



A new predictor of U.S. real economic activity: The S&P 500 option implied risk aversion

DOI:

[10.1287/mnsc.2018.3049](https://doi.org/10.1287/mnsc.2018.3049)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Faccini, R., Konstantinidi, E., Skiadopoulos, G., & Sarantopoulou-Chiourea, S. (2018). A new predictor of U.S. real economic activity: The S&P 500 option implied risk aversion. *MANAGEMENT SCIENCE*, 65(10), 4451-4949. <https://doi.org/10.1287/mnsc.2018.3049>

Published in:

MANAGEMENT SCIENCE

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



A New Predictor of U.S. Real Economic Activity: The S&P 500 Option Implied Risk Aversion[§]

Renato Faccini^a, Eirini Konstantinidi^b,
George Skiadopoulos^c and Sylvia Sarantopoulou-Chiourea^d

January 5, 2018

Abstract

We propose a new predictor of U.S. real economic activity (REA), namely the representative investor's implied relative risk aversion (IRRA) extracted from S&P 500 option prices. IRRA is forward-looking and hence, it is expected to be related to future economic conditions. We document that U.S. IRRA predicts U.S. REA both in- and out-of-sample once we control for well-known REA predictors and take into account their persistence. An increase (decrease) in IRRA predicts a decrease (increase) in REA. We extend the empirical analysis by extracting IRRA from the South Korea, UK, Japanese and German index option markets. We find that South Korea IRRA predicts the South Korea REA both in- and out-of-sample, as expected given the high liquidity of its index option market. We show that a parsimonious yet flexible production economy model calibrated to the U.S. economy can explain the documented negative relation between risk aversion and future economic growth.

JEL Classification: E44, G13, G17

Keywords: Option prices, Risk aversion, Risk-neutral moments, Real Economic Activity, Production economy model.

[§]We are grateful to Lauren Cohen (the department editor), an anonymous associate editor, and two anonymous referees for their constructive, stimulating, and thorough comments. We would also like to thank Carol Alexander, Michael Anthopoulos, Michael Brennan, Chris Brooks, Peter Christoffersen, Gonzalo Cortazar, Charoula Daskalaki, Jin-Chuan Duan, Marcelo Fernandes, Jens Jackwerth, Anastasios Kagkadis, Sang Baum Kang, Alexandros Kostakis, Jesper Lindé, Lorian Mancini, Haroon Mumtaz, Michael Neumann, Stavros Panageas, Katerina Panopoulou, George Skoulakis, Christos Stefanadis, John Tsoukalas, Dimitri Vayanos, Andrea Vedolin, Raman Uppal, Costas Xiouros, Weiqi Zhang, and participants at the 2014 Multinational Finance Society Symposium (Larnaca), 2017 Financial Management Association European Conference (Lisbon), 2017 European Financial Management Association Annual Meeting (Athens), 2017 Conference on Research on Economic Theory and Econometrics (Milos), 2017 Xenophon Conference (Athens), and the Bank of Greece, ICMA Centre's University of Reading, Sveriges Riksbank and University of Sussex seminar series for many useful comments and discussions. Financial support from the Alliance MBS Strategic Research Investment Fund and the Research Centre of the University of Piraeus is gratefully acknowledged. Any remaining errors are our responsibility. George Skiadopoulos dedicates this paper to the memory of his beloved father, mentor and friend, Spyros Skiadopoulos.

^aSchool of Economics and Finance, Queen Mary, University of London, UK, r.faccini@qmul.ac.uk

^bAlliance Manchester Business School, University of Manchester, UK, eirini.konstantinidi@mbs.ac.uk

^cCorresponding author. School of Economics and Finance, Queen Mary, University of London, UK, and Department of Banking and Financial Management, University of Piraeus, Greece. Also Associate Research Fellow with Cass Business School and Warwick Business School, g.skiadopoulos@qmul.ac.uk, gskiado@unipi.gr

^dIndependent Authority for Public Revenue, s.sylvanna@gmail.com

1. Introduction

The question whether the growth of real economic activity (REA) can be predicted is of particular importance to policy makers, firms and investors. Monetary and fiscal policy as well as firms' business plans and investors' decisions are based on REA growth forecasts. There is an extensive literature which studies whether REA growth can be predicted by employing a number of financial variables (for a review, see Stock and Watson, 2003). This literature has become even more topical recently when the 2007 turbulence in the financial markets was followed by a significant economic recession which caught investors and academics by surprise (Gourinchas and Obstfeld, 2012). These facts highlight the link between financial markets and the real economy as well as the need to develop new accurate REA predictors based on financial markets' information (for a discussion on this, see also Ng and Wright, 2013).

In this paper, we propose a new predictor of U.S. REA. We investigate whether the representative investor's relative risk aversion (RRA) extracted from the S&P 500 market option prices (implied RRA, IRRA) predicts the growth of U.S. REA. The motivation for the choice of our predictor stems from the informational content that market prices of liquid index options are expected to possess. This is because S&P 500 options are inherently forward-looking contracts. Their payoff depends on the future state of the economy because the underlying stock index is a broad one that eliminates idiosyncratic risk. In addition, evidence suggests that informed traders tend to prefer option markets rather than the underlying spot market to exploit their informational advantage (e.g., Easley et al., 1998, Pan and Poteshman, 2006, and references therein), thus making option-based measures of highly liquid options, such as the S&P 500 ones, even more appealing for forecasting REA.

We extract U.S. IRRA's time series over July 1998-August 2015 via the Kang et al. (2010) formula. The formula proxies the difference between the risk-neutral and physical variance as a function of the representative investor's RRA by assuming a power utility function. It employs the S&P 500 risk-neutral volatility, risk-neutral skewness, risk-neutral kurtosis and the physical variance as inputs. We calculate the risk-neutral moments via Bakshi et al. (2003) method which uses the cross-section of traded S&P 500 option prices. Hence, IRRA incorporates information from all traded options by construction. The extracted IRRA values

are within the range of values reported by previous literature.

Next, we investigate whether U.S. IRRA predicts future U.S. REA. To this end, we use a number of alternative REA proxies. We test IRRA's forecasting ability across different forecasting horizons (up to one year) controlling for a large set of variables documented by the previous literature to predict REA. We conduct statistical inference carefully to cope with the persistence of regressors. We employ the recently developed instrumental variable test of Kostakis et al. (2015) designed to deal with the question of predictability in the case of multiple predictors whose order of persistence is unknown. We find that IRRA is a statistically significant predictor of REA over and above the set of control variables, i.e. IRRA contains information that has not already been incorporated by other financial predictors. An increase (decrease) in IRRA predicts a decrease (increase) in future U.S. REA. We document the predictive ability of IRRA both in- and out-of-sample. Application of Kelly and Pruitt's (2015) factor-based approach to forecasting corroborates our results.

We repeat our empirical analysis for South Korea, UK, Japan and Germany. The index option markets of these countries differ in their liquidity and hence, this additional analysis will shed more light on the properties of IRRA as a predictor of REA. We extract IRRA separately from options written on the KOSPI 200 (South Korea), FTSE 100 (UK), Nikkei 225 (Japan) and DAX (Germany) index. We find that IRRA predicts the respective REA both in- and out-of-sample only in the case of South Korea; an increase (decrease) in South Korea's IRRA predicts a decrease (increase) in South Korea's REA. This comes as no surprise since KOSPI 200 options have become the most actively traded option contracts in the world since their inauguration in 1997 (Ryu, 2015); the more liquid the index option market is, the richer its informational content for future REA is expected to be. Our findings also supplement Stock and Watson (2003) who document that the predictive ability of a given a REA predictor may differ across countries.

We explain the negative relation between IRRA and future REA by modelling a parsimonious yet flexible production economy in the spirit of the real business cycle (RBC) literature. The RBC framework is a natural candidate to explain our findings because it allows exploring the interactions of key macroeconomic variables that arise endogenously from the intertem-

poral optimization problem of households and firms within a general equilibrium setting. The model is standard: on the production side, we assume a representative firm operating in perfectly competitive markets for both the output and inputs of production. On the household side, we assume that the representative agent has preferences over consumption dictated by a power utility function with habits. The key difference with respect to the baseline RBC setting though, is that we abstract from shocks to technology on the firms' side and instead we focus on the real effects of shocks to households' risk aversion.

We calibrate the steady state solution of our model to the U.S. economy. We confirm that the model yields a negative relation between RRA and future REA by (i) investigating the impulse response function of output to an exogenous shock in RRA, and (ii) running predictive regressions that employ simulated values of REA and RRA generated by our model. The intuition for the model's predictions is that a negative shock in RRA makes agents decrease consumption and hence increase savings and thus investment. This boosts REA via the accumulation of capital.

Related literature: Our paper ties four strands of literature. The first strand has to do with the use of financial variables to predict REA. The rationale is that financial markets reflect investors' perceptions about the future state of the economy and hence they can predict REA. The term spread (Estrella and Hardouvelis, 1991) and default spread (Stock and Watson, 2003) are two prominent predictors of REA. More recently, other financial variables such as asset pricing factors (Liew and Vassalou, 2000), the TED spread (Chiu, 2010), the Baltic dry index (Bakshi et al., 2012), commodity futures open interest (Hong and Yogo, 2012), and commodity-specific factors (Bakshi et al., 2014) have been found to predict REA. Bakshi et al. (2011) also document empirically that option implied information, namely the forward variance extracted from S&P 500 index options, predicts REA. Our paper differs in that we use index option prices to estimate IRRA whose predictive ability over REA is not only documented empirically but is also theoretically explained within a RBC framework.

The second strand of literature has to do with the estimation of the representative agent's risk aversion from index options market prices (Ait-Sahalia and Lo, 2000, Jackwerth, 2000, Rosenberg and Engle, 2002, Bliss and Panigirtzoglou, 2004, Bakshi and Madan, 2006, Kang

and Kim, 2006, Kang et al., 2010, Kostakis et al., 2011, Barone-Adesi et al., 2014, and Duan and Zhang, 2014). This is possible due to the theoretical relation of risk aversion to the ratio of the risk-neutral distribution and the subjective distribution of the option's underlying index; the former can be recovered from option prices (for a review, see Jackwerth, 2004). We choose Kang et al. (2010)'s methodology to extract IRRA because it is parsimonious in terms of the required inputs. Most importantly, these inputs can be estimated accurately from the cross-section of market option prices which are readily available.

The third strand of literature uses the informational content of market option prices to address a number of topics in economics and finance. The rationale is that market option prices convey information which can be used for policy making (Söderlind and Svensson, 1997), risk management (Chang et al., 2012, Buss and Vilkov, 2012), market timing (Kostakis et al., 2011) and stock selection purposes (DeMiguel et al., 2013, for reviews see Giamouridis and Skiadopoulos, 2012, Christoffersen et al., 2013). Surprisingly, there is a paucity of research on whether the information embedded in index option prices can be used to predict REA, too. To the best of our knowledge, Bakshi et al. (2011) is the only paper which explores this and documents that forward variances extracted from index options forecast REA.

The fourth strand has to do with the use of RBC models in the finance literature. So far, RBC models have been used to address pricing puzzles by considering the effects of technology shocks (e.g., Jerman, 1998, Boldrin et al., 2001, Kaltenbrunner and Lochstoer, 2010, Papanikolaou, 2011). We deviate from previous literature and we investigate the effects of exogenous shocks to risk-aversion to future REA.

The rest of the paper is structured as follows. Section 2 describes the U.S. data. Section 3 explains how we extract U.S. IRRA. Section 4 presents the empirical evidence on the U.S. IRRA as a predictor of U.S. REA. Section 5 presents further evidence on the predictive content of IRRA in the case of other countries. Section 6 presents the RBC model and discusses its results. Section 7 concludes.

2. U.S. Data

2.1. *Real economic activity data*

We obtain monthly data for six alternative measures to proxy U.S. REA over July 1998 to August 2015. This is a rich period because it includes events of importance such as the August 1998 Russian crisis, the early 2000s recession and the subsequent bullish U.S. stock market, the 2007-2009 financial crisis and the great economic recession as well as the 2008-2014 U.S. quantitative easing era.

First, we use industrial production (IPI) which measures the amount of the industries output. Second, we consider non-farm payroll employment (NFP) defined as the number of employees in the non-farm sectors in the U.S. economy. Third, we employ real retail sales including food services sales as a proxy for retail sales (RS). Fourth, we use housing starts (HS) which measures the total newly started privately owned housing units. We use the monthly logarithmic growth rates for these four REA proxies. We obtain these data from the Federal Reserve Economic Data (FRED) database maintained by the Federal Reserve Bank of St. Louis (FRED).

Fifth, we consider the Chicago Fed National Activity Index (CFNAI). CFNAI is a weighted average of 85 existing monthly indicators of national economic activity. It is constructed to have an average value of zero and a standard deviation of one. Since economic activity tends toward trend growth rate over time, a positive (negative) index value corresponds to growth above (below) trend. The 85 economic indicators that are included in the CFNAI are drawn from four broad categories of data: production and income; employment, unemployment, and hours; personal consumption and housing; and sales, orders, and inventories. We obtain CFNAI from FRED.

Finally, we use the Aruoba-Diebold-Scotti (ADS, Aruoba et al., 2009) business conditions index. ADS is compiled based on six economic indicators: weekly initial jobless claims, monthly payroll employment, industrial production, personal income less transfer payments, manufacturing and trade sales and quarterly real GDP. It blends high- and low-frequency information, as well as stock and flow data. The average value of the ADS index is zero. Positive (negative) values indicate better-(worse-) than-average conditions. We obtain ADS

from the Philadelphia Fed webpage.

2.2. IRRA inputs: S&P 500 options and 5-minute spot data

We use the following data to estimate U.S. IRRA at any point in time. First, we obtain S&P 500 European style index option data (closing quotes calculated based on the midpoint of bid and ask prices), their corresponding implied volatilities, the closing price of the S&P 500 and the continuously paid dividend yield for January 1996 to August 2015 from the Ivy DB OptionMetrics database. Implied volatilities are calculated based on the midpoint of bid and ask quotes using Merton's (1973) model. We use the obtained implied volatilities to compute the S&P 500 risk-neutral moments with a τ -month horizon ($\tau = 1$ month) via Bakshi et al. (2003)'s model-free method (see Appendix A). As a proxy for the risk-free rate, we use the zero-coupon curve also provided by Ivy DB. We filter implied volatilities to remove any noise by applying a number of filtering criteria to their corresponding option prices. We only consider out-of-the-money and at-the-money options with time-to-maturity 7 to 90 days. We also discard options with zero open interest and zero trading volume. Furthermore, we retain option contracts that do not violate Merton's (1973) no-arbitrage condition and have implied volatilities less than 100%. We also eliminate options that form vertical and butterfly spreads with negative prices, as well as option contracts with zero bid prices and premiums below 3/8\$.

Second, we obtain 5-minute intra-day S&P 500 prices from Thomson Reuters Tick History to estimate the S&P 500 physical variance with a τ -month horizon ($\tau = 1$ month). We assume that the physical variance follows a random walk in line with Andersen and Bollerslev (1998). The τ -month physical variance, $\sigma_{p,t}^2(\tau)$, equals the realized variance from $t - \tau$ to t , $RV_{t-\tau,t}$, computed as the sum of the daily realized variances plus the sum of the overnight squared returns (OR) of the S&P 500 over the last one month, i.e.

$$RV_{t-\tau,t} = \sum_{i=t-\tau}^t \sigma_i^2 + \sum_{i=t-\tau}^t OR_i^2 \quad (1)$$

where σ_t^2 is the realized variance on day t and OR_t is the overnight return. We calculate overnight returns as the log difference of each day's opening price minus the closing price of

the previous day: $OR = \ln S_t^{Op} - \ln S_{t-1}^{Cl}$, where S^{Op} and S^{Cl} are the opening and the closing prices of the S&P 500 index, respectively.

2.3. Control variables and a large macroeconomic dataset

We collect data on a number of variables documented to predict REA by previous literature; these will be used as control variables in the subsequent predictive regressions. Data span the same period that IRRA is extracted for, i.e. July 1998-August 2015. First, we obtain monthly data from the FRED website to measure the term spread (TERM, difference between the ten-year Treasury bond rate and the three-months Treasury bill rate), default spread (DEF, difference between the yields of the Moody's AAA and BAA corporate bonds) and TED spread (difference between the three-months U.S. Libor rate and the three-months Treasury bill rate). Second, we obtain monthly data on the monthly Fama-French (1996) high minus low (HML) and small minus big (SMB) factors from Wharton Research Data Services (WRDS). Third, we collect data on the Baltic Dry Index (BDI) from Bloomberg.

Fourth, we obtain data on 22 individual commodity futures from Bloomberg to construct the three Daskalaki et al. (2014) commodity-specific factors, namely hedging-pressure (HP), momentum (MOM) and basis factors (BASIS); Appendix B describes the construction of these factors. Table 1 lists the employed commodities categorized in five sectors (grains and oilseeds, energy, livestock, metals and softs). In addition, we construct a commodity futures open interest variable (OI) in line with Hong and Yogo (2012). First, we compute the growth rate of open interest for each commodity futures. Then, at any given point in time, we compute the median of the growth rates of open interest for all commodities futures of each sector. Last, we compute the equally weighted average of the medians growth rates of all sectors.

Fifth, we use the options data discussed in Section 2.2 to compute at time t the forward variance $FV_{t,t+1}$ between t and $t+1$, i.e. the forward variance with a one-month horizon. To this end, we follow Bakshi et al. (2011). Finally, we obtain the McCracken and Ng (2016) large macroeconomic dataset from FRED. This dataset consists of 134 monthly macroeconomic U.S. indicators and we will use it in the out-of-sample tests in Section 4.2.

3. Extracting risk aversion from option prices

Bakshi and Madan (2006) derive a formula which can be used to extract the risk aversion of the representative agent from European options market prices. By assuming that a power utility function describes the representative agent's preferences, RRA is extracted from the following equation:

$$\frac{\sigma_{q,t}^2(\tau) - \sigma_{p,t}^2(\tau)}{\sigma_{p,t}^2\tau} \approx -\gamma\sigma_{p,t}(\tau)\theta_{p,t}(\tau) + \frac{\gamma^2}{2}\sigma_{p,t}^2(\kappa_{p,t}(\tau) - 3) \quad (2)$$

where γ is the RRA coefficient, $\sigma_{q,t}^2(\tau)$ is the risk-neutral variance of the continuously compounded return distribution at time t with horizon τ , and $\sigma_{p,t}^2(\tau)$, $\theta_{p,t}(\tau)$ and $\kappa_{p,t}(\tau)$ are the *physical* variance, skewness and kurtosis of the continuously compounded return distribution at time t with horizon τ , respectively.

Equation (2) shows that the RRA extraction requires estimation of the higher order physical moments (skewness and kurtosis) which is a challenging task. On the one hand, a long time series is required to estimate higher order physical moments accurately and on the other hand, a small sample size is needed to capture their time variation (Jackwerth and Rubinstein, 1996). To bypass the estimation of the physical higher order moments, we use Kang et al. (2010) formula which is a variant of equation (2), i.e.

$$\frac{\sigma_{p,t}^2(\tau) - \sigma_{q,t}^2(\tau)}{\sigma_{q,t}^2\tau} \approx \gamma\sigma_{q,t}(\tau)\theta_{q,t}(\tau) + \frac{\gamma^2}{2}\sigma_{q,t}^2(\kappa_{q,t}(\tau) - 3) \quad (3)$$

where $\theta_{q,t}(\tau)$ and $\kappa_{q,t}(\tau)$ is the *risk-neutral* skewness and kurtosis of the continuously compounded return distribution at time t with horizon τ , respectively. Kang et al. (2010) derive equation (3) by also assuming that the representative agent's preferences are described by a power utility function. Then, they use the moment generating functions of the risk-neutral and physical probability distributions and they truncate their expansion series appropriately.

The extraction of IRRA from either equation (2) or (3) is model-dependent. However, the advantage of extracting RRA from equation (3) rather than from equation (2) is twofold. First, the former equation requires the risk-neutral skewness and kurtosis rather than their physical counterparts as inputs. Hence, it circumvents the above mentioned challenges of

estimating higher order physical moments. This is because the estimation of the higher order risk-neutral moments is model-free (e.g., Bakshi et al., 2003, Jiang and Tian, 2005, Carr and Wu, 2009). Therefore, even though we use a model-dependent method to back out RRA, three out of the four required inputs are model-free in contrast to equation (2).¹ Second, the risk-neutral moments are forward-looking (they can be computed at time t from the market option prices observed at time t) whereas the physical moments estimates are backward-looking (they rely on past historical data). This makes equation (3) the natural choice for the purposes of our study.

We use the 30-days realized variance calculated from 5-minute S&P 500 prices as an estimate of the physical variance. We compute the S&P 500 risk-neutral moments with a horizon of $\tau = 1$ -month by implementing the Bakshi et al. (2003) formulae (see Appendix A).² In line with Bakshi and Madan (2006), Kang et al. (2010) and Duan and Zhang (2014), we use the generalized method of moments (GMM, Hansen, 1982) to estimate RRA. We minimize the following objective function with respect to γ :

$$\begin{aligned} J_T &\equiv \min_{\gamma} g_T' H_T g_T \\ g_T &\equiv \frac{1}{T} \sum_{t=1}^T \epsilon_t \otimes Z_t \\ \epsilon_t &\equiv \frac{\sigma_{p,t}^2(\tau) - \sigma_{q,t}^2(\tau)}{\sigma_{q,t}^2(\tau)} - \gamma \sigma_{q,t}(\tau) \theta_{q,t}(\tau) - \frac{\gamma^2}{2} \sigma_{q,t}^2(\kappa_{q,t}(\tau) - 3) \end{aligned} \quad (4)$$

¹Inevitably, any method to extract RRA from option prices will be model-dependent. For instance, an alternative way to extract IRRA would be the Bliss and Panigirtzoglou (2004) method, which uses the relation between the ratio of the risk-neutral to the physical probability density function and the stochastic discount factor. However, that method is model-dependent, too, because it requires an assumption on the utility function of the representative agent as well as further parametric transformations and assumptions. Given that there is not a model-free method to back out IRRA, the "first best" (i.e. use a model-free method to estimate IRRA) cannot be attained. However, the choice of the Kang et al. (2010) formula attains the "second best" (i.e. get as many parameters as possible estimated in a model-free way): three out of its four required inputs can be estimated in a model-free way.

²The risk-neutral and physical variances should not be annualized when used as inputs in equations (2) and (3). To prove this statement, we multiply and divide equation (3) by 252,

$$\frac{\sigma_{p,t}^{*2}(\tau) - \sigma_{q,t}^{*2}(\tau)}{\sigma_{q,t}^{*2} \tau} \approx \frac{\gamma}{\sqrt{252}} \sigma_{q,t}^*(\tau) \theta_{q,t}(\tau) + \frac{1}{2} \frac{\gamma^2}{252} \sigma_{q,t}^{*2}(\kappa_{q,t}(\tau) - 3) = \gamma^* \sigma_{q,t}^*(\tau) + \frac{1}{2} \gamma^{*2} \sigma_{q,t}^{*2}(\kappa_{q,t}(\tau) - 3) \theta_{q,t}$$

where $*$ denotes the annualized values. Hence, if we use the annualized instead of the raw variance as input, we shall estimate the annualized risk aversion coefficient, $\gamma^* = \frac{\gamma}{\sqrt{252}}$, which differs from the raw risk aversion estimate γ . Hence, we use the raw values of the variances as inputs to estimate the risk aversion coefficient.

where J_T is the objective function, g_T denotes the sample mean estimate of the orthogonality condition of the instruments, H_T is the inverse of the variance-covariance matrix of the function g_T and Z_T are the instruments. In equation (4), there are as many moment conditions as instruments. In line with the three above mentioned studies, we use three different sets of instruments to assess whether the choice of instruments affects the extracted IRRA. The first set consists of a constant and one lag of the risk-neutral variance $[\sigma_{q,t-1}^2(\tau)]$. The second set consists of a constant and two lags of the risk-neutral variance $[\sigma_{q,t-1}^2(\tau), \sigma_{q,t-2}^2(\tau)]$. The third set contains a constant and three lags of the risk-neutral variance $[\sigma_{q,t-1}^2(\tau), \sigma_{q,t-2}^2(\tau), \sigma_{q,t-3}^2(\tau)]$. We apply a two-step GMM.

In line with the three above cited studies, we extract IRRA for a constant time horizon $\tau = 1$ month (=30 days). We record the risk-neutral moments and realized variance at the last trading day of each month. We use equation (3) to extract the monthly IRRA series with a rolling GMM estimation using a rolling window of size 30 months. This yields an IRRA time series for July 1998 - August 2015 given that our option dataset spans January 1996 to August 2015.

Figure 1 shows IRRA's monthly time variation for each one of the three sets of instruments extracted from the rolling GMM. Four remarks are in order. First, IRRA values range from 2.27 to 9.55. These fall within the range of IRRA estimates reported by the previous literature. Ait-Sahalia and Lo (2000) report a full-sample IRRA of 12.7, Rosenberg and Engle (2002) report values from 2.26 to 12.55, Bakshi et al. (2003) report values between 1.76 and 11.39, Bliss and Panigirtzoglou (2004) report a full sample estimate of 4.08, Bakshi and Madan (2006) report values from 12.71 to 17.33, Kang and Kim (2006) report values between 2 and 4, Kang et al. (2010) 1.2 to 1.4, Barone-Adesi et al. (2014) report values between -0.5 and 3, and Duan and Zhang (2014) obtain values from 1.8 to 7.1.

Second, IRRA's time variation is similar across all three sets of instruments. In the remainder of the paper, we report results for the case of the IRRA estimated by the first set of instruments comprising the constant and one lag of the risk-neutral variance. Third, we can see that the U.S. IRRA is not affected by the 1998 Russian crisis and the early March 2001-November 2001 U.S. recession whereas it increases significantly up to the 2008 financial

crisis. Interestingly, it starts decreasing thereafter; this pattern may be a result of the 2008-2014 quantitative easing monetary policy exercised by the Fed which might have alleviated U.S. agents' concerns. Finally, IRRA is persistent ($\rho = 0.986$). We will take this persistence into account in the subsequent analysis where we will explore whether IRRA predicts REA.

4. Predicting U.S. REA

We examine whether the U.S. IRRA predicts U.S. REA growth first in-sample and then out-of-sample.

4.1. In-sample evidence

To identify whether IRRA predicts REA growth over h forecasting horizons, we regress each one of the employed REA measures on IRRA by controlling for a set of variables documented to predict REA. We estimate the predictive regression:

$$REA_{i,t+h} = c_i + \beta_{1,i} REA_t + \beta_{2,i} IRRA_t + \beta'_{3,i} \mathbf{x}_t + \varepsilon_{i,t+h} \quad (5)$$

where $REA_{i,t+h}$ denotes the growth of the $i - th$ REA proxy ($i = 1$ for IPI, 2 for NFP, 3 for RS, 4 for HS, 5 for CFNAI, 6 for ADS) over the period t to $t + h$, $REA_{i,t}$ is the growth of the $i - th$ REA proxy over the period $t - h$ to t , $IRRA_t$ is the implied risk aversion at time t and $\mathbf{x}_t$ is a (11×1) vector which contains a set of control variables. We compute the h -month overlapping log growth rates of IPI, NFP, RS, and HS. The values of CFNAI and ADS signify growth or recession by construction and hence, there is no need to compute the growth rates for these two REA proxies. We set $h = 1, 3, 6, 9, 12$ months.

We consider the following control variables: term spread (Estrella and Hardouvelis, 1991), default spread (Faust et al., 2013), TED spread that proxies for funding liquidity (Chiu, 2010), SMB and HML Fama-French (1996) factors (Liew and Vassalou, 2000), Baltic dry index (BDI, Bakshi, et al., 2012), forward variances (FV, Bakshi et al., 2011), commodity-specific factors (hedging-pressure, momentum, and basis, Bakshi et al., 2014), and the growth rate of the commodity futures market open interest (Hong and Yogo, 2012). The sample spans July

1998 - August 2015 (206 observations).

We conduct inference by taking the high degree of IRRA’s persistent into account. This is because statistical inference is flawed once conducted by standard/Newey-West t -statistic in the case where predictors are persistent (e.g., Kostakis et al., 2015). More specifically, we use the IVX-Wald test statistic (Kostakis et al., 2015) to test the predictive ability of IRRA (for a description of the test, see Appendix C). The IVX-Wald test is robust to the unknown time series properties of the predictors. In particular, it does not assume a-priori knowledge of the degree of persistence and it allows for different classes of persistence of the predictor variables, ranging from purely stationary to purely non-stationary. It also allows conducting inference in the case of multiple predictors whereas previous tests related to predictors persistence are developed only for single predictor models (e.g., Campbell and Yogo’s test, 2006).

Table 2 reports the results from estimating equation (5) for forecasting horizons $h = 1, 3, 6, 9$ and 12 months. Panel A shows the standardized ordinary-least-squares (OLS) coefficient estimates, the Newey-West and IVX-Wald test p -values of each one of the predictors. One, two and three asterisks denote rejection of the null hypothesis of a zero coefficient based on the p -values of the IVX-Wald test at the 1%, 5% and 10% level, respectively. Panel B shows the in-sample adjusted R^2 for any given REA proxy. Two remarks can be drawn in the case of a one-month forecasting horizon. First, we can see that U.S. IRRA predicts all but one (i.e. RS) U.S. REA proxies.³ Second, the sign of the IRRA coefficient is negative in all cases. This suggests that an increase in IRRA predicts a decrease in REA. Extending the evidence from the one-month results, we can see that IRRA predicts most of the REA proxies at longer horizons as well. IRRA is significant for four out of six REA proxies for longer horizons. More specifically, IRRA predicts NFP and HS at all horizons $h > 1$ month. It also predicts RS for

³We explore whether U.S. IRRA predicts U.S. REA when we extract U.S. IRRA by an alternative method to that of Kang et al. (2010). We use the Kostakis et al. (2011) U.S. IRRA dataset extracted from the Bliss and Panigirtzoglou (2004) approach. We are grateful to Kostakis et al. (2011) who have kindly shared their IRRA estimates with us. They have extracted U.S. IRRA from futures S&P 500 options over July 1998 - May 2010 by assuming a representative agent whose preferences are described by a power utility function. The results, presented in the online appendix, show that the evidence of predictability is weaker in this case. Even though the period under scrutiny in this explorative exercise is significantly shorter from the one employed in the main body of the paper and hence results are not comparable, the weaker performance of the Bliss and Panigirtzoglou U.S. IRRA as a U.S. REA predictor may be attributed to the way the Bliss and Panigirtzoglou IRRA is extracted. The extraction of the Bliss and Panigirtzoglou (2004) IRRA requires a conversion of the implied to the subjective PDF. The Kang et al (2010) IRRA does not require such a transformation.

horizons up to nine months and IPI for longer horizons ($h = 9$ and 12 months). R^2 range from 0.2 to 0.85 across the various REA proxies and forecasting horizons.

4.2. Out-of-sample evidence

In Section 4.1 we documented that U.S. IRRA forecasts U.S. REA in an in-sample setting. In this section, we assess the forecasting ability of U.S. IRRA in a real time out-of-sample setting over October 2007 - August 2015. This is also a period of particular interest because it includes the onset and development of the recent financial crisis and the subsequent significant economic recession (also termed Great Recession) and the quantitative easing conducted by the U.S. Fed. This is a period of time where standard predictors failed to forecast the forthcoming realized recession. As a result, we have posed a high hurdle for our IRRA predictor to overcome. For each REA proxy, we estimate equation (5) recursively by employing an expanding window; the first estimation sample window spans July 1998 - September 2007. Then, at each point in time, we form $h = 1, 3, 6, 9, 12$ months-ahead REA forecasts.

We use the out-of-sample R^2 (Campbell and Thompson, 2008) to evaluate the out-of-sample forecasting performance of IRRA. The out-of-sample R^2 shows whether the variance explained by a full model (which contains IRRA in the set of predictors) is greater or smaller than the variance explained by a restricted model (which does not contain IRRA within the set of predictors). Then, the out-of-sample R_i^2 obtained from predicting the $i - th$ REA proxy is defined as:

$$R_{i,h}^2 = 1 - \frac{\text{var} [E_t (REA_{i,t+h}^{Full}) - REA_{i,t+h}]}{\text{var} [E_t (REA_{i,t+h}^{Restricted}) - REA_{i,t+h}]} \quad (6)$$

where $E_t (REA_{i,t+h}^{Full})$ and $E_t (REA_{i,t+h}^{Restricted})$ denote h -month ahead forecasts from the full and restricted model, respectively. A positive (negative) out-of-sample R^2 suggests that the full model outperforms (underperforms) the restricted model and hence, IRRA has (has no) out-of-sample predictive ability.

We consider two alternative model specifications. First, we obtain forecasts from the regression model described by equation (5). In this case, the forecasts $E_t (REA_{i,t+h}^{Full})$ for the $i - th$ REA proxy obtained from the full model are:

$$E_t (REA_{i,t+h}^{Full}) = b_0 + b_1 REA_t + b_2 IRRA_t + b_3' \mathbf{x}_t \quad (7)$$

and the forecasts $E_t (REA_{i,t+h}^{Restricted})$ from the restricted model are:

$$E_t (REA_{i,t+h}^{Restricted}) = b_0 + b_1 REA_t + b_3' \mathbf{x}_t \quad (8)$$

Second, we consider Kelly and Pruitt's (2015) three-pass regression filter (3PRF). 3PRF is developed within the factor-based approach to forecasting advocated by Stock and Watson (2002a, 2002b). Hence, it is a dimension reduction method that is suitable in the case where the number of potentially useful for prediction variables is large and the number of observations is relatively small. In contrast to previous factor-based forecasting methods, 3PRF identifies factors that are relevant to the variable that we wish to forecast; these factors may be a strict subset of the factors driving the predictor variables. For the purposes of implementing the 3PRF approach, we need to consider a large dataset. Hence, we consider a dataset which includes IRRA and the 134 McCracken and Ng (2016) macroeconomic variables. Following McCracken and Ng (2016), we transform the original time series into stationary and we remove outliers; outliers are defined as observations that deviate from the sample median by more than ten interquartile ranges. Then, we standardize the transformed variables. Following Kelly and Pruitt (2015), we extract one 3PRF factor and we take care to avoid any look-ahead bias given that the estimation of the first two steps uses the full sample (for a description, see Appendix D). For any given REA proxy to be predicted, we construct the 3PRF factor by removing the variables from the McCracken and Ng (2016) dataset which measure the same notion of economic activity as the REA proxy does. Once we construct the 3PRF factor, our 3PRF model is

$$REA_{t+h}^i = \gamma_0 + \gamma_1 F_t + u_{t+h} \quad (9)$$

where F_t is the 3PRF factor. We obtain forecasts for the $i - th$ REA proxy from the full and restricted 3PRF models defined as:

$$E_t (REA_{i,t+h}^{Full}) = \gamma_0 + \gamma_1 F_t^{Full} \quad \text{and} \quad E_t (REA_{i,t+h}^{Restricted}) = \gamma_0 + \gamma_1 F_t^{Restricted}$$

respectively, where we extract F_t^{Full} from a large set of variables which includes IRRA and the McCracken and Ng (2016) macroeconomic variables and $F_t^{Restricted}$ from a large set of variables which includes only McCracken and Ng (2016) variables.

Table 3 shows the out-of-sample R^2 for the case of forecasts obtained from the regression predictive models [equations (7) and (8), Panel A] and from the two 3PRF models (Panel B). We can see that the out-of-sample R^2 is positive in most cases, i.e. the full model performs better than the restricted model. This implies that the inclusion of IRRA is statistically significant in an out-of-sample setting, too. In the case of forecasts obtained by the regression models, the evidence on this is somewhat weaker for longer horizons. For RS and HS, the full model outperforms the restricted model across all predictive horizons. The out-of-sample R^2 is also positive for short and intermediate forecasting horizons ($h = 1, 3, 6$ months) in the case of NFP and CFNAI. In addition, IRRA predicts IPI and ADS for short horizons ($h = 1, 3$ months and $h = 1$ month, respectively). In the case of the 3PRF model, the out-of-sample R^2 is positive in all but one cases. The only exception occurs at a one-month horizon for CFNAI.^{4,5}

Finally, we examine the stability of the IRRA coefficients over the out-of-sample period. To this end, we estimate equation (5) at each point in time over October 2007-August 2015 by employing an expanding window for each one of the employed forecasting horizons. Figure 2 shows the standardized IRRA coefficients in the case of the one-month forecasting horizon where we use IPI, NFP, RS, HS, CFNAI and ADS as REA proxies (Panels A, B, C, D, E and F, respectively). We can see that the time series evolution of the estimated IRRA coefficient is stable over time. The sign of the estimated IRRA coefficient is negative at each estimation time step suggesting that a decrease in IRRA predicts an increase in future REA. This is in line with the results obtained from the in-sample analysis (see Table 2). The time series

⁴We investigate the robustness of our U.S. in- and out-of-sample results by estimating U.S. IRRA for alternative rolling window sizes, namely 45 and 60 months. Results are qualitatively similar and we report them in the online appendix.

⁵We consider two additional robustness checks of our U.S. out-of-sample results. First, we examine alternative out-of-sample periods by reducing the size of the sample used to initialize the out-of-sample experiment. This delivers January 2004 - August 2015, January 2005 - August 2015, January 2006 - August 2015, and January 2007 - August 2015 as alternative out-of-sample periods. Second, we consider an alternative benchmark model to calculate the out-of-sample R^2 , namely the moving average of past 30-month REA values for the U.S. The results for both robustness tests are similar to those reported in the main body of the paper, and we present them in the online appendix.

evolution of the estimated IRRA coefficient is stable over time for the longer horizons, too, and hence due to space limitations we do not report additional figures.

5. Further evidence from other countries

We examine whether IRRA predicts REA in South Korea, Germany, Japan and UK. South Korea is the most active index option markets in the world (Ryu, 2015). Options written on the South Korea Composite Stock Price Index (KOSPI) 200 were introduced in July 1997. Since then, KOSPI options have become one of the most actively traded contracts in the world. In 2014, the aggregate trading volume was 462 million contracts. This figure is greater than the 2014 aggregate trading volume of each one of the other index options that we consider (S&P 500: 223 million, Nikkei 225: 26.7 million, DAX: 20.8 million and FTSE 100: 8.6 million). Therefore, the chosen panel of countries can shed further light on the conditions under which IRRA may predict future REA since the informational content of market option prices may vary with the option market's liquidity.

5.1. Data

To extract IRRA for South Korea, UK, Japan and Germany we obtain European index option data (closing prices) written on KOSPI 200 options data from the Korea Exchange (KRX) and on FTSE 100, Nikkei 225 and DAX from Optionmetrics Ivy DB Global Indices and their respective implied volatilities spanning January 2004 to June 2015. We filter the options data as follows. We only consider out-of-the-money and at-the-money options. We retain options with time-to-maturity 7 to 60 days for South Korea and 7 to 270 days for the UK, Japan and Germany. We remove options with zero open interest and zero trading volume.⁶ We also discard option contracts which have an implied volatility that is less than zero and

⁶We do not apply the open interest filter in the case of FTSE 100 index options for the year 2006. This is because the open interest of FTSE 100 index options is zero across all option observations for the whole year 2006. This is attributed to an error in recording the open interest variable for 2006, following discussions with Optionmetrics. We do not expect the non-application of the open interest constraint to affect our results though: The zero open interest filter removes only an extra 5.46% of the observations for FTSE 100 index options data spanning January 2007-June 2015 on top of the ones already removed when all other filters are applied.

greater than 100%; we use Merton’s (1973) model to back out the implied volatility.⁷ We do not consider South Korean options with premiums less than 0.02. Finally, we retain option contracts that do not violate Merton’s (1973) no arbitrage bounds.

To ensure a common set of REA proxies across countries, we use monthly data for two alternative measures to proxy REA for the South Korea, UK, Japan and Germany from June 2006 to June 2015. First, we obtain data on industrial production (IP, proxied by industrial production index) from the Bank of Korea for South Korea and from Bloomberg for UK, Japan and Germany. Second, we obtain data on unemployment (U, proxied by unemployment rate) from Bloomberg for all countries. We use logarithmic growth rates for all REA proxies.

We use the following control variables in the predictive regressions: TERM, DEF, TED, BDI, and FV.⁸ These span June 2006 to June 2015. The construction of the spread variables (i.e. TERM, DEF and TED) is shown in Table 4 and is dictated by data availability. Finally, we obtain BDI from Bloomberg and we construct FV using the index options data for each country as discussed above.

5.2. Results

We estimate IRRA for South Korea, UK, Japan and Germany by GMM using a rolling window of 30 monthly observations and three separate sets of instruments, just as we did in the case with the U.S. IRRA in Section 3. This delivers the IRRAs of the respective countries over June 2006 to June 2015. We use monthly data on KOSPI 200, FTSE 100, Nikkei 225 and DAX options to compute the one-month horizon risk-neutral moments for South Korea, UK, Japan and Germany, respectively. To estimate IRRA at any point in time, we estimate the one-month physical variance. For South Korea, we use monthly data on KOSPI 200 and we

⁷Two remarks are at place for South Korea. First, we use the 91 days certificate of deposit (CD) rate as the risk free rate which is the standard practice for the Korea market (e.g., Kim and Kim, 2005). This is because the South Korea treasury bill market is not liquid. For the remaining countries, we use zero yield curve that IvyDB Global Indices provides. Second, we set the continuous dividend yield equal to zero. This is because we have no access to data on South Korea dividends. However, the effect of the value of dividend yield on the risk-neutral moments is small. This is because we use OTM options to calculate the risk-neutral moments. OTM options have a small delta and therefore any effect of dividends on the underlying index price and hence on the option price will be small. For the remaining countries, we use the respective dividend yield provided by IvyDB Global Indices.

⁸We do not consider the South Korea, UK, Japan and Germany analogues of the U.S. Fama-French factors and Daskalaki et al. (2014) commodity-related variables as controls. This is because the corresponding Fama-French factors are not available for the entire time period under consideration and there are no commodity futures contracts traded in Korea, UK, Japan and Germany.

construct the one-month physical variance as one-step ahead forecasts from a GARCH(1,1) model using a rolling window of 30 observations of KOSPI 200.⁹ For the UK, Japan and Germany we use 5-minute intra-day FTSE 100, Nikkei 225 and DAX prices obtained from Thomson Reuters Tick History to estimate the FTSE 100, DAX and Nikkei 225 physical variance with a one-month horizon, respectively.

Figure 3 shows the evolution of the monthly Korea, UK, Japanese and German IRRA over June 2006-June 2015. IRRA's time variation is similar across all three sets of instruments; in the remainder of this section we will employ IRRA extracted from the first set of instruments just as we did in the U.S. case. We can see that the Korea, UK, Japan and German IRRA's spike upwards over the 2007-2008 U.S. financial crisis just as the U.S. IRRA does; this is another manifestation of the interconnectedness of financial markets across the globe. Interestingly, the Korea IRRA is negative over April 2011 to February 2012. This suggests that the Korea representative agent exhibits a risk-loving behaviour over this period.

Next, we examine whether IRRA predicts REA in South Korea, UK, Japan and Germany in-sample. To this end, we estimate equation (5) in-sample across the full sample June 2006 - June 2015. To alleviate multicollinearity concerns, we orthogonalize the control variables that have a correlation with IRRA bigger than 0.50. In the UK, we orthogonalize TERM on IRRA ($\rho = -0.66$) and FV on IRRA ($\rho = 0.75$). In Japan, we orthogonalize TERM on IRRA ($\rho = 0.66$) and TED on IRRA ($\rho = 0.69$). In South Korea and Germany, all control variables have pairwise correlations lower than 0.50.

Table 5 Panel A reports results from estimating equation (5) in-sample for $h = 1, 3, 6, 9$ and 12 months for South Korea, UK, Japan and Germany. We report the standardized OLS coefficient estimates and the IVX-Wald test's p -value (within brackets) for the IRRA predictor. We can draw two main findings. First, we can see that IRRA predicts REA in the case of South Korea and UK in most of the cases. More specifically, it predicts the South Korea IPI for $h > 3$ months and U for $h > 1$ month. It also predicts the UK IPI and U across all horizons. On the other hand, there is weak evidence of in-sample predictability for Japan,

⁹We do not estimate the one-month physical variance using high frequency data as we did in the U.S. case. This is because the intra-day KOSPI 200 futures data are significantly contaminated with measurement errors and typographical errors; the provided documentation does not allow correcting them.

where IRRA predicts only U for $h = 9$ months. Interestingly, IRRA does not predict any of the German REA proxies. Second, we can see that an increase (decrease) in IRRA predicts a decrease (increase) in REA; the estimated IRRA coefficient is negative in the case of IPI, whereas it is positive in the case of U implying that an increase in IRRA predicts an increase in the rate of unemployment and hence, a decrease in future REA.

Next, we examine IRRA's predictive ability in an out-sample setting over January 2009 - June 2015 for South Korea and UK where we found strong evidence of in-sample predictability. The out-of-sample period for these two countries does not start in October 2007 as it was the case for U.S. This is because such a choice would yield a sample with only 15 observations to be used for the estimation of the predictive model in the first out-of-sample estimation step (versus 111 observations used in the U.S. in the first out-of-sample estimation step). Table 5 Panel B reports the out-of-sample R^2 .¹⁰ We can see that the in-sample predictive ability of IRRA also carries over to an out-of-sample setting for South Korea only. In particular, IRRA predicts South Korea REA in the case of IPI for $h > 3$ months and U for $h > 1$ month. Overall, our results from the analysis beyond U.S., document that only South Korea IRRA predicts South Korea REA. Our findings suggest that IRRA predicts future REA in the case where it is extracted from highly liquid options; South Korea KOSPI 200 index options have the highest trading volume in the panel of considered countries over the period under scrutiny and therefore the informational content of these market option prices is expected to be significant.

As a final robustness test, for the countries where IRRA is documented to predict REA both in- and out-of-sample (i.e. US and South-Korea), we have also examined whether this

¹⁰Application of the Kelly and Pruitt (2015) 3PRF is not possible in the case of the additional countries because an analogous to McCracken and Ng (2016) large macroeconomic dataset is not available for them.

finding prevails once IRRA is extracted from longer maturity (three-month) options.^{11, 12} The results, presented in the online appendix, show that U.S. IRRA does not predict U.S. REA neither in-sample nor out-of-sample when IRRA is estimated from options with three months to maturity. This is in line with the intuition that the prices of shorter maturity options may be more informative than those of longer maturity ones because the former are more liquid. On the other hand, South Korea IRRA predicts South Korea REA in-sample across most forecasting horizons; an increase (decrease) in IRRA predicts a decrease (increase) in REA. Notably, South Korea IRRA predicts South Korea REA out-of-sample, too. The ability of IRRA to predict REA when it is extracted from longer maturity in the case of South Korea comes as no surprise since the KOSPI 200 index options are the most actively traded contracts. For instance, when we extract IRRA using three-month constant maturity risk-neutral moments, we consider KOSPI 200 index option with 7 to 270 days to maturity that have a trading volume equal to 319 million in 2014 (versus 71 million for the respective S&P 500 index options in the U.S.).

6. Explaining empirical evidence: A production model

We take the textbook version of the real business cycle model presented in Miao (2014), and augment it with shocks to the coefficient of RRA. We show that this very standard and parsimonious production economy produces the same predictive high-frequency relation between risk aversion and future REA that we have identified in the empirical analysis.

¹¹Typically, the literature uses short maturity options to extract IRRA for two reasons. First, the liquidity is greater in shorter maturity options and hence, the informational content of IRRA is expected to be richer. Second, from a theoretical perspective, utility is defined over real consumption. The extraction of IRRA implicitly assumes that utility is defined over end-of-period wealth, i.e. nominal consumption rather than real consumption. In the latter case, researchers proxy wealth (i.e. dollar-denominated consumption) by a stock index like the S&P 500 for the case of U.S. To be accurate though, nominal consumption (i.e. wealth in this case) should have been converted to real consumption by taking inflation into account. Such a conversion is not possible though since IRRA is extracted from option prices. However, this is not a serious issue in the case where short horizons, i.e. short maturity options are considered (see Ait-Sahalia and Lo, 2000, for an excellent discussion on this).

¹²IRRA cannot be extracted from options with maturities longer than three months because of data constraints due to the lower liquidity in the longer maturity options. For instance, in the case of South Korea, we can estimate risk-neutral moments for six-months time-to-maturity only after 2014.

6.1. The economic environment

Households

The economy is populated by an infinitely lived representative household endowed with one unit of time in each period. Time is divided between work N_t and leisure L_t , so that $N_t + L_t = 1$. The household derives utility from the consumption good C_t and disutility from the fraction of time spent working. It maximizes expected (discounted) lifetime utility with respect to $\{C_t, N_t\}_{t=0}^{\infty}$

$$\max_{\{C_t, N_t\}_{t=0}^{\infty}} E_t \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t$$

where

$$\mathcal{U}_t = \frac{(C_t - h\bar{C}_{t-1})^{1-\gamma_t}}{1-\gamma_t} - \chi \frac{N_t^{1+\phi}}{1+\phi}, \quad (10)$$

$\mathcal{U}_t$ is the utility function and β is the subjective discount factor. We assume that $\mathcal{U}_t$ is separable over time and over consumption versus labour choices. $\mathcal{U}_t$ posits that households enjoy utility from the level of their own consumption C_t adjusted for habits, which in turn depends on aggregate consumption $\bar{C}_{t-1}$. γ_t is a variable driving time-variation in risk aversion and $h \in [0, 1)$ is a parameter governing external habits.¹³ ϕ governs the (Frisch) elasticity of labour supply to the real wage, and χ is a scale parameter to be assigned in the calibration. We assume that γ_t follows an autoregressive stochastic process of order one parameterized in logs:

$$\ln \gamma_t = \ln \gamma + \rho (\ln \gamma_{t-1} - \ln \gamma) + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma^2), \quad (11)$$

where ϵ_t is an exogenous innovation to risk aversion.

The household's relative risk aversion is given by

$$RRA_t = -\frac{C_t (\partial^2 \mathcal{U}_t / \partial C_t^2)}{\partial \mathcal{U}_t / \partial C_t} = \gamma_t \frac{C_t}{C_t - h\bar{C}_{t-1}}, \quad (12)$$

¹³We have introduced habits in consumption following a suggestion by the referee. This feature of the model amplifies the transmission of shocks quantitatively, but it does not affect the qualitative results of the analysis compared to the case where there are no habits. For recent theoretical and empirical analysis of the impact of habit formation on individual consumption behaviour see Baucells and Sarin (2010) and Acland and Levy (2015). Christoffersen et al. (2017) also use habits in consumption to explain stylized facts of credit spreads and their relation with option markets.

where $(C_t - h\bar{C}_{t-1})/C_t$ is the consumption surplus ratio. Notice that the empirical IRRA time series extracted from option prices is consistent with the presence of habits because its time variation may reflect the time variation in the consumption surplus ratio in equation (12).

The household receives a real wage W_t in exchange for supplying labour services and accumulates physical capital, K_t which rents to the firms at the net rate of return R_{t-1} . Capital accumulation follows the law of motion:

$$K_{t+1} = (1 - \delta) K_t + I_t, \quad (13)$$

where I_t denotes investment, δ is a constant rate of depreciation and K_t is predetermined at time t .

The intertemporal problem of the household is to maximize current and future expected utility (equation (10)) subject to the budget constraint

$$C_t + I_t = R_{t-1}K_t + W_tN_t,$$

and the law of motion for capital [equation (13)] where R_{t-1} is the real return on capital at $t - 1$. The first order conditions for dynamic optimality with respect to C_t , N_t , and K_{t+1} deliver a standard Euler equation:

$$\lambda_t = \beta E_t \lambda_{t+1} [R_t + 1 - \delta], \quad (14)$$

where λ_t denotes the marginal utility of consumption

$$\lambda_t = (C_t - h\bar{C}_{t-1})^{-\gamma_t},$$

and an equation for labour supply,

$$\lambda_t W_t = \chi N_t^\phi.$$

The above expression equalizes the marginal disutility from work χN_t^ϕ to the return from a

marginal increase in labor supply in utility units, $\lambda_t W_t$.

Firms

A representative perfectly competitive firm produces a homogeneous good Y_t using a standard Cobb-Douglas technology

$$Y_t = AK_t^{1-\alpha} N_t^\alpha. \quad (15)$$

We interpret the level of production Y_t as the theoretical analogue of the various proxies of real economic activity used in the empirical analysis. We assume that total factor productivity A is constant because we are only interested in the dynamics generated by the shock to risk aversion and hence we abstract from technology shocks. At every time t , firms minimize the cost of their inputs subject to the production technology in equation (15). The markets for capital and labour are assumed to be perfectly competitive which implies that the real return on capital and the real wage equal the marginal product of capital and labour, respectively:

$$R_t = (1 - \alpha) AK_t^{-\alpha} N_t^\alpha, \quad (16)$$

$$W_t = \alpha AK_t^{1-\alpha} N_t^{\alpha-1}. \quad (17)$$

Equilibrium

The competitive equilibrium is a sequence of quantities $\{C_t, N_t, K_{t+1}, I_t, Y_t\}_{t=0}^\infty$, and prices $\{R_t, W_t\}_{t=0}^\infty$ such that 1) given the prices and the exogenous stochastic process for γ_t , the vector of quantities satisfies the household's conditions for dynamic optimality, i.e. the Euler equation

$$(C_t - hC_{t-1})^{-\gamma_t} = E_t \beta (C_{t+1} - hC_t)^{-\gamma_{t+1}} [(1 - \alpha) K_{t+1}^{-\alpha} N_{t+1}^\alpha + 1 - \delta], \quad (18)$$

and the labour supply equation

$$W_t = \chi N_t^\phi (C_t - hC_{t-1})^{\gamma_t}, \quad (19)$$

the feasibility constraint

$$Y_t = C_t + I_t,$$

the production function (15) and the law of motion for capital (13);

2) The price system solves the firm's first order conditions (16) and (17);

3) The exogenous stochastic process for the coefficient of risk aversion obeys equation (11).

Given that the equilibrium of the model has no closed form solution, we solve it numerically as follows. First, we assign parameter values to pin down the steady-state of the model. Then, we take a second order approximation of the model's equations around the steady state. Finally, we solve for the policy functions using the Kim et al. (2008) algorithm. See Appendix E for an explanation of the model's solution.

6.2. Calibration

The deterministic steady-state of the model can be solved in closed form by assigning parameter values in a particular order (see Miao, 2014). Parameter values are calibrated so as to match key statistics for the post-war U.S. economy, under the conventional normalization that one period in the model corresponds to a quarter. The parameter values and calibration targets are reported in Table 6. Given that the calibration is standard and follows Miao (2014) very closely, we relegate the details to Appendix E.

Instead, here we devote attention to the calibration of the parameters governing the stochastic process for γ_t , which marks the departure from the literature. We assign values to γ , ρ and σ by simulating the model's RRA via equation (12) so that we match the mean, standard deviation and autocorrelation of the simulated RRA_t with the empirical mean, standard deviation and autocorrelation of the U.S. IRRA series estimated in Section 3.¹⁴ We perform the simulation over 100,000 quarters by drawing 100,000 respective ϵ_t . We perform the matching by a trial and error iterative approach. Finally, we note that the habit parameter is set to 0.6, which is in line with the habit estimates in Christiano et al. (2005).

To assess whether the proposed model explains the empirically documented relation bet-

¹⁴The IRRA series in Section 3 is estimated at monthly frequencies and hence, it needs to be converted into a quarterly frequency for the purposes of the calibration. The mean, standard deviation and autocorrelation of this quarterly series are 5.8, 1.38 and 0.966 respectively.

ween IRRA and future REA, in the next Section we will (i) examine the impulse responses of the endogenous variables to a shock in γ_t , and (ii) simulate time series of γ_t and Y_t and examine their predictive relation.

6.3. Results and discussion

Inspecting the mechanism through impulse responses

We inspect impulse responses of the model's endogenous variables to a negative exogenous innovation to risk-aversion. Specifically, we perturb the steady-state equilibrium *once* with a single innovation ϵ_t at time t which generates a deviation of γ_t at time t relative to its steady-state value at $t - 1$ (see Appendix E for details). This allows us to provide intuition for the mechanism by which shocks to risk aversion propagate to the real economy and explore the impact of an exogenous change in risk-aversion at time t on the future growth rates of production, i.e., $\ln Y_{t+h} - \ln Y_t$, for $h = 1, \dots, 40$ quarters. Figure 4 reports impulse responses to a one standard deviation negative innovation to γ_t . All variables in this figure are expressed in log deviations from the steady-state, with the exception of risk aversion, which is expressed in level deviations, and GDP growth in the last panel, which is measured in log deviations relative to the impact period t , i.e. $\ln Y_{t+h} - \ln Y_t$.

Figure 4 shows that a decrease in RRA yields a subsequent increase in Y_{t+h} over a number of subsequent quarters. This is in line with the previously provided empirical evidence. The impulse response function reveals the channel via which this causal effect occurs. The first panel of the figure shows that RRA_t drops as soon as the shock appears, subsequently exhibits hump-shaped dynamics due to habits, and then gradually returns to its long run average, as dictated by the mean reverting process in equation (11). The marginal utility of consumption $\lambda_t = (C_t - hC_{t-1})^{-\gamma_t}$, reported in the next panel, also falls following the decrease in γ_t .¹⁵ Intuitively, periods when γ_t is low are times when the marginal utility of consumption is low and hence consumption is valued less, so consumption falls, as reported in the third panel.

In turn, real wages rise. This is explained by the first order condition for labour supply

¹⁵Notice that the response of the marginal utility of consumption to the impact of an exogenous change to γ_t is $\partial (c_t - hc_{ss})^{-\gamma_t} / \partial \gamma_t = -\ln(c_t - hc_{ss}) c_t^{-\gamma_t} > 0$, where c_{ss} denotes the steady-state value of consumption. $\ln(c_t - hc_{ss})$ is negative and it increases in absolute value with the value of h because $0 < c_t - hc_{ss} < 1$ for plausible parameterizations.

in equation (19), $W_t = \chi \frac{N_t^\phi}{\lambda_t}$ given the initial fall in λ_t . Intuitively, for a given labour supply N_t , the disutility of work N_t^ϕ increases relative to the marginal utility of consumption, hence the workers require a higher real wage.

Given that real wage increases and capital at time t is predetermined, equation (17) shows that employment N_t must fall to equalize the marginal product of labour to the real wage (see the fourth panel in Figure 4). Equation (16) shows that the lower level of N_t generates a decrease in the marginal product of capital and hence in the rate of return on capital, as shown in the fifth panel of Figure 4. Intuitively, the return on capital falls because savings increase since consumption falls. Given that the model has no financial intermediaries, any increase in savings translates in an increase in investments as shown in the sixth panel in Figure 4. The rise of investment I_t at time t implies an increase in capital K_{t+1} at time $t + 1$ [equation (13)] and this leads to an increase in Y_{t+1} (last panel).

At this point, a remark is in order. At time t , GDP decreases as a response to the contemporaneous shock on RRA (see the eighth panel), reflecting the fall in employment. This is because $Y_t = AK_t^{1-\alpha} N_t^\alpha$. Given that K_t is predetermined at time t , the fall in employment directly translates into an fall in GDP. However, capital starts increasing thereafter because investment increases as we described. This delayed increase in capital taken together with the reversion of employment to its steady-state value, implies that output growth is positive between time t and $t + 1$ (see the last panel). In the following quarters, capital accumulation and the increase in employment continue driving output growth, leading to a negative relation between γ_t and GDP growth, $\ln(Y_{t+h}) - \ln(Y_t)$, which remains positive for various quarters h , extending far beyond the 1-year horizon considered in the empirical section.

Simulations

To provide further evidence that the model reproduces the predictive relation identified in the empirical analysis, we simulate time series for output Y_t and risk aversion RRA_t . We draw a random sequence of 100,000 innovations for ϵ_t , which leads to exogenous variation in γ_t according to equation (11). In turn, shocks to risk-aversion engenders fluctuations in all the endogenous variables, including Y_t . We collect a vector of 100,000 artificial observations

for both RRA_t and Y_t , and run the regression:

$$\Delta Y_{t,t+h} = c + bRRA_t + \varepsilon_{t+h} \quad (20)$$

where $\Delta Y_{t,t+h}$ is measured as $\ln(Y_{t+h}) - \ln(Y_t)$ for $h = 1, 2, 3$ and 4 quarters. Table 7 reports the estimated RRA coefficient along with the Newey-West and IVX-Wald p -values.

We can see that the estimated coefficients are negative and significant at all horizons between one quarter and one-year. This confirms that the model reproduces the same predictive relation that we have identified in the empirical analysis of Section 4.

7. Conclusions

We propose a new predictor of U.S. real economic activity (REA), namely the U.S. representative agent's implied relative risk aversion (IRRA) extracted from S&P 500 index option market prices. We find that an increase (decrease) in U.S. IRRA predicts a decrease (increase) in future U.S. REA both in- and out-of-sample. Interestingly, we document that the predictive ability of market option prices for future REA is not confined to the U.S. economy. We extract IRRA from the highly liquid South Korea KOSPI 200 options market and we find that it predicts the South Korea future REA, too. On the other hand, we find that the IRRA extracted from less liquid option markets such as the German, Japanese and UK does not forecast the respective countries' future REA. Our results imply that the U.S. and South Korea IRRAs extracted from highly liquid index options should be added to the existing list of REA predictors for these two countries.

We explain the negative predictive relation between risk aversion and future REA by invoking a production economy model, where an exogenous fall in risk aversion generates protracted growth in REA relative to the impact period. While the model is consistent with the estimated predictive relation, there might as well exist alternative, non mutually exclusive mechanisms that give rise to the same predictive relation. For instance, it is possible that expectations of future growth in REA decrease current risk-aversion (Guiso et al., 2017). Under rational expectations, future expected REA equals, on average, future realized REA. In this

case, the estimated negative relation between current IRRA and future actual REA could reflect the causal effect of a change in expectations about future REA onto current IRRA. Disentangling the contribution of these different mechanisms to the estimated predictive relation falls beyond the scope of this paper, and we leave it for future analysis.

References

- Acland, D., and M. R. Levy (2015). Naivete, projection bias, and habit formation in gym attendance, *Management Science* 61, 146-160.
- Ait-Sahalia, Y., and A. W. Lo (2000). Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9-51.
- Andersen, T. G., and T. Bollerslev (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885-905.
- Aruoba, S. B., F.X. Diebold, and C. Scotti (2009). Real-time measurement of business conditions. *Journal of Business and Economic Statistics* 27, 417-427.
- Baba, N., and I. Shim (2011). Dislocations in the Won-Dollar swap markets during the crisis of 2007-2009, Bank of International Settlements, Working Paper.
- Bakshi, G., N. Kapadia, and D. Madan (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101-143.
- Bakshi, G., and D. Madan (2006). A theory of volatility spreads, *Management Science* 52, 1945-1956.
- Bakshi, G., G. Panayotov, and G. Skoulakis (2011). Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios, *Journal of Financial Economics* 100, 475-495.
- Bakshi, G., G. Panayotov, and G. Skoulakis (2012). The Baltic dry index as a predictor of global stock returns, commodity returns, and global economic activity, Chicago Meetings Paper, AFA.
- Bakshi, G., X. Gao, and A. Rossi (2014). A better specified asset pricing models to explain the cross-section and time-series of commodity returns, University of Maryland, Working Paper.

- Barone-Adesi, G., L. Mancini, and H. Shefrin (2014). Sentiment, risk aversion and time preference, University of Lugano, Working Paper.
- Baucells, M., and R. K. Sarin (2010). Predicting utility under satiation and habit formation, *Management Science* 56, 286-301.
- Bliss, R. R., and N. Panigirtzoglou (2004). Option-implied risk aversion estimates, *Journal of Finance* 59, 407-446.
- Boldrin, M., L. Christiano, and J. Fisher (2001). Habit persistence, asset returns and the business cycle, *American Economic Review* 91, 149-166
- Buss, A., and G. Vilkov (2012). Measuring equity risk with option-implied correlations, *Review of Financial Studies* 25, 3113-3140.
- Campbell J. Y., and S.B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509-1531.
- Campbell J. Y., and M. Yogo (2006). Efficient tests of stock return predictability, *Journal of Financial Economics* 81, 27-60.
- Carr, P., and L. Wu (2009). Variance risk premiums, *Review of Financial Studies* 22, 1311-1341.
- Christiano, L. J., M. Eichenbaum, and C.L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy, *Journal of Political Economy* 113, 1-45.
- Christoffersen, P., D. Du and R. Elkamhi (2017). Rare disasters, credit, and option market puzzles, *Management Science* 63, 1341–1364.
- Christoffersen, P., B.Y. Chang, and K. Jacobs (2013), Forecasting with option implied information. In *Handbook of Economic Forecasting*, Volume 2, Part A, 2013, Pages 581–656, Elliott and Timmermann (Eds).
- Chang, B. Y., P. Christoffersen, K. Jacobs, and G. Vainberg (2012). Option-implied measures of equity risk, *Review of Finance* 16, 385-428.

- Chetty, R., A. Guren, D. Manoli, and A. Weber (2012). Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities, *NBER Macroeconomics Annual* 27, 1-56.
- Chiu, C.W. (2010). Funding liquidity, business cycles and asset pricing in the U.S., Working Paper, Duke University.
- Daskalaki C., A. Kostakis, and G. Skiadopoulos. (2014). Are there common factors in individual commodity futures returns?, *Journal of Banking and Finance* 40, 346-363.
- DeMiguel, V., Y. Plyakha, R. Uppal, and G. Vilkov (2013). Improving portfolio selection using option-implied volatility and skewness, *Journal of Financial and Quantitative Analysis* 48, 1813-1845.
- Duan, J. C., and W. Zhang (2014). Forward-looking market risk premium, *Management Science* 60, 521-538.
- Easley, D., M. O'Hara, and P. Srinivas (1998). Option volume and stock prices: Evidence on where informed traders trade, *Journal of Finance* 53, 431-465.
- Estrella, A., and G. A. Hardouvelis (1991). The term structure as a predictor of real economic activity, *Journal of Finance* 46, 555-576.
- Fama, E. F., and K. French (1996). Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55-84.
- Faust, J., S. Gilchrist, J.H. Wright, and E. Zakrajšek (2013). Credit spreads as predictors of real economic activity: A Bayesian model-averaging approach, *Review of Economics and Statistics* 95, 1501-1519.
- Giamouridis, D., and G. Skiadopoulos (2012). The informational content of financial options for quantitative asset management: A review. In *Handbook of Quantitative Asset Management*, 243-265, B. Scherer and K. Winston (Editors). Oxford University Press.
- Gourinchas, P-O., and M. Obstfeld (2012). Stories of the twentieth century for the twenty-first, *American Economic Journal: Macroeconomics* 4, 226-65.

- Guiso, L., Sapienza, P. and L. Zingales (2017). Time varying risk aversion, *Journal of Financial Economics*, forthcoming.
- Hansen, L. P. (1982). Large sample properties of Generalized Method of Moments estimators, *Econometrica* 50, 1029-1084.
- Hong, H., and M. Yogo (2012). What does futures market interest tell us about the macroeconomy and asset prices?, *Journal of Financial Economics* 105, 473-490.
- Jackwerth, J. C., and M. Rubinstein (1996). Recovering probability distributions from option prices, *Journal of Finance* 51, 1611-1631.
- Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433-451.
- Jackwerth, J., (2004). Option-implied risk-neutral distributions and risk aversion, Research Foundation of CFA Institute.
- Jermann, U. (1998). Asset pricing in production economies, *Journal of Monetary Economics* 41, 257-275.
- Jiang, G. J., and Y. S. Tian (2005). The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305-1342.
- Kaltenbrunner, G., and L. Lochstoer (2010). Long run risk through consumption smoothing, *Review of Financial Studies* 23, 3190-3224
- Kang, B. J., and T. S. Kim (2006). Option-implied risk preferences: An extension to wider classes of utility functions, *Journal of Financial Markets* 9, 180-198.
- Kang, B. J., T. S. Kim, and S-J. Yoon (2010). Information content of volatility spreads, *Journal of Futures Markets* 30, 533-558.
- Kelly, B., and S. Pruitt (2015). The three-pass filter regression filter: A new approach to forecasting using many predictors, *Journal of Econometrics* 186, 294-316.

- Kim, I.Y. and S. Kim (2005). Is it important to consider the jump component for pricing and hedging short-term options?, *Journal of Futures Markets* 25, 989-1009.
- Kim, J, S. Kim, E. Schaumburg, C.A. Sims (2008). Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models, *Journal of Economic Dynamics and Control* 32, 3397-3414.
- Kostakis, A., N. Panigirtzoglou, and G. Skiadopoulos (2011). Market timing with option-implied distributions: A forward-looking approach, *Management Science* 57, 1231-1249.
- Kostakis, A., T. Magdalinos, and M. P. Stamatogiannis (2015). Robust econometric inference for stock return predictability, *Review of Financial Studies* 28, 1506-1553.
- Liew, J., and M. Vassalou (2000). Can book-to-market, size and momentum be risk factors that predict economic growth?, *Journal of Financial Economics* 57, 221-245.
- McCracken, M.W., and S. Ng (2016). FRED-MD: A monthly database for macroeconomic research, *Journal of Business and Economic Statistics* 34, 574-589.
- Merton, R. C. (1973). Theory of rational option pricing, *Bell Journal of Economics and Management Science* 4, 141-183.
- Miao, J. (2014). *Economic dynamics in discrete time*, MIT Press.
- Neumann, M., and G. Skiadopoulos (2013). Predictable dynamics in higher order risk-neutral moments: Evidence from the S&P 500 options, *Journal of Financial and Quantitative Analysis* 48, 947-977.
- Newey, W., and K. West (1994). Automatic lag selection in covariance matrix estimation, *Review of Economic Studies* 61, 631-653.
- Ng, S., and J.H. Wright (2013). Facts and challenges from the great recession for forecasting and macroeconomic modeling, *Journal of Economic Literature* 51, 1120-1154.

- Pan, J., and A. Poteshman (2006). The information in options volume for future stock prices, *Review of Financial Studies* 19, 871-908.
- Papanikolaou, D. (2011). Investment shocks and asset prices, *Journal of Political Economy* 119, 639-685.
- Rosenberg, J. V., and R. F. Engle (2002). Empirical pricing kernels, *Journal of Financial Economics* 64, 341-372.
- Ryu, D. (2015). The information content of trades: An analysis of KOSPI 200 index derivatives. *Journal of Futures Markets* 35, 201-221.
- Söderlind, P., and L. Svensson (1997). New techniques to extract market expectations from financial instruments, *Journal of Monetary Economics* 40, 383-429.
- Stock, J. H., and M. H. Watson (2003). Forecasting output and inflation: The role of asset prices, *Journal of Economic Literature* 41, 788-829.
- Stock, J. H., and M. H. Watson (2002a). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association* 97, 1167-1179.
- Stock, J. H., and M. H. Watson (2002b). Macroeconomic forecasting using diffusion indexes. *Journal of Business and Economic Statistics* 20, 147-162.
- Yashiv, E. (2016). Capital values and job values, *Review of Economic Dynamics* 19, 190-209.

Appendix A: Computation of the risk-neutral moments

We compute the S&P 500 risk-neutral moments from market option prices following Bakshi et al. (2003) methodology. The advantage of this methodology is that it is model-free because it does not require any specific assumptions for the underlying's asset price stochastic process.

Let $S(t)$ be the price of the underlying asset at time t , r the risk-free rate and $R(t, \tau) \equiv \ln[S(t + \tau)] - \ln S(t)$ the τ -period continuously compounded return. The computed at time t model-free risk-neutral volatility $[\sigma_{q,t}(\tau)]$, skewness $[\theta_{q,t}(\tau)]$ and kurtosis $[\kappa_{q,t}(\tau)]$ of the log-returns $R(t, \tau)$ distribution with horizon τ are given by:

$$\sigma_{q,t}(\tau) = \sqrt{E_t^Q [R(t, \tau)^2] - \mu(t, \tau)^2} = \sqrt{V(t, \tau)e^{r\tau} - \mu(t, \tau)^2} \quad (\text{A.1})$$

$$\begin{aligned} \theta_{q,t}(\tau) &= \frac{E_t^Q \left[\left(R(t, \tau) - E_t^Q R(t, \tau) \right)^3 \right]}{E_t^Q \left[\left(R(t, \tau) - E_t^Q R(t, \tau) \right)^2 \right]^{3/2}} \\ &= \frac{e^{r\tau} W(t, \tau) - 3\mu(t, \tau)e^{r\tau} V(t, \tau) + 2\mu(t, \tau)^3}{[e^{r\tau} V(t, \tau) - \mu(t, \tau)^2]^{3/2}} \end{aligned} \quad (\text{A.2})$$

$$\kappa_{q,t}(\tau) = \frac{E_t^Q \left[\left(R(t, \tau) - E_t^Q R(t, \tau) \right)^4 \right]}{\left\{ E_t^Q \left[\left(R(t, \tau) - E_t^Q R(t, \tau) \right)^2 \right] \right\}^2} \quad (\text{A.3})$$

where $V(t, \tau)$, $W(t, \tau)$ and $X(t, \tau)$ are the fair values of three artificial contracts (volatility, cubic and quartic contract) defined as:

$$V(t, \tau) = E_t^Q [e^{-r\tau} R(t, \tau)^2], \quad W(t, \tau) \equiv E_t^Q [e^{-r\tau} R(t, \tau)^3] \quad \text{and} \quad X(t, \tau) \equiv E_t^Q [e^{-r\tau} R(t, \tau)^4]$$

and $\mu(t, \tau)$ is the mean of the log return over the period τ defined as:

$$\mu(t, \tau) \equiv E_t^Q \{\ln(S_{t+\tau}/S_t)\} \approx e^{r\tau} - 1 - \frac{e^{r\tau}}{2} V(t, \tau) - \frac{e^{r\tau}}{6} W(t, \tau) - \frac{e^{r\tau}}{24} X(t, \tau)$$

The prices of the three contracts can be computed as a linear combination of out-of-the-money call and put options:

$$\begin{aligned}
V(t, \tau) = & \int_{S_t}^{\infty} \frac{2(1 - \ln(K/S_t))}{K^2} C(t, \tau; K) dK + \\
& + \int_0^{S_t} \frac{2(1 + \ln(S_t/K))}{K^2} P(t, \tau; K) dK
\end{aligned} \tag{A.4}$$

$$\begin{aligned}
W(t, \tau) = & \int_{S_t}^{\infty} \frac{6 \ln(K/S_t) - 3 \ln(K/S_t)}{K^2} C(t, \tau; K) dK + \\
& + \int_0^{S_t} \frac{6 \ln(S_t/K) + 3 \ln(S_t/K)}{K^2} P(t, \tau; K) dK
\end{aligned} \tag{A.5}$$

$$\begin{aligned}
X(t, \tau) = & \int_{S_t}^{\infty} \frac{12 [\ln(K/S_t)]^2 - 4 [\ln(K/S_t)]^3}{K^2} C(t, \tau; K) dK + \\
& + \int_0^{S_t} \frac{12 [\ln(S_t/K)]^2 + 4 [\ln(S_t/K)]^3}{K^2} P(t, \tau; K) dK
\end{aligned} \tag{A.6}$$

where $C(t, \tau; K)$ ($P(t, \tau; K)$) are the call and put prices with strike price K and time to maturity τ .

Equations (A.4), (A.5) and (A.6) show a continuum of out-of-the-money calls and puts across strikes is required to compute the risk-neutral moments. However, options trade for discrete strikes. We also need constant-maturity risk-neutral moments to extract IRRA corresponding to a 30-days constant horizon. We estimate the constant-maturity risk-neutral moments of the S&P 500 returns distribution in line with Jiang and Tian (2005), Carr and Wu (2009), Chang et al. (2013), and Neumann and Skiadopoulos (2013). First, we keep only maturities for which there are at least two out-of-the-money puts and two out-the-money calls. In addition, to ensure that the options span a wide range of moneyness regions, we also discard maturities for which there are no options with deltas below 0.25 and above 0.75; we calculate deltas by using the implied volatility of the closest-to-the-money option. Next, for

any given maturity and date t , we convert strikes into moneyness ($K/S(t)$) levels. Then, we interpolate using a cubic spline across the implied volatilities provided by Optionmetrics to obtain a continuum of implied volatilities as a function of moneyness levels. To compute constant maturity moments, for each moneyness level, we interpolate across implied volatilities in the time dimension using a cubic spline. We keep the implied moments with a constant one-month maturity. Finally, we convert implied volatilities to option prices using Merton's (1973) model. Using trapezoidal approximation, we compute the prices for the three contracts which we then use to compute the risk- neutral moments.

Appendix B: Construction of the commodity factors

We construct the three commodity risk factors (hedging-pressure, basis and momentum risk factors) along the lines of Daskalaki et al. (2014).

B.1 Hedging-pressure risk factor

We denote as $HP_{i,t}$ the hedging pressure for any commodity i at time t defined to be the number of short hedging positions minus the number of long hedging positions, divided by the total number of hedgers in the respective commodity market. Risk averse speculators take futures positions only if they receive compensation and they share the price risk with hedgers (hedging pressure hypothesis). So, if $HP_{i,t}$ is positive (negative), hedgers are net short (long) in the futures contract. Speculators are willing to take the long (short) position only if they receive a positive risk premium. At any given month t , we construct a zero cost mimicking portfolio in line with the above strategy. First, we calculate $HP_{i,t}$ for each futures contract. Then, we construct two portfolios: portfolio H that contains all commodities with positive HP and portfolio L that contains all commodities with negative HP . At time t , we construct the high minus low HP risk factor by going long in portfolio H and short in portfolio L. Finally, at time $t + 1$, i.e. the next month, we calculate the realized mimicking portfolio return realized over t to $t + 1$. We construct a time series of our factor by repeating the above steps throughout our sample.

B.2 Momentum risk factor

According to Gordon et al. (2012), a negative shock to inventories leads to an increase in prices which is then followed by a short period of high expected futures returns for the respective commodity. This occurs because demand exceeds the supply for the commodity for that period and thus a price momentum is created. At any point in time t , we construct two portfolios: portfolio H that contains all commodities with positive prior 12-month average return and portfolio L that contains those with negative prior 12-month average return. Then at t , we construct the high minus low momentum zero-cost risk factor, by going long in portfolio H and short in portfolio L. Finally, at time $t + 1$, i.e. the next month, we calculate the realized

mimicking portfolio return realized over t to $t + 1$. We construct a time series of our factor by repeating the above steps throughout our sample.

B.3 Basis risk factor

According to the theory of storage, a positive basis is associated with low inventories for any given commodity. In addition, Gordon, et al. (2012) find that a portfolio of commodities with a high basis outperforms the portfolio of commodities with a low basis. At any point in time t , we construct two portfolios: portfolio H that contains all commodities with positive basis and portfolio L that contains all commodities with negative basis. Then, we construct the zero-cost high minus low basis risk factor by going long in portfolio H and short in portfolio L. Finally, at time $t + 1$, i.e. the next month, we calculate the realized mimicking portfolio return realized over t to $t + 1$. We construct a time series of our factor by repeating the above steps throughout our sample.

Appendix C: The IVX-Wald test (Kostakis et al., 2015)

C.1 The IVX estimator

Consider the following predictive regression:

$$y_{t+1} = c + Ax_t + \varepsilon_{t+1} \quad (\text{C.1})$$

where A is a $(m \times r)$ coefficient matrix and:

$$x_{t+1} = R_n x_t + u_{t+1} \quad (\text{C.2})$$

with $x_t = (x_{1t}, x_{2t}, \dots, x_{rt})$ being the vector of predictors employed in (C.1), $R_n = I_r + \frac{C}{n^\alpha}$ for some $\alpha \geq 0$, $C = \text{diag}(c_1, \dots, c_r)$ and n being the sample size. The IVX methodology does not require a-priori knowledge of the predictors' degree of persistence. In fact, it allows for various classes of persistence through the autocorrelation matrix R_n ; the accommodated classes of persistence vary from purely stationary ($c_i < 0$ for all i and $\alpha = 0$) to purely non-stationary ($C = 0$ or $\alpha > 1$).

We estimate equation (C.1) via two-stage least squares based on the near-stationary instruments $\tilde{z}_t$ and not the initial predictors x_t :

$$\begin{aligned} \tilde{A}_{IVX} &= \underline{Y}' \tilde{Z} \left(\underline{X}' \tilde{Z} \right)^{-1} \\ &= \sum_{t=1}^n (y_t - \bar{y}_n) \tilde{z}'_{t-1} \left[\sum_{j=1}^n (x_j - \bar{x}_{n-1}) \tilde{z}'_{j-1} \right]^{-1} \end{aligned} \quad (\text{C.3})$$

where $\bar{y}_n = 1/n \sum_{t=1}^n y_t$, $\bar{x}_{n-1} = 1/n \sum_{t=1}^n x_{t-1}$, $\tilde{Z} = (\tilde{z}'_0, \dots, \tilde{z}'_{n-1})$ is the instrument matrix, and $\underline{Y} = (Y'_1, \dots, Y'_n)$ and $\underline{X} = (X'_0, \dots, X'_{n-1})$ are the demeaned predictive regression matrices; we take the demeaned predictive regression matrices because we allow for a constant in the predictive regression given in equation (C.1). Following Kostakis et al. (2015), we choose $C_Z = -I_r$ and $\beta = 0.95$.

The intuition behind the IVX methodology is to construct an instrumental variable with a known degree of persistence from the initial predictors x_t which has an unknown degree of persistence. Once we have done that, we apply standard instrumental variable estimation.

To construct the near-stationary instrumental variable $\tilde{z}_t$, we first estimate equations (C.1) and (C.2) with ordinary-least squares. Then, we construct $\tilde{z}_t$, initialized at $\tilde{z}_0 = 0$, as follows:

$$\tilde{z}_t = R_{nz}\tilde{z}_{t-1} + \Delta x_t \quad (\text{C.4})$$

where $R_{nz} = I_r + \frac{C_z}{n^\beta}$ is an artificial autoregressive matrix with specified persistence, $\beta \in (0, 1)$ and $C_z < 0$.

C.2 The IVX-Wald test

We test for the predictive ability of x_{it} , i.e. we test the null hypothesis:

$$H_0 : Hvec(A) = 0$$

where H is a known $r \times r$ matrix whose (i, i) entry is one and the remaining entries are zero, i.e. we test for the significance of each predictor separately.

The IVX-Wald test statistic for testing the H_0 is:

$$W_{IVX} = \left(Hvec\tilde{A}_{IVX} \right)' Q_H^{-1} \left(Hvec\tilde{A}_{IVX} \right) \stackrel{H_0}{\sim} \chi^2(1) \quad (\text{C.5})$$

where:

$$Q_H = H \left[(\tilde{Z}'\underline{X})^{-1} \otimes I_m \right] \mathbb{M} \left[(\underline{X}'\tilde{Z})^{-1} \otimes I_m \right] H' \quad (\text{C.6})$$

$$\mathbb{M} = \tilde{Z}'\tilde{Z} \otimes \hat{\Sigma}_{\varepsilon\varepsilon} - n\bar{z}_{n-1}\bar{z}_{n-1}' \otimes \hat{\Omega}_{FM} \quad (\text{C.7})$$

$$\hat{\Omega}_{FM} = \hat{\Sigma}_{\varepsilon\varepsilon} - \hat{\Omega}_{\varepsilon u} \hat{\Omega}_{uu}^{-1} \hat{\Omega}_{\varepsilon u}' \quad (\text{C.8})$$

To calculate the test statistic in (C.5), we need to construct the following short-run and

long-run covariance matrices:

$$\hat{\Sigma}_{\varepsilon\varepsilon} = \frac{1}{n} \sum_{t=1}^n \hat{\varepsilon}_t \hat{\varepsilon}_t', \quad \hat{\Sigma}_{\varepsilon u} = \frac{1}{n} \sum_{t=1}^n \hat{\varepsilon}_t \hat{u}_t', \quad \hat{\Sigma}_{uu} = \frac{1}{n} \sum_{t=1}^n \hat{u}_t \hat{u}_t' \quad (\text{C.9})$$

$$\hat{\Lambda}_{uu} = \frac{1}{n} \sum_{i=1}^{M_n} \left(1 - \frac{i}{M_n + 1}\right) \sum_{t=i+1}^n \hat{u}_t \hat{u}_{t-i}', \quad \hat{\Omega}_{uu} = \hat{\Sigma}_{uu} + \hat{\Lambda}_{uu} + \hat{\Lambda}_{uu}' \quad (\text{C.10})$$

$$\hat{\Lambda}_{u\varepsilon} = \frac{1}{n} \sum_{i=1}^{M_n} \left(1 - \frac{i}{M_n + 1}\right) \sum_{t=i+1}^n \hat{u}_t \hat{\varepsilon}_{t-i}', \quad \hat{\Omega}_{\varepsilon u} = \hat{\Sigma}_{\varepsilon u} + \hat{\Lambda}_{u\varepsilon}' \quad (\text{C.11})$$

where $\hat{\varepsilon}_t$ and $\hat{u}_t$ are the ordinary least squares residuals from equations (C.1) and (C.2), respectively, and M_n is a bandwidth parameter satisfying $M_n \rightarrow \infty$ and $M_n/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$. Following Kostakis et al. (2015), we choose $M_n = n^{1/3}$; the choice of the bandwidth parameter does not affect the properties of the IVX-Wald test statistic.

Appendix D: The three-pass regression filter (Kelly and Pruitt, 2015)

The three-pass regression filter (3PRF) is a dimension reduction method. It identifies factors that are relevant to the variable that we wish to forecast (forecast target, y). These factors may be a strict subset of the factors driving the predictor variables (X). To extract the factors we use a set of proxies which are variables that are related to the forecast target.

To fix ideas, we consider the following variables. First, $y = (y_1, y_2, \dots, y_T)'$ is a $(T \times 1)$ vector of the forecast target where T is the number of time series observations in the in-sample period. Second, X is a $(T \times N)$ matrix of the standardized predictor variables where N is the number of predictors. We denote with x_{it} the (t, i) -th element of the X matrix, i.e. the t -th time series observation of the i -th predictor ($i = 1, 2, \dots, N$ and $t = 1 - h, 2 - h, \dots, T - h$). Third, Z is the $(T \times L)$ matrix of proxies, i.e. variables which are driven by target relevant factors. Note that L is the number of proxies. We denote with z_{lt} the (t, l) -th element of the Z matrix, i.e. the t -th time series observation of the l -th proxy ($l = 1, 2, \dots, L$ and $t = 1, 2, \dots, T$).

Following Kelly and Pruitt (2015), we extract the 3PRF factor using one proxy ($L = 1$), namely the forecast target ($z = y_T$). To fix ideas, standing at time T , we construct the h -month out-of-sample forecast as follows. First, we run N time-series regressions:

$$\begin{aligned} x_{i,T-h} &= \phi_{0,i} + z' \phi_i + \varepsilon_{i,t} \\ &= \phi_{0,i} + \phi_i y_T + \varepsilon_{i,t} \quad \text{for } i = 1, 2, \dots, N \end{aligned} \quad (\text{D.1})$$

Next, we retain the estimated $\hat{\phi}_i$ and we estimate cross-sectional regressions at times $t = 1, 2, \dots, T - h$ and at time T :

$$x_{it} = \gamma_{0,i} + \hat{\phi}_i' F_t + \epsilon_{i,t} \quad \text{for } t = 1, 2, \dots, T - h \text{ and } T \quad (\text{D.2})$$

This yields the factor estimates $\hat{F}_1, \hat{F}_2, \dots, \hat{F}_{T-h}$ and $\hat{F}_T$. Then, we use $\hat{F}_1, \hat{F}_2, \dots, \hat{F}_{T-h}$ to estimate the third-pass regression:

$$y_T = \beta_0 + \hat{F}_{T-h}' \beta + \eta_{t+1} \quad (\text{D.3})$$

Finally, we use the estimated coefficients and the estimated factor $\hat{F}_T$ to construct our forecast:

$$E_T(y_{T+h}) = \beta_0 + \hat{F}_T' \beta \tag{D.4}$$

Note that in line with Kelly and Pruitt (2015), we take care to avoid any look-ahead bias by using information up to time T to estimate the factor from equation (D.2) and to estimate β_0 and β from equation (D.3). In the latter case, this requires we estimate (D.3) using observations on the factor up to $T - h$ (i.e. up to F_{T-h}).

Appendix E: Model solution, impulse responses, simulations and calibration

Model solution

The solution to the full system of non-linear dynamic equations listed in the characterization of the competitive general equilibrium is a list of equations, called policy functions. These relate the vector of all current period endogenous variables $\vec{x}_t$ only to the current exogenous shock to risk-aversion γ_t and the past state of a subset of endogenous variables $\vec{x}_{t-1}^-$, called ‘state variables’. This subset $\vec{x}_{t-1}^-$ includes the variables whose value at time t is predetermined, like K_t , and the variables that appear with a lag, C_{t-1} and γ_{t-1} . So, for example, the policy function for consumption is a function $C_t = C(\vec{x}_{t-1}^-, \gamma_t; \Omega)$, where Ω is the set of parameter values to be assigned in the calibration stage.

Given that the model has no closed form solution, we solve it numerically as follows. First, we assign parameter values to pin down the deterministic steady-state of the model, which is the stationary point $\vec{x}_t = \vec{x}_{t-1} = \vec{x}$. Then, we approximate the model up to a second order approximation around the steady state. Finally, we solve for the policy functions using the Kim et al. (2008) algorithm.

Impulse responses

To calculate the impulse responses, we perturb the steady-state equilibrium *once* with a single innovation ϵ_t at time t which generates a deviation of γ_t at time t relative to its steady-state value at $t - 1$. Given γ_t , we obtain the value for the vector of endogenous variables $\vec{x}_t$ at time t via the policy functions $\vec{x}_t = f(\vec{x}_{t-1}^-, \gamma_t; \Omega)$, where $\vec{x}_{t-1}^-$ is the vector of steady-state values for the state variables. For the subsequent periods, we compute the values of γ_{t+h} , for $h = 1, 2, \dots, \mathcal{T}$ periods ahead. To this end, we take the exponent of γ_t in equation (11) and iterate forward, under the assumption that the realized $\epsilon_{t+h} = 0$ for $h = 1, 2, \dots, \mathcal{T}$. We iterate on the policy functions to simulate the dynamics of the endogenous variables $\vec{x}_{t+h} = f(\vec{x}_{t+h-1}^-, \gamma_{t+h}; \Omega)$ by using the obtained time series for γ_{t+h} , for $h = 1, 2, \dots, \mathcal{T}$.

Simulations

To obtain a simulated time series of γ_t and Y_t , we perturb the stationary equilibrium with a random *sequence* of N innovations to risk aversion, i.e., we produce a vector of $\{\epsilon_{t+h}\}_{h=0}^N$.

Iterating on the law of motion for γ_t [equation 11], we generate a time series for this exogenous variable. Given the values of $\{\gamma_{t+h}\}_{h=0}^N$, we iterate on the policy functions to produce a path for the vector of endogenous variables $\vec{x}_{t+h} = f(\vec{x}_{t+h-1}^-, \gamma_{t+h}; \Omega)$.

Calibration

Following the RBC literature, we assume that one period in the model corresponds to a quarter. In line with Miao (2014), we set the labour share of income $\alpha = 0.67$, and the discount factor $\beta = 0.99$ (which implies a real rate of return of about 1% per quarter). We set the depreciation rate of capital to the conventional value of $\delta = 0.025$ in line with estimates for the US economy (see Yashiv 2016). We normalize the long run value of total factor productivity, A to one. Solving the Euler equation (18) for the capital-labour ratio and evaluating it at the steady state yields


$$\frac{K}{N} = \left[\frac{1 - \beta(1 - \delta)}{\beta(1 - \alpha)} \right]^{-\frac{1}{\alpha}} \quad (\text{D.5})$$

Hence, we can recover the capital-labour ratio once we assign parameter values for α , β and δ . In turn, given the capital-labour ratio and α , we can compute the return to labour, W , and the return to capital $R + 1 - \delta$, using equations (16) and (17), respectively. We then normalize employment to the standard value of $N = 0.33$, which implies that households work 8 out of 24 hours a day. This allows us to compute the stock of capital by solving for K equation (D.5), the level of investment as $I = \delta K$, output Y by making use of the production function in (15) and consumption as $C = Y - I$. The marginal utility of consumption can be recovered as $[C(1 - h)]^{-\gamma}$ once we assign a value to the habit parameter h . We select a value of $h = 0.6$, which is in line with the habit estimates in Christiano et al. (2005). Finally, the value of the scale parameter χ in the utility function is implied by the intersection of the labour demand and supply equations (17) and (19), respectively:

$$\chi = \lambda W N^{-\phi}$$


where the inverse Frisch elasticity of labor supply, ϕ , is set to the value of 2, in line with the evidence in Chetty et al (2012).

Figure 1: Evolution of the U.S. implied risk aversion




The figure shows the evolution of the U.S. implied risk aversion (IRRA) over July 1998 - August 2015. We extract the IRRA time series via Kang et al. (2010) formula by performing a generalized-method-of-moments (GMM) rolling window estimation. We use a rolling window with size 30 months and three sets of instruments to obtain three respective U.S. IRRA time series.


Figure 2: Stability of U.S. IRRA coefficients




Panel A: Predicting IPI




Panel B: Predicting NFP




Panel C: Predicting RS



Panel D: Predicting HS




Panel E: Predicting CFNAI




Panel F: Predicting ADS

The figure shows the standardized U.S. IRRA coefficients from the estimated multiple predictor regression [equation (5)] for various U.S. real economic activity (REA) proxies and for a one-month horizon over the out-of-sample period October 2007 - August 2015. The REA proxies considered are: industrial production (IPI), non-farm payrolls (NFP), retail sales (RS, proxied by real retail sales), housing starts (HS), the Chicago Fed National Activity Index (CFNAI) and the Aruoba-Diebold-Scotti business conditions index (ADS). The multiple predictor model includes the lagged REA and implied relative risk aversion (IRRA) as predictors and is augmented by a set of control variables: term spread (TERM), default spread (DEF), TED spread (TED), Fama-French (1996) Small-Minus-Big factor (SMB), Fama-French (1996) High-Minus-Low factor (HML), Baltic Dry Index (BDI), forward variance (FV), hedging pressure commodity factor (HP), momentum commodity factor (MOM), basis commodity factor (BASIS), commodities open interest (OI). We estimate IRRA by the generalized-method-of-moments (GMM) with a 30-months rolling window using equation (3)


Figure 3: Evolution of the South Korea, UK, Japanese and German implied risk aversion




Panel A: South Korea IRRA



Panel B: UK IRRA




Panel C: Japanese IRRA



Panel D: German IRRA

The figure shows the evolution of the South Korea, UK, Japanese and German implied risk aversion (IRRA) over June 2006 - June 2015. For each country, we extract the IRRA time series via Kang et al. (2010) formula by performing a generalized-method-of-moments (GMM) rolling window estimation. We use a rolling window with size 30 months and three sets of instruments to obtain three respective IRRA time series for each country.

Figure 4: Production economy model: Impulse responses to a risk-aversion shock



The figure shows the impulse responses generated by the production economy model regarding the impact of a shock to risk-aversion for the model with habits on the set of the model's endogenous variables. All impulse responses are expressed in log deviations from the steady-state, except for relative risk aversion, which is expressed in level deviations and GDP growth in the last panel, which is measured in log deviations relative to the impact period t , i.e. $\ln Y_{t+h} - \ln Y_t, h = 1, 2, \dots, 40$ quarters.

Table 1: List of commodity futures

Sector	Commodities
Grains and Oilseeds	Corn Kansas Wheat Oats Soybean Meal Soybean Oil Soybeans Wheat
Energy	Crude Oil Heating Oil
Livestock	Feeder Cattle Pork Bellies Lean Hogs Live Cattle
Metals	Copper Gold Palladium Platinum Silver
Softs	Cocoa Coffee Cotton Sugar

Entries report the twenty-two commodity futures categorized in five broad sectors (grains and oilseeds, energy, livestock, metals and softs). These are used to construct the three Daskalaki et al. (2014) commodity-specific factors (hedging pressure, momentum and basis factors).

Table 2: In-sample prediction of U.S. REA with U.S. IRRA

	IPI	NFP	RS	HS	CFNAI	ADS
Panel A: Standardized IRRA coefficient						
$h = 1M$	-0.133*** (0.097) [0.075]	-0.131** (0.002) [0.012]	-0.148 (0.014) [0.107]	-0.253* (0.000) [0.003]	-0.123*** (0.024) [0.060]	-0.062*** (0.153) [0.087]
$h = 3M$	-0.097 (0.282) [0.169]	-0.120* (0.021) [0.008]	-0.245** (0.009) [0.022]	-0.401** (0.000) [0.015]	-0.064 (0.329) [0.729]	-0.050 (0.474) [0.676]
$h = 6M$	-0.133 (0.279) [0.151]	-0.156* (0.023) [0.002]	-0.310*** (0.006) [0.098]	-0.659* (0.000) [0.000]	-0.122 (0.302) [0.969]	-0.156 (0.241) [0.304]
$h = 9M$	-0.227* (0.146) [0.060]	-0.224*** (0.023) [0.068]	-0.377*** (0.007) [0.078]	-0.785* (0.000) [0.000]	-0.194 (0.204) [0.787]	-0.254 (0.143) [0.950]
$h = 12M$	-0.327* (0.066) [0.003]	-0.268** (0.023) [0.041]	-0.470 (0.001) [0.174]	-0.823* (0.000) [0.000]	-0.294 (0.078) [0.572]	-0.310 (0.068) [0.967]
Panel B: In-sample R^2						
$h = 1M$	0.205	0.725	0.138	0.205	0.617	0.856
$h = 3M$	0.482	0.797	0.246	0.283	0.579	0.640
$h = 6M$	0.419	0.759	0.409	0.483	0.458	0.452
$h = 9M$	0.414	0.685	0.469	0.597	0.356	0.379
$h = 12M$	0.460	0.645	0.532	0.720	0.320	0.351

Entries report results from the in-sample estimated multiple predictor regressions for various U.S. real economic activity (REA) proxies and for an h month horizon ($h = 1, 3, 6, 9$ and 12 months). The REA proxies considered are: industrial production (IPI), non-farm payrolls (NFP), retail sales (RS, proxied by real retail sales), housing starts (HS), the Chicago Fed National Activity Index (CFNAI) and the Aruoba-Diebold-Scotti business conditions index (ADS). The multiple predictor model includes the lagged REA and implied relative risk aversion (IRRA) as predictors and is augmented by a set of control variables: term spread (TERM), default spread (DEF), TED spread (TED), Fama-French (1996) Small-Minus-Big factor (SMB), Fama-French (1996) High-Minus-Low factor (HML), Baltic Dry Index (BDI), forward variance (FV), hedging pressure commodity factor (HP), momentum commodity factor (MOM), basis commodity factor (BASIS), and commodities open interest (OI). To construct our IRRA measure, we estimate equation (3) via the generalized-method-of-moments (GMM) with a 30-months rolling window. Panel A shows the standardized ordinary-least-squares (OLS) coefficient estimates, Newey-West (within brackets) and IVX-Wald (within squared brackets) p -values of each one of the predictors. One, two and three asterisks denote rejection of the null hypothesis of a zero coefficient based on the IVX-Wald test statistic at the 1%, 5% and 10% level, respectively. Panel B shows the in-sample adjusted R^2 for any given model. The sample spans July 1998 - August 2015.

Table 3: Out-of-sample predictability of U.S. REA

	IPI	NFP	RS	HS	CFNAI	ADS
Panel A: Out-of-sample R^2 from predictive regressions						
$h = 1M$	0.019	0.037	0.052	0.034	0.062	0.024
$h = 3M$	0.017	0.027	0.102	0.078	0.016	-0.033
$h = 6M$	-0.057	0.014	0.167	0.363	0.021	-0.009
$h = 9M$	-0.164	-0.042	0.164	0.572	-0.049	-0.042
$h = 12M$	-0.239	-0.108	0.171	0.640	-0.092	-0.113
Panel B: Out-of-sample R^2 from Kelly and Pruitt (2015) three-pass regression filter						
$h = 1M$	0.003	0.011	0.010	0.006	-0.012	0.013
$h = 3M$	0.005	0.016	0.011	0.038	0.010	0.013
$h = 6M$	0.010	0.016	0.026	0.067	0.017	0.018
$h = 9M$	0.014	0.019	0.032	0.101	0.028	0.028
$h = 12M$	0.018	0.022	0.042	0.097	0.030	0.016

Entries report the out-of-sample R^2 for U.S. Panel A shows the out-of-sample R^2 obtained from the predictive model in equation (5) versus the benchmark model that considers only lagged REA and the control variables as predictors. Panel B shows the out-of-sample R^2 obtained from Kelly and Pruitt (2015) three-pass regression filter in equation (9) applied to the set of variables consisting of IRRA and a large set of 135 macroeconomic variables compiled by McCracken and Ng (2016) versus the benchmark model that is the Kelly and Pruitt (2015) three-pass regression filter applied to the 135 McCracken and Ng (2016) macroeconomic variables. For each U.S. REA proxy, we estimate equations (5) and (9) for the respective full and benchmark models recursively by employing an expanding window; the first estimation sample window spans July 1998 - September 2007. At each point in time, we form $h = 1, 3, 6, 9, 12$ months-ahead REA forecasts. The positive (negative) sign of the out-of-sample R^2 indicates that the full model which includes IRRA as a predictor outperforms (underperforms) the restricted model.

Table 4: Construction of spread variables for South Korea, UK, Japan and Germany

Panel A: Term spread (TERM)			
South Korea	Ten-year government bond rate (<i>Bloomberg</i>)	<i>minus</i>	91-days certificate of deposit (CD) rate (<i>Korea Exchange</i>)
UK	Long-term interest rate (<i>OECD</i>)	<i>minus</i>	Short-term interest rate (<i>OECD</i>)
Japan	Long-term interest rate (<i>OECD</i>)	<i>minus</i>	Short-term interest rate (<i>OECD</i>)
Germany	Long-term interest rate (<i>OECD</i>)	<i>minus</i>	Short-term interest rate (<i>OECD</i>)
Panel B: Default spread (DEF)			
South Korea	AAA South Korea corporate bond yields (<i>Bloomberg</i>)	<i>minus</i>	BBB+ South Korea corporate bond yields (<i>Bloomberg</i>)
UK	Ten-year Euro swap rate (<i>Bloomberg</i>)	<i>minus</i>	Long-term interest rate (<i>OECD</i>)
Japan	BBB Japanese corporate bond yields (<i>Bloomberg</i>)	<i>minus</i>	AAA Japanese corporate bond yields (<i>Bloomberg</i>)
Germany	Ten-year Euro swap rate (<i>Bloomberg</i>)	<i>minus</i>	Long-term interest rate (<i>OECD</i>)
Panel C: TED spread (TED)			
South Korea	One-year KORIBOR (<i>Bloomberg</i>)	<i>minus</i>	One-year monetary stabilization South Korea bond rate (<i>Bloomberg</i>)
UK	British Pound Sterling LIBOR rate (<i>Bloomberg</i>)	<i>minus</i>	Short-term interest rate (<i>OECD</i>)
Japan	Japanese Yen LIBOR rate (<i>Bloomberg</i>)	<i>minus</i>	Short-term interest rate (<i>OECD</i>)
Germany	Euro LIBOR rate (<i>Bloomberg</i>)	<i>minus</i>	Short-term interest rate (<i>OECD</i>)

Entries define the spread variables (term, default and TED spread) used in the predictive regressions for South Korea, UK, Japan and Germany. The data source to construct each variable, is reported within parenthesis for each country and variable.

Table 5: Prediction of REA with IRRA: Evidence from other countries

Panel A: In-sample IRRA coefficient				
	South Korea		UK	
	IPI _{t+h}	U _{t+h}	IPI _{t+h}	U _{t+h}
$h = 1M$	0.016 (0.881)	0.049 (0.103)	-0.071** (0.048)	0.240* (0.000)
$h = 3M$	0.054 (0.791)	0.108* (0.000)	-0.097** (0.016)	0.263* (0.000)
$h = 6M$	-0.066** (0.015)	0.463* (0.000)	-0.243* (0.002)	0.392* (0.000)
$h = 9M$	-0.133** (0.012)	0.507* (0.006)	-0.326* (0.000)	0.527* (0.000)
$h = 12M$	-0.242* (0.001)	0.534* (0.000)	-0.352* (0.000)	0.653* (0.000)
	Japan		Germany	
	IPI _{t+h}	U _{t+h}	IPI _{t+h}	U _{t+h}
$h = 1M$	-0.084 (0.736)	0.018 (0.965)	0.025 (0.830)	-0.190 (0.706)
$h = 3M$	-0.157 (0.638)	0.070 (0.843)	0.165 (0.362)	-0.239 (0.300)
$h = 6M$	-0.245 (0.485)	0.161 (0.505)	0.244 (0.660)	-0.366 (0.567)
$h = 9M$	-0.177 (0.339)	0.295*** (0.069)	0.297 (0.756)	-0.419 (0.954)
$h = 12M$	-0.081 (0.937)	0.390 (0.869)	0.270 (0.881)	-0.459 (0.608)
Panel B: Out-of-sample R^2 from predictive regression				
	South Korea		UK	
	IPI	U	IPI	U
$h = 1M$	-0.026	-0.013	-0.005	-0.029
$h = 3M$	-0.016	0.056	-0.200	-0.104
$h = 6M$	0.028	0.208	-0.102	-0.089
$h = 9M$	0.551	0.378	0.210	-0.124
$h = 12M$	0.262	0.620	-0.051	-0.610

Panel A reports the standardized coefficient for IRRA obtained by estimating equation (5) by OLS for two South Korea, UK, Japan and Germany real economic activity (REA) proxies and for various forecasting horizons. The REA proxies considered are: industrial production (IPI) and unemployment rate (U). The horizons considered are $h = 1, 3, 6, 9$ and 12 months. The multiple predictor model includes the lagged REA and implied relative risk aversion (IRRA) as predictors and is augmented by a set of predictors as control variables: term spread (TERM), default spread (DEF), TED spread (TED), Baltic Dry Index (BDI) and forward variance (FV). To construct our IRRA measure, we estimate equation (3) with the generalized-method-of-moments (GMM) with a 30-months rolling window. We report the standardized ordinary-least-squares (OLS) coefficient estimates and IVX-Wald (within parentheses) p -values of each one of the predictors for any given model. One, two and three asterisks denote rejection of the null hypothesis of a zero coefficient based on the IVX-Wald test statistic at the 1%, 5% and 10% level, respectively. All predictor variables span June 2006 - June 2015.

Panel B reports the out-of-sample R^2 in the case of South Korea and UK for which in-sample predictability has previously been documented. For each REA proxy, we estimate equation (5) and the benchmark model recursively by employing an expanding window; the first estimation sample window contains observations spanning June 2006 - December 2008. At each point in time, we form h month-ahead REA forecasts ($h = 1, 3, 6, 9$ and 12 months). The benchmark model considers only lagged REA and the control variables as predictors.

Table 6: RBC model's calibrated parameters and implied steady-state values

Panel A: RBC model's calibrated parameters			
<i>Description</i>	<i>Parameter</i>	<i>Value</i>	<i>Source/Target</i>
Discount factor	β	0.99	1% interest rate
Capital depreciation rate	δ	0.025	Yashiv (2016)
Elasticity of output to labor	α	0.67	US labor share of income
Habits	h	0.6	Christiano et al. (2005)
Inverse Frisch elasticity	ϕ	2	Chetty et al. (2012)
Disutility of labor	χ	266.5	Share of hours worked 0.33
Coefficient of risk aversion	γ	2.241	Average IRRA
Autocorr. risk aversion shock	ρ	0.893	Autocorrelation IRRA
St. dev. risk aversion shock	σ	0.118	standard deviation IRRA
Panel B: Implied steady state values			
<i>Definition</i>	<i>Variable</i>	<i>Value</i>	
Consumption	C	0.76	
Investment	I	0.23	
Output	Y	0.99	
Hours (share)	N	0.33	
Real interest rate	R	0.0101	
Investment/capital ratio	I/K	0.025	
Capital/output ratio	K/Y	28.34	

Entries report the real business cycle (RBC) model's calibrated parameters (Panel A) and implied steady-state values (Panel B). Calibration is performed to the U.S. economy. We assign the values for the parameters γ , ρ and σ governing the stochastic process for γ_t , in equation (12) to match the mean, autocorrelation and standard deviation of the RRR_t series generated by simulating the model over 100,000 quarters with the empirical mean, standard deviation and autocorrelation of the IRRA time series estimated in Section 3.

Table 7: Predictive regressions using simulated data from the RBC model

	REA_{t+3}	REA_{t+6}	REA_{t+9}	REA_{t+12}
RRA_t	-0.004*	-0.008*	-0.012*	-0.016*
	(0.000)	(0.000)	(0.000)	(0.000)
	[0.000]	[0.000]	[0.000]	[0.000]

Entries report results from the predictive regression of output growth on RRA in equation (20). The regression has been performed on 100,000 simulated observations for output and risk aversion obtained by simulating the real business cycle (RBC) model presented in Section 6, using the parameter values reported in Table 6. We report the ordinary-least-squares (OLS) coefficient estimate, and Newey-West (within brackets) and IVX-Wald (within squared brackets) p -values. One asterisk denotes rejection of the null hypothesis of a zero coefficient on RRA based on the IVX-Wald test statistic at a 1% significance level.