left(\left(\sum_i w_iR_{i,t+1}\right)^{-\gamma}R_{j,t+1},R_{j,t+1}\right).$$

Define $f:\mathbb{R}^N\to\mathbb{R}$ by

$$f(R_{1,t+1},R_{2,t+1},...R_{N,t+1})=\left(\sum_i w_iR_{i,t+1}\right)^{-\gamma}R_{j,t+1}.$$

An analysis parallel to that in Section 2.3 yields the following first order approximation

$$Cov_t(\left(\sum_i w_i R_{i,t+1}\right)^{-\gamma} R_{j,t+1}, R_{j,t+1}) \approx (\mathrm{E}_t R_{m,t+1})^{-\gamma} (Var_t(R_{j,t+1}) - \gamma Cov_t(R_{m,t+1}, R_{j,t+1}))\,.$$

Thus, up to a first order approximation, $Cov_t(M_{t+1}R_{j,t+1}, R_{j,t+1})$ is non-positive if and only if $\delta_{j,t} \leq \gamma$ as before.

Proposition 7 *Up to a first order approximation, the NCC holds for asset j whenever $\delta_{j,t} \leq \gamma$. For such assets, $\frac{var^*(R_{j,t+1})}{R_{f,t}}$ serves as a lower bound on the asset's expected excess return between times t and $t+1$.*

Table 1 Summary Statistics for δ_j

This table reports summary statistics for $\delta_j = \frac{Var(R_j)}{Cov(R_j, R_m)}$ calculated using return data of CRSP common stocks (share code 10 and 11) included in the S&P index starting from the year 2005 until the end of 2014. In Panel A, δ_j is calculated based on monthly stock returns between January 1995 and December 2014. Beta is the regression coefficient from time-series regressions of the stock excess return on CRSP value weighted excess return. Size is the average of log market-cap over the sample period. B/M is the average over the sample period of the log of the ratio of equity book value to market value at the end of the fiscal year obtained from Compustat. In Panel B, δ_j and beta are calculated using a 60-month rolling window. P-values are reported in parentheses below correlation estimates.

Panel A: Delta Calculation Based on the Entire Sample Period

Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
692	6.4615	8.4451	2.3093	2.8901	3.1465	3.9307	5.2051	7.1486	10.0060	12.6492	22.2188
Corr with		Beta			Size			B/M			
		-0.0703 (0.0645)*			-0.1122 (0.0031)***			-0.0557 (0.1436)			

Panel B: Delta Calculation Based on a 60-Month Rolling Window

Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
55332	8.12	110.22	1.5873	1.9266	2.2452	2.9413	4.2990	6.6650	10.5213	14.7998	40.0904
Corr with		Beta			Size			B/M			
		-0.0380 (0.0000)***			-0.0069 (0.1478)			-0.0181 (0.0002)***			

Table 2 Summary Statistics for δ_j by Different Groups

This table reports the summary statistics for $\delta_j = \frac{Var(R_j)}{Cov(R_j, R_m)}$ for the conservative, moderate, liberal, and very liberal groups of stocks defined in (21). The calculation of δ_j is based on a 60-month rolling window of monthly stock returns from January 2000 to December 2014, and the market return is taken as the CRSP value weighted return. Obs is the average number of monthly observations in each group. We also report the mean and standard deviations of the firm's beta, size and book-to-market ratio. Beta is the regression coefficient from 60-month rolling window time-series regressions of the stock excess return on CRSP value weighted excess return. Size is the average of market-cap over the sample period measured in billions of dollars. B/M is the average over the sample period of the ratio of equity book value to market value at the end of the fiscal year obtained from Compustat. ρ^2 is the R-squared from 60-month rolling window time-series regressions of the stock excess return over CRSP value weighted excess return.

Panel A: Summary Statistics for δ_j

	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
Conservative	189.08	2.8270	0.6621	1.4723	1.7430	1.8964	2.3224	2.8389	3.3660	3.7293	3.8670	3.9735
Moderate	133.24	5.2513	0.8380	4.0204	4.1019	4.2126	4.5319	5.1254	5.9124	6.5178	6.7599	6.9470
Liberal	48.40	8.2533	0.8395	7.0188	7.1009	7.1979	7.5299	8.1462	8.9130	9.5239	9.7466	9.9436
Very Liberal	46.32	34.7729	328.8407	10.0495	10.2333	10.5094	11.3717	13.9376	20.5537	37.6301	65.2923	300.6199

Panel B: Firm Characteristics

	Beta	Size($ billions)	B/M	ρ^2
	Mean (Std. Dev)	Mean (Std. Dev)	Mean (Std. Dev)	Mean (Std. Dev)
Conservative	1.0401	28.2250	0.4520	0.3766
	(0.4384)	(54.2959)	(0.2769)	(0.1487)
Moderate	1.3395	15.5207	0.5090	0.2616
	(0.7598)	(30.6535)	(0.4346)	(0.1453)
Liberal	1.4713	12.4862	0.5132	0.1833
	(1.0773)	(21.6667)	(0.5097)	(0.1342)
Very Liberal	0.8998	11.6684	0.4319	0.0717
	(0.9511)	(20.4505)	(0.5058)	(0.0825)

Table 3 Summary Statistics for Estimated Lower Bound

This table reports summary statistics for the estimated lower bound obtained from (23). The calculation is based on options with underlying securities included in the S&P index starting from the year 2005 and subject to the filters described in Section 4. We estimate the lower bound for each stock and each day during the sample period of January 2005 to August 2014. The summary statistics are reported for the entire sample and for each of the four groups defined in (21) separately. The table also reports an estimated lower bound for the market premium using option prices on the S&P 500 (SPX). All reported lower bound estimates are annualized.

Panel A: Entire Sample Period

	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
Entire Sample	575197	0.1258	0.1783	0.0165	0.0248	0.0318	0.0495	0.0828	0.1416	0.2392	0.3485	0.7821
Conservative	302058	0.0841	0.1094	0.0157	0.0223	0.0271	0.0393	0.0607	0.0972	0.1521	0.2043	0.4310
Moderate	174645	0.1597	0.2142	0.0172	0.0354	0.0465	0.0699	0.1098	0.1767	0.3036	0.4326	0.8921
Liberal	54964	0.1925	0.2430	0.0208	0.0376	0.0494	0.0802	0.1290	0.2126	0.3605	0.5459	1.1904
Very Liberal	43530	0.1952	0.2265	0.0227	0.0376	0.0505	0.0806	0.1382	0.2211	0.3701	0.5489	1.1444
Market Premium	2424	0.0455	0.0342	0.0170	0.0192	0.0203	0.0232	0.0366	0.0540	0.0791	0.1221	0.1921

Panel B: Sample Period Without Years 2008 and 2009

	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
Entire Sample	444484	0.0938	0.0992	0.0156	0.0231	0.0290	0.0441	0.0706	0.1131	0.1753	0.2290	0.4203
Conservative	253903	0.0671	0.0485	0.0152	0.0215	0.0257	0.0366	0.0549	0.0841	0.1209	0.1523	0.2331
Moderate	126629	0.1201	0.1007	0.0154	0.0310	0.0423	0.0631	0.0960	0.1452	0.2155	0.2857	0.5004
Liberal	37320	0.1523	0.1945	0.0190	0.0328	0.0437	0.0713	0.1106	0.1779	0.2642	0.3539	0.8007
Very Liberal	26632	0.1416	0.1531	0.0202	0.0314	0.0430	0.0660	0.1066	0.1709	0.2487	0.3400	0.7497
Market Premium	1922	0.0349	0.0175	0.0168	0.0188	0.0200	0.0217	0.0277	0.0418	0.0590	0.0729	0.0942

Table 4: Fama MacBeth Analysis of the Bounds

This table presents the results of monthly cross-sectional Fama-MacBeth regressions. The sample period is January 2005 to August 2014. We first estimate a monthly lower bound for each stock by averaging the daily lower bounds reported in Table 3 over the month. For each month we then run cross-section regressions in which the dependent variable is the average monthly lower bound and the independent variables are beta, size, book-to-market ratio, and momentum. Obs is the time-series average of the number of firms in the cross-sectional regressions. Beta is calculated as the coefficient estimate from regressing each stock's monthly excess returns on the market excess returns during a 60-month rolling window. Firm size is the log of the firm's market-cap at the end of previous year. B/M is the log of the ratio of equity book value to market cap at the end of the preceding fiscal year. Momentum is given by the cumulative return over the 12 months period prior the current month. The table reports the time-series average of the cross-sectional coefficient estimates as well as Newey-West standard errors with 4 lags. The analysis is performed both for the entire sample and for each of the four groups defined in (21) separately. Asterisks denote statistical significance at the 1% (***), 5% (**) and 10% (*) level

	Obs	β	Size	B/M Ratio	Momentum
Entire Sample	314.3	0.0498	-0.0203	0.0078	-3.8011
		(0.0085)***	(0.0017)***	(0.0048)	(1.6426)**
Conservative ($1 \leq \delta < 4$)	171.2	0.0493	-0.0051	0.0035	-1.8912
		(0.0083)***	(0.0011)***	(0.0037)	(1.2731)
Moderate ($4 \leq \delta < 7$)	97.8	0.0427	-0.0246	0.0091	-6.8896
		(0.0094)***	(0.0049)***	(0.0064)	(1.8835)***
Liberal ($7 \leq \delta \leq 10$)	29.4	0.0344	-0.0295	0.0132	-5.1119
		(0.0126)***	(0.0067)***	(0.0141)	(2.6330)*
Very Liberal ($\delta > 10$)	21.8	-0.4922	-1.0438	-1.1901	45.7667
		(0.6021)	(0.9620)	(1.0910)	(47.8518)

Table 5: Predictive Value of the Bound

This table reports the average realized returns sorted by the estimated lower bound on expected excess returns. Each month t during our sample period we classify stocks into 10 deciles accourding to the monthly average of their estimated lower bound. The table reports the average of the realized returns in the next month/6 months/12 months for each decile for the entire sample and for each of the four groups defined in (21) separately. The bottom row reports the average returns of a portfolio which is long in Decile 10 and short it Decile 1 along with Newey-West standard errors with 3 lags in parenthesis. Asterisks denote statistical significance at the 1% (***), 5% (**) and 10% (*) level.

Panel A: Average Future Return of Stocks Grouped by Lower Bound Estimates — One-Month Horizon

Decile	Entire Sample	Conservative ($1 \leq \delta < 4$)	Moderate ($4 \leq \delta < 7$)	Liberal ($7 \leq \delta \leq 10$)	Very Liberal ($\delta > 10$)
1	0.0063	0.0081	0.0070	0.0074	-0.0000
2	0.0066	0.0069	0.0124	0.0108	0.0076
3	0.0101	0.0048	0.0119	0.0054	0.0003
4	0.0092	0.0062	0.0125	0.0121	-0.0015
5	0.0086	0.0038	0.0095	0.0135	0.0017
6	0.0087	0.0068	0.0162	0.0112	0.0127
7	0.0107	0.0030	0.0163	0.0329	0.0017
8	0.0103	0.0084	0.0136	0.0189	0.0134
9	0.0160	0.0070	0.0142	0.0200	0.0051
10	0.0197	0.0130	0.0175	0.0338	0.0268
10-1	0.0134	0.0049	0.0104	0.0264	0.0268
	(0.0098)	(0.0110)	(0.0117)	(0.0187)	(0.0241)

Panel B: Average Future Return of Stocks Grouped by Lower Bound Estimates — Six-Month Horizon

Decile	Entire Sample	Conservative ($1 \leq \delta < 4$)	Moderate ($4 \leq \delta < 7$)	Liberal ($7 \leq \delta \leq 10$)	Very Liberal ($\delta > 10$)
1	0.0373	0.0366	0.0566	0.0193	0.0252
2	0.0485	0.0402	0.0707	0.0606	0.0153
3	0.0667	0.0545	0.0774	0.0812	0.0484
4	0.0646	0.0547	0.0804	0.0738	0.0685
5	0.0619	0.0502	0.0766	0.0769	0.0383
6	0.0607	0.0454	0.0917	0.1220	0.0495
7	0.0689	0.0429	0.0959	0.1532	0.0953
8	0.0932	0.0454	0.0753	0.1119	0.0956
9	0.1171	0.0591	0.1161	0.0826	0.1056
10	0.1516	0.0931	0.2009	0.1539	0.1443
10-1	0.1143	0.0565	0.1443	0.1346	0.1191
	(0.0685)*	(0.0633)	(0.0843)*	(0.0724)*	(0.0586)**

Panel C: Average Future Return of Stocks Grouped by Lower Bound Estimates — Twelve-Month Horizon

Decile	Entire Sample	Conservative ($1 \leq \delta < 4$)	Moderate ($4 \leq \delta < 7$)	Liberal ($7 \leq \delta \leq 10$)	Very Liberal ($\delta > 10$)
1	0.0857	0.0771	0.1294	0.0774	0.1201
2	0.1066	0.0855	0.1725	0.1198	0.0468
3	0.1360	0.1119	0.1776	0.1000	0.0609
4	0.1241	0.1192	0.1712	0.1976	0.2084
5	0.1410	0.1105	0.1730	0.1612	0.1071
6	0.1554	0.1167	0.1853	0.2544	0.1463
7	0.1556	0.1043	0.2196	0.3465	0.1988
8	0.1965	0.1052	0.1663	0.2050	0.2066
9	0.2513	0.1197	0.2610	0.1813	0.1769
10	0.3604	0.2457	0.4393	0.3961	0.2011
10-1	0.2747	0.1686	0.3099	0.3187	0.0810
	(0.1513)*	(0.1390)	(0.1667)*	(0.1762)*	(0.0682)

Table 6: Type I and Type II Errors for Taylor Approximations

This table reports the type I (false positive) and type II (false negative) error probabilities resulting from using the first order approximation to test the validity of the NCC. We assume a utility function of the form $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$, where γ ranges between 2 and 10. The NCC holds if $cov(u'(R_m) R_j, R_j) \leq 0$ and the first order approximation holds if $\gamma \geq \delta_j$, where $\delta_j = \frac{Var(R_j)}{Cov(R_j, R_m)}$ calculated using a 60-month rolling window. Our sample period for this test is January 2005 to December 2014. For the market return R_m we use monthly CRSP value weighted returns and for the stock return (R_j) we use monthly return data. The analysis is restricted to S&P 500 constituents starting from the year 2005.

γ	Type I Error	Type II Error
2	0.0018	0.1977
3	0.0076	0.1180
4	0.0159	0.0922
5	0.0299	0.0767
6	0.0522	0.0550
7	0.0695	0.0510
8	0.1050	0.0417
9	0.1499	0.0339
10	0.1954	0.0290

Table 7 Summary Statistics for Estimated Lower Bound Using Synthetic European Option Prices

This table reports summary statistics for the estimated lower bound obtained from (23) using synthetic European option prices calculated following the approximation presented in Barone-Adesi and Whaley (1987). The estimation is based on options with underlying securities included in the S&P index starting from the year 2005 and subject to the filters described in Section 4. We estimate the lower bound for each stock and each day during the sample period of January 2005 to August 2014. The summary statistics are reported for the entire sample and for each of the four groups defined in (21) separately. All reported lower bound estimates are annualized.

Panel A: Entire Sample Period

	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
Entire Sample	575197	0.1155	0.1356	0.0154	0.0239	0.0307	0.0477	0.0796	0.1347	0.2223	0.3153	0.6608
Conservative	302058	0.0802	0.0924	0.0149	0.0216	0.0263	0.0383	0.0589	0.0936	0.1445	0.1943	0.4042
Moderate	174645	0.1449	0.1565	0.0155	0.0328	0.0438	0.0675	0.1055	0.1675	0.2738	0.3840	0.7417
Liberal	54964	0.1652	0.1648	0.0182	0.0346	0.0467	0.0757	0.1196	0.1948	0.3123	0.4526	0.8909
Very Liberal	43530	0.1795	0.1792	0.0205	0.0357	0.0479	0.0767	0.1339	0.2136	0.3461	0.4828	0.9026

Panel B: Sample Period Without Years 2008 and 2009

	Obs	Mean	Std.Dev	1%	5%	10%	25%	50%	75%	90%	95%	99%
Entire Sample	444484	0.0865	0.0700	0.0146	0.0223	0.0280	0.0426	0.0680	0.1081	0.1638	0.2092	0.3434
Conservative	253903	0.0643	0.0423	0.0144	0.0208	0.0250	0.0357	0.0534	0.0811	0.1159	0.1435	0.2158
Moderate	126629	0.1104	0.0782	0.0141	0.0282	0.0397	0.0606	0.0927	0.1387	0.1984	0.2490	0.3959
Liberal	37320	0.1265	0.0988	0.0167	0.0300	0.0410	0.0669	0.1023	0.1589	0.2280	0.2873	0.5075
Very Liberal	26632	0.1284	0.1028	0.0186	0.0299	0.0405	0.0630	0.1013	0.1621	0.2347	0.3099	0.5585

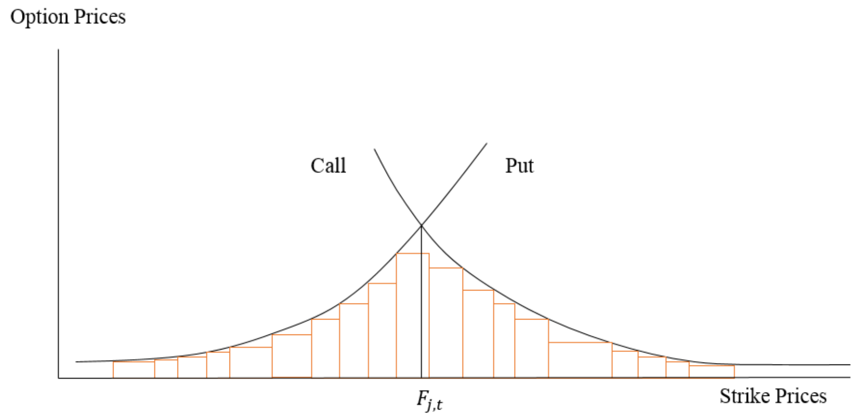

Figure 1: Illustration of the numerical estimation of the lower bound.

Panel A: Estimated Lower Bound for Microsoft

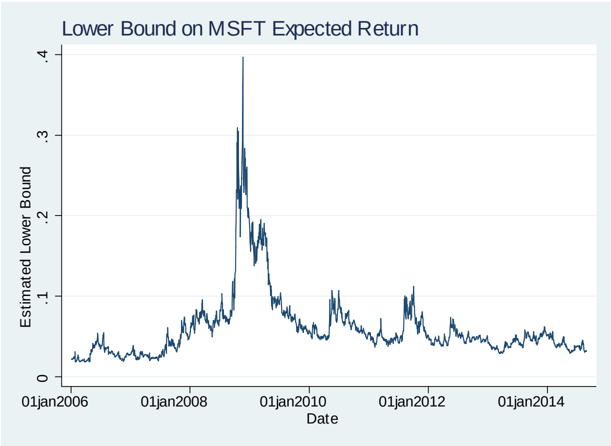

Panel B: Estimated Lower Bound for P&G

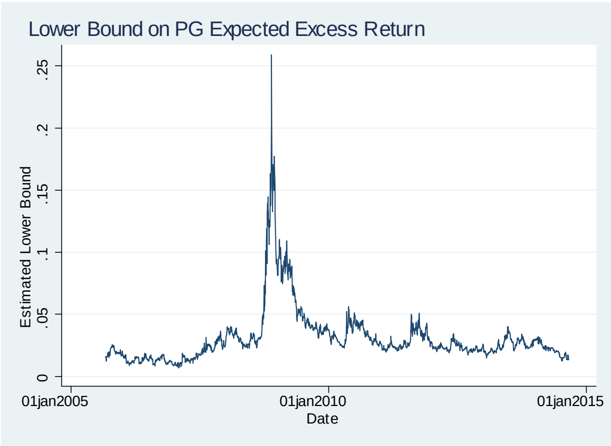

Panel C: Estimated Lower Bound for Apple

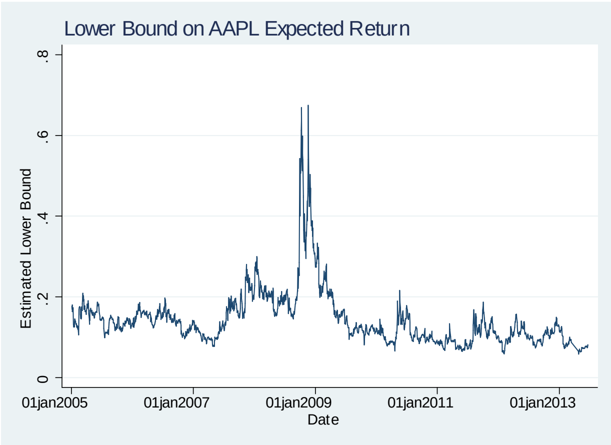

Figure 2: Time series of estimated lower bound for Microsoft, P&G and Apple between January 2005 and August 2014.

Panel A: Conservative Group

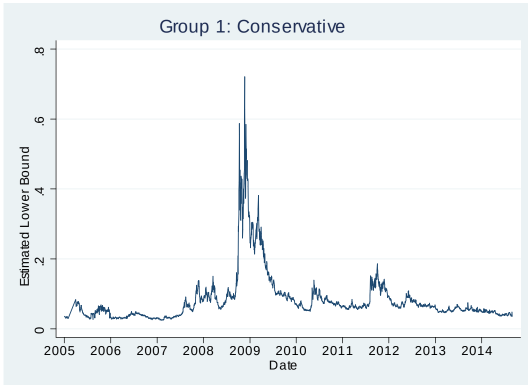

Panel B: Moderate Group

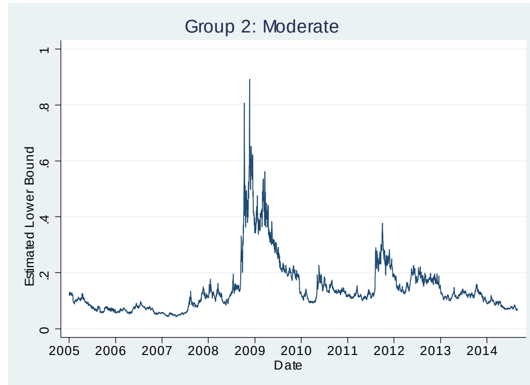

Panel C: Liberal Group

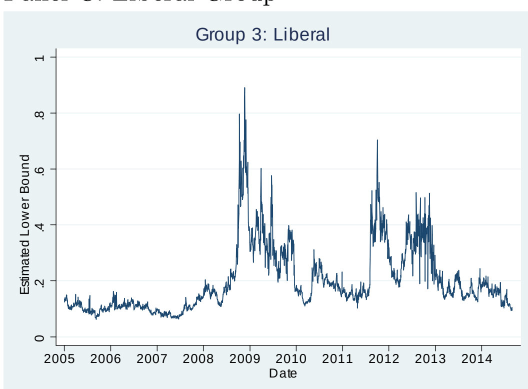

Panel D: Very Liberal Group

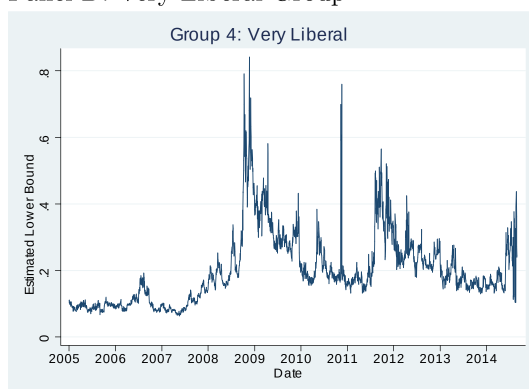

Panel E: Whole Sample

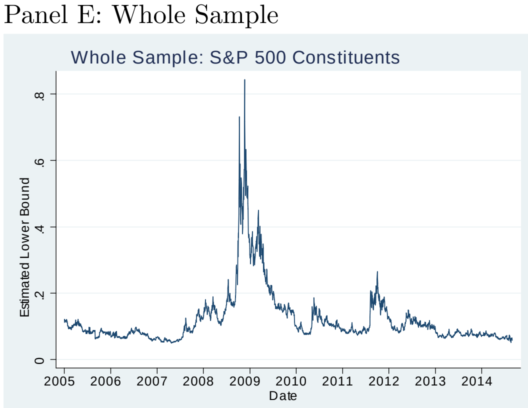

Panel F: Market Premium

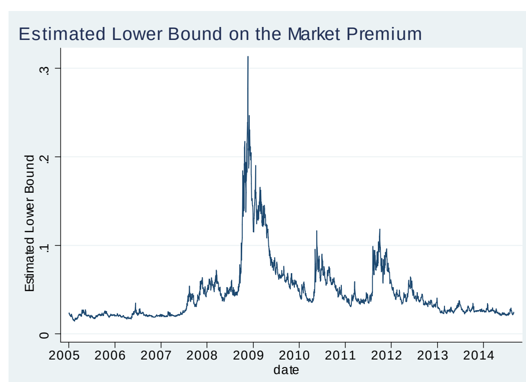

Figure 3: Time series of average estimated lower bound for each group, the whole sample, and the market premium between January 2005 and August 2014.

References

[1] Aıt-Sahalia, Yacine, and Andrew W. Lo, 2000. "Nonparametric risk management and implied risk aversion." Journal of econometrics 94, no. 1: 9-51.

[2] Bansal, Ravi, Dana Kiku, and Amir Yaron, 2009,. "An empirical evaluation of the long-run risks model for asset prices." No. w15504. National Bureau of Economic Research.

[3] Bansal, Ravi, and Amir Yaron, 2004, "Risks for the long run: A potential resolution of asset pricing puzzles." The Journal of Finance 59: 1481-1509.

[4] Barone-Adesi, Giovanni and Robert E. Whaley, 1987, "Efficient analytic approximation of American option values." The Journal of Finance 42.2: 301-320.

[5] Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2011, "Explaining asset pricing puzzles associated with the 1987 market crash." Journal of Financial Economics 101.3: 552-573.

[6] Berk, Jonathan and van Binsbergen, Jules H., 2015, "Assessing Asset Pricing Models Using Revealed Preference." Journal of Financial Economics, forthcoming.

[7] Black, Fischer, and Myron Scholes, 1973. "The pricing of options and corporate liabilities." The Journal of Political Economy: 637-654.

[8] Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, "Option-implied risk aversion estimates." The Journal of finance 59: 407-446.

[9] Borovička, Jaroslav, Lars P. Hansen, and José A. Scheinkman, 2014. "Misspecified recovery." Working paper, National Bureau of Economic Research.

[10] Brav, Alon, Reuven Lehavy, and Roni Michaely, 2005. "Using expectations to test asset pricing models." Financial management 34, no. 3: 31-64.

[11] Breeden, Douglas T., and Robert H. Litzenberger, 1978, "Prices of state-contingent claims implicit in option prices." Journal of business: 621-651.

[12] Epstein, L. G., & Zin, S. E., 1989, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework." Econometrica 937-969.

[13] Fama, Eugene F., and Kenneth R. French, 1992, "The cross-section of expected stock returns." the Journal of Finance 47: 427-465.

[14] Fama, Eugene F., and James D. MacBeth, 1973, "Risk, return, and equilibrium: Empirical tests." The Journal of Political Economy, 607-636.

[15] Figlewski, Stephen, 2010, "Estimating the implied risk-neutral density of the U.S. market portfolio." Volatility and Time Series Econometrics. Oxford University Press.

[16] Jackwerth, Jens Carsten, 2000. "Recovering risk aversion from option prices and realized returns." Review of Financial Studies 13, no. 2: 433-451.

[17] 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.

[18] Martin, Ian, 2013. "Simple variance swaps." Working Paper, National Bureau of Economic Research.

[19] Martin, Ian, 2015. "What is the Expected Return on the Market?" Working paper, London School of Economics.

[20] Mehra, Rajnish, and Edward C. Prescott, 1985, "The equity premium: A puzzle." Journal of monetary Economics 15: 145-161.

[21] MacMillan, Lionel W., 1986, "Analytic Approximation for the American Put Option." Advances in Futures and Options Research 1: 119-139.

[22] Ross, Steve, 2015 "The recovery theorem." The Journal of Finance 70.2: 615-648.

[23] Vissing-Jorgensen, Annette, and Orazio P. Attanasio, 2003, "Stock-market participation, inter-temporal substitution, and risk-aversion." American Economic Review: 383-391.