Gold, Platinum, and Expected Stock Returns

Darien Huang*

April 2015

Abstract

I show that the ratio of gold to platinum prices (GP) reveals variation in risk and proxies for an important economic state variable. GP predicts future stock returns in the time-series and explains variation in average stock returns in the cross-section. GP outperforms existing predictors and similar patterns are found in international markets. Contrary to conventional views of gold as a hedge, gold prices *fall* in recessions, albeit by less than platinum prices. GP is persistent and significantly correlated with option-implied tail risk measures. An equilibrium model featuring recursive preferences, time-varying tail risk, and shocks to preferences for gold and platinum can quantitatively account for the asset pricing dynamics of equity, gold, and platinum markets, rationalize the return predictability, and explain why gold prices fall in bad times.

*Darien Huang (darienh@wharton.upenn.edu) is at the Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104. I thank my advisors Amir Yaron, Franklin Allen, and Ivan Shaliastovich - I am immensely grateful for their help, support, and advice. I also thank Andy Abel, Erik Gilje, Philipp Illeditsch, Petri Jylha, Marcin Kacperczyk, Andrew Karolyi, Mete Kilic, Robert Kosowski, Christian Opp, George Panayotov, Cesare Robotti, Nick Roussanov, Ronnie Sadka, Krista Schwarz, Rob Stambaugh, Luke Taylor, Jessica Wachter, and seminar participants at Wharton, Boston College, Cornell, HKU, HKUST, Imperial College, Indiana, Notre Dame, Rochester, Toronto, UNC, and York University for helpful discussions. All errors are my own.

1 Introduction

"As gold's unquenchable beauty shines like the sun, people have turned to it to protect themselves against the darkness ahead."

— Bernstein (2012), *The Power of Gold: The History of an Obsession*

Gold is one of the most important assets in financial markets and the global economy. As the author Peter Bernstein summarizes above, gold is viewed as two things: it is a consumption good (mostly jewellery) and it is also seen as something valuable in times of severe distress. Platinum, on the other hand, is a precious metal with similar uses as gold in consumption. Therefore, the ratio of gold to platinum prices should be largely insulated from shocks to consumption and jewelry demand, and should instead reveal variation in risk and proxy for an important economic state variable. In this paper, I examine this hypothesis by investigating three main questions.

First, I ask whether the ratio of gold to platinum prices (GP) predicts future stock returns in the time-series and explains variation in average stock returns in the cross-section. I show empirically that GP is a strong predictor of future stock returns. A one standard deviation increase in GP predicts a 6.4% increase in U.S. stock market excess returns over the following year. GP outperforms nearly all existing return predictors, and is robust to various econometric inference concerns highlighted in the literature. Gold and platinum are actively traded around the world, and similar patterns of stock return predictability are found in international markets. GP risk is priced in the cross-section of stock returns and commands a negative market price of risk.

After discussing the main empirical results, examining the mechanism which drives the results leads to my second question: Is gold a hedge? More specifically, do gold prices go up in bad times?[1] The answer - contrary to conventional wisdom - is no. Figure 1 plots real gold (top panel) and platinum (bottom panel) prices alongside stock market valuations and NBER recession indicators from 1975 - 2013.[2] We see in the data that gold prices fall in recessions, albeit by less than platinum prices. For example, in the 1981 - 1982 recession, real gold prices fell 32% peak to trough, and in the recent 2008 - 2009 financial crisis real gold prices fell 22%. Unlike index put options or VIX futures, gold futures would not have helped investors hedge downside risks during the crises. Not by coincidence, the real price of platinum fell by 39% and 59% over the same periods, respectively.

To the extent that some investors view shocks to gold prices as short-lived, flight-to-liquidity phenomena, I find that this is not the case. Shocks to GP do not correlate with shocks to transient measures of liquidity risk such as the Pastor and Stambaugh (2003) factor, and instead have a much longer half-life. Furthermore, GP is significantly related to measures

[1] See e.g., Erb and Harvey (2013), Barro and Misra (2013).

[2] I focus exclusively on the post-gold standard era, where gold prices vary freely by a market mechanism. While the "Nixon shock" of 1971 temporarily suspended convertibility of U.S. dollars into gold at $35 per oz, a new peg was later put in place at $38 per troy oz, followed by $42.22. Gold convertibility was only completely abolished by November 1973 (Lannoye (2011)). Executive Order 6102, put in place by President Franklin Roosevelt in 1933, banned gold trading within the United States. This act was repealed by President Gerald Ford in 1974 and took effect on December 31st, 1974. See Public Law 93-373.

of economic tail risk including the slope of the implied volatility curve for S&P500 index options, and the Bakshi, Kapadia, and Madan (2003) model-free risk-neutral skewness.[3] These findings lead to my final question, which is whether an extension of the time-varying disaster risk model (Wachter (2013)), which features recursive preferences and stochastic disaster probabilities, can quantitatively explain the time variation and return predictability of GP while simultaneously accounting for the asset pricing dynamics of equity, gold, and platinum markets. The model is motivated by the fact that, under no arbitrage, investors are indifferent between buying gold or leasing gold in perpetuity (Barro and Misra (2013)).

I adopt a three-good model where agents derive utility from nondurable consumption as well as service flows from gold and platinum, which are non-depreciating durable goods with negligible outlays relative to nondurable consumption. In normal times, service flows from gold and platinum (which can be thought of as jewellery) complement nondurable consumption and are highly procyclical. However, when the probability of a consumption disaster is high, agents display an increased preference for gold relative to platinum. This is motivated by both historical and institutional reasons, since gold is viewed as financial collateral and is formally recognized as such by the Basel Accords.[4]

The countercyclical benefits to physical ownership of gold and platinum are modeled in reduced-form, using a pair of stochastic processes which are proportional to the probability of a consumption disaster; gold is calibrated to have greater countercyclical benefits than platinum, which is both consistent with the historical and institutional facts and also allows the model to rationalize the low gold lease rate and risk premium observed in the data.[5] In the model, GP is insulated from shocks to consumption since they affect gold and platinum prices equally. Increases in disaster probabilities raise risk premia, leading to higher discount rates and lower stock prices. Gold and platinum prices fall as well because of strong discount rate effects, although gold prices fall by less than platinum prices due to the higher countercyclical component of its service flow. As a result, GP is high when stock prices are low and the equity risk premium is high, giving GP the power to predict future stock market excess returns. The model quantitatively captures the key moments of gold and platinum returns, while remaining consistent with standard asset pricing moments such as the equity premium and risk-free rate. This is achieved by linking gold and platinum valuations to state variables in the time-varying disaster risk model, which suggests that gold and platinum prices can largely be explained by the same risk factors affecting stocks and bonds.

Barro and Misra (2013) study gold returns in a Lucas (1978) endowment economy with rare consumption disasters. The authors match the low gold risk premium using a high elasticity of substitution between gold service flow and nondurable consumption. This assumption is can be improved for two reasons. First, viewing gold as jewellery suggests a complementary rather than substitutable relationship (one cannot wear jewellery in place of

[3] Tail risk is also known as jump risk or disaster risk depending on the literature.

[4] See e.g., Basel I (1988), Basel II (2004), and Basel III (2012). Gold is also accepted as collateral by major derivatives exchanges and clearinghouses such as the CME and ICE Clear Europe, as well as large broker dealers such as JP Morgan.

[5] Platinum is not eligible collateral under the Basel Accords, central banks are not known to hold platinum reserves, and major financial institutions do not accept platinum as collateral.

consuming food, but jewellery is highly valued when food is plentiful). Second, optimality conditions reveal that the elasticity of substitution is inversely proportional to the degree of consumption leverage. Following analysis similar to Wachter (2013), substitutability results in the counterfactual prediction that gold lease rates fall (gold prices rise) when disaster probabilities increase.

This paper contributes to the literature on return predictability by demonstrating that GP, a model-free measure available in real-time, is robust to, and in most cases outperforms, existing forecasting variables including equity valuation ratios (in various forms), the default spread, term spread, inflation, implied cost of capital, consumption-wealth ratio, and variance premium.[6] The predictive power of GP is stable both out-of-sample and over sub-samples, which alleviates concerns raised by studies such as Goyal and Welch (2008), who show that many predictors such as valuation ratios have low forecasting performance out-of-sample and unstable forecasting ability over sub-samples.

This paper extends the growing literature on gold and gold lease rates. To my knowledge, Barro and Misra (2013) is the only other paper to value gold in an equilibrium model. Fama and French (1988) analyze the behavior of metals prices over the business cycle based on the Brennan (1958) theory of storage. While base metals such as aluminum and copper behave as the theory of storage predicts, precious metals such as gold seem unresponsive; Fama and French hypothesize that this is due to low storage costs for precious metals. Tufano (1996) studies risk management practices in the gold mining industry. Schwartz (1997), Casassus and Collin-Dufresne (2005), and Le and Zhu (2013) study gold lease rates (known as "convenience yields" in the commodities literature) using dynamic term structure models. Erb and Harvey (2013) examine various theories regarding gold returns, including whether gold prices appreciate when stock prices fall. The authors find that many of the largest S&P500 declines were associated with falling gold prices.

Finally, this paper draws on the literature examining the impact of heavy-tailed shocks to economic state variables on asset prices. Examples from the option pricing literature include Bates (2000), Duffie, Pan, and Singleton (2000), Pan (2002), and Broadie, Chernov, and Johannes (2007). Jurek (2014) discusses the impact of crash risk on currency carry trade returns. Examples from the general equilibrium literature include the rare disasters framework (Rietz (1988), Barro (2006), Gabaix (2012), Gourio (2012), Wachter (2013), Nowotny (2011), Seo and Wachter (2014)), as well as extensions of the Bansal and Yaron (2004) long-run risks framework incorporating jumps in economic fundamentals (Eraker and Shaliastovich (2008), Bansal and Shaliastovich (2011), Benzoni, Collin-Dufresne, and Goldstein (2011), Drechsler and Yaron (2011)).

The paper proceeds as follows; data sources are discussed as the relevant sections are presented. Section 2 presents the empirical results on stock return predictability, the cross-

[6] Valuation ratios include the price-dividend ratio, price-earnings ratio, and net payout yield. References include Campbell and Shiller (1988), Hodrick (1992), and Boudoukh, Michaely, Richardson, and Roberts (2007). References for predictability using business cycle variables include Lintner (1975), Campbell (1987), and Fama and French (1989). The implied cost of capital is studied by Pastor, Sinha, and Swaminathan (2008) and Li, Ng, and Swaminathan (2013). The consumption-wealth ratio (CAY) is from Lettau and Ludvigson (2001). References for variance premium predictability include Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011).

sectional evidence, and the relationship between GP and tail risk measures. Section 3 discusses key aspects of gold and platinum markets, focusing on sources of demand for each metal, the leasing markets, and return dynamics. Section 4 presents the model. Section 5 discusses the model calibration and simulation results. Section 6 concludes.

2 Empirical Results

2.1 Data Description

Gold and platinum prices are the monthly average of daily fixing prices from the London Bullion Market Association (LBMA) and London Platinum and Palladium Market (LPPM), respectively, from 1975 to 2013.[7] Platinum fixing prices are available from April 1990; prior to this, I use dealer prices from the U.S. Geological Survey.[8] The log GP ratio is calculated as the natural logarithm of the ratio of gold to platinum prices.[9] My measure of U.S. stock returns is the CRSP value-weighted index. The risk-free rate is the 1-month U.S. Treasury bill rate. I compare the performance of GP to various forecasting variables proposed in the literature.

- Price-Dividend Ratio ($\log PD$) is the log ratio of aggregate stock market price divided by the sum of the past twelve months of dividends. Dividends are computed from the difference between the CRSP value-weighted index return including and exluding dividends.

- Price-Earnings Ratio ($\log PE$) is the cyclically-adjusted log ratio of aggregate stock market price divided by past earnings, obtained from Robert Shiller's website.

- Net Payout Ratio ($\log PNY$) is the log ratio of total market capitalization divided by the sum of dividends, repurchases, and share issuance, as described in Boudoukh et al. (2007) and obtained from Michael Roberts's website. The series is available until December 2010.

- Implied Cost of Capital (ICC) is rate of return which solves the long horizon dividend discount model, constructed from I/B/E/S analyst earnings per share forecasts, as described in Li et al. (2013). The series starts from January 1977.

- Default Spread ($DFSP$) is the percentage difference in yield between Moody's Baa and Aaa rated corporate bonds and obtained from the Federal Reserve Bank of St. Louis (FRED) website.

- Term Spread ($TMSP$) is the percentage difference in yield between 10 year U.S. government bonds and 3 month U.S. Treasury bills and obtained from FRED.

[7] I use prices from the a.m. fixing, which is conducted at 9:45 a.m. GMT (for platinum) and 10:30 a.m. GMT (for gold).

[8] The results are nearly unchanged using platinum prices directly obtained from Platts, which is a large data vendor for the metals markets.

[9] I use the terms "GP" and "log GP" interchangeably unless otherwise noted.

- Inflation ($INFL$) is the log growth rate of the *Consumer Price Index (All Urban Consumers: All Items)*, in percentages, from FRED.
- Consumption-Wealth Ratio (CAY) is the Lettau and Ludvigson (2001) measure of the consumption-wealth ratio, obtained from Martin Lettau's website. Monthly observations are computed by interpolating the quarterly observations. The series is available until March 2013.
- Variance Premium (VRP) is the difference between model-free implied variance computed from S&P500 option prices (VIX^2) and realized variance computed from 5-minute tick data over the past 30 days. The data for the VIX is obtained from the CBOE website, and the data for realized variance is from Hao Zhou's website. The series starts from January 1990.

Figure 2 plots the time-series of GP (solid line) along with the price-dividend ratio (dashed line). The average level of GP is below zero; gold trades at a 20% discount to platinum on average, consistent with platinum being a much scarcer metal. GP is strongly countercyclical and peaks during times of economic and financial distress including all NBER recessions between 1975 - 2013, as well as the October 1987 stock market crash, 1998 Russian default and LTCM crisis, and 2011 U.S. debt ceiling crisis. Table 1 presents summary statistics for all the predictors. With the exception of the variance premium and inflation, all other predictors are quite persistent. The AR(1) coefficient for GP is 0.98, which is inside the unit circle. Formally, a Dickey and Fuller (1979) stationarity test rejects the null of a unit root for $\log GP$ at the 5% level.[10] The high persistence of GP is in contrast to the view that shocks to gold prices reflect transient phenomena. Innovations in GP are uncorrelated with Pastor and Stambaugh (2003) liquidity factor innovations.

GP is countercyclical and strongly negatively correlated with equity valuation ratios; GP is high when stock prices are low. The strong positive correlation between GP and the default spread suggests that GP is high when firms with low credit ratings are more likely to default, which raises the required yield on their corporate bonds. GP is positively correlated with ICC since the cost of capital for firms is high in adverse economic conditions. High values of CAY are associated with high risk premia, and accordingly we see a positive correlation between GP and CAY. GP is not correlated with $INFL$; this is expected, since inflation equally affects both the numerator and denominator of the GP ratio and cancels out.

2.2 Stock Return Predictability

My measure of U.S. stock returns is the CRSP value-weighted index. The risk-free rate is the 1 month U.S. Treasury bill rate. I compare the performance of GP to various forecasting variables proposed in the literature. Table 2 shows the main predictability result of the paper. I run the regression:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h}. \tag{1}$$

[10]The optimal number of lags is chosen based on the Ng and Perron (1995) sequential t-test.

Long-horizon returns are constructed from overlapping monthly returns. The top panel uses ordinary least squares regression with Newey and West (1987) HAC robust standard errors.[11] At the 1 month horizon, the degree of predictability is fairly low with an R^2 just above 1%; however, the estimated slope is statistically significant with a 2.82 t-statistic. We see similar patterns of predictability up to the 1 year horizon, which has an R^2 of 16.57%. The bottom panel uses the vector autoregression (VAR) framework as in Hodrick (1992), which is potentially more conservative for overlapping returns, although it imposes parametric assumptions. The point estimates are very similar, although the R^2 is lower (yet still very large at 10.89% for the 1 year horizon) using the VAR. For longer horizons of 2 to 5 years, the estimated coefficients are still significant although the magnitude is decreasing. The estimated coefficient on $\log GP$ for the one year horizon is 0.243, the standard deviation of $\log GP$ is 0.266, so a one standard deviation increase in $\log GP$ is associated with a 6.4% increase in U.S. stock market excess returns over the following year. For all horizons from 1 month to 5 years, the estimated slopes are statistically significant.

Table 3 shows the results of univariate predictability regressions for each of the predictors. For short horizon returns (1 and 3 months), only GP, and VRP are statistically significant at conventional levels, with ICC significant at the 10% level. At the intermediate 1 year horizon, GP, ICC, and VRP are strongly significant, while $TMSP$ and $INFL$ are marginally significant. At this frequency, GP has the highest R^2 of all predictors. For long horizon (e.g. 5 year) returns, GP is still significant, while valuation ratios, CAY, and $TMSP$ are also significant.

How does GP stack up against other predictors in a horse race? The regression is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_1 \log GP_t + \beta_2 X_t + \epsilon_{t+h} \tag{2}$$

where X_t is another predictor. Table 4 shows the results for 1 and 3 month horizons.[12] At short horizons, VRP is known to be a strong predictor, as seen in Table 3. This supports the findings of Bollerslev et al. (2009) and Drechsler and Yaron (2011) on more recent data. However, in Table 4 we see that GP is still significant, even after controlling for VRP. Similar results hold for the 3 month horizon. Table 5 shows the results for long horizon returns. In most cases, GP drives out the significance of the other predictor, with the exception of $INFL$ at the 1 year horizon and ICC, $TMSP$, and CAY (marginally) at the 5 year horizon. Including alternative predictors leaves the magnitude of the coefficient on GP largely unchanged. The evidence suggests that GP is robust to and in most cases outperforms other forecasting variables proposed in the literature.

Goyal and Welch (2008) argue that predictors such as the price-dividend ratio do not perform well out-of-sample. I test out-of-sample robustness using Out-of-Sample R^2. If

[11] I also conduct a robust check using Hodrick (1992) standard errors and confirm that the results are not sensitive to the choice of standard errors.

[12] For bivariate regressions as well, the results using the Hodrick (1992) VAR methodology are similar.

GP is a robust predictor, Out-of-Sample R^2 should be significantly greater than zero and similar to its in-sample counterparts. The statistic is given by:

$$R^2_{OS} = 1 - \frac{\sum_{k=1}^{T-m} (r^e_{m+k} - \widehat{r^e}_{m+k})^2}{\sum_{k=1}^{T-m} (r^e_{m+k} - \overline{r^e}_{m+k})^2}. \tag{3}$$

We can calculate R^2_{OS} using either an expanding window (use all data available from month 1 to month m, so the regression sample expands at each time step), or a rolling window of length m (use only the past m months of data at each time step). In both cases, I estimate equation (1) in the estimation period, compute the squared prediction error over the next period and increment my time step. An expanding window uses more available data, while a rolling window better accounts for potential time variation in the predictive relationship. I consider windows of length 120 months and 180 months to estimate betas, and predict the return in the next period. The p-values are from the Clark and West (2007) adjusted-MSPE statistic:

$$f_{t+1} = (r_{t+1} - \overline{r}_{t+1})^2 - \left[(r_{t+1} - \widehat{r}_{t+1})^2 - (\overline{r}_{t+1} - \widehat{r}_{t+1})^2\right]$$

which is regressed against a constant and the test is a one-sided test of whether $R^2_{OS} > 0$.

Table 6 shows the results from the out-of-sample analysis. With the exception of the 1 month horizon rolling 10 year window regressions, all other combinations of forecast horizons and methods give large, positive, and significant out-of-sample R^2 values. The pattern of R^2 as we increase the forecast horizon are similar between rolling and expanding methods. Goyal and Welch (2008) find that for predictors such as the price-dividend ratio, the predictive ability is diminished in out-of-sample tests. For GP, the out-of-sample R^2 are significantly greater than zero and similar to the in-sample R^2. Figure 3 shows the estimated slopes and 95% error bands for both rolling and expanding methods with either 120 or 180 month windows for 2 year-ahead predictive regressions. We see that the estimates are stable, never change signs, and are statistically significant in nearly all sub-samples. For comparison, Figure 4 plots the same sub-sample betas for log PD. We see a lot of variation in the estimated coefficients with numerous sign changes and weak statistical significance in sub-samples (95% intervals straddle zero). The evidence suggests return predictability by GP is robust both out-of-sample and over sub-samples.

In Appendix A1, I show that GP is robust to finite-sample bias (Stambaugh (1999)) and size distortions (Torous, Valkanov, and Yan (2004)). In Appendix A2, I show that realized utility gains are high for mean-variance investors using GP for portfolio allocation.

Gold and platinum are globally traded assets. This suggests that GP should also predict future stock returns in international markets. I run the same predictive regressions as in equation (1) using the MSCI World Index, which is a U.S. dollar denominated index composed of stocks from 23 Developed Markets countries covering approximately 85% of the free float-adjusted market capitalization in each country. Since the index is dollar denominated, I use the U.S. Treasury bill rate as the risk-free rate. Table 7 shows that the patterns of predictability are very similar to the U.S. results: high GP predicts high future excess returns, although the coefficients are somewhat smaller in magnitude than for U.S. returns. Since there may be some concern that the world portfolio consists of a large proportion of U.S. stocks, I also run the same predictability regressions for other developed countries. Panel B of Table 7 reports the results for the U.K., Switzerland, Japan, and Sweden. I use the MSCI country indices for each of these countries, denominated in the local currency. The risk-free rate is the local currency treasury bill rate. The results for the U.K., Switzerland, and Sweden are nearly the same as for the U.S., while Japan shows significant predictability in terms of the magnitude of estimated slopes, albeit smaller t-statistics (significant at the 10% level) and somewhat lower R^2. The results suggest that GP predicts future excess stock market returns for the U.S. market as well as international markets, which also mitigates potential concerns about data snooping (Ang and Bekaert (2007)).

2.3 Dividend Growth Predictability

I have argued that stock return predictability by GP is driven by time variation in risk premia and not from news about future dividend growth rates. Some may argue that platinum has a characteristic not shared by gold: it is demanded by the automotive industry for catalytic converters.[13] Is it possible that it is actually bad news about the future cash flows of car makers (GP is low when platinum is expensive, which is bad news for future cash flows of car makers) that drives the predictability through a cash-flow channel? I run standard dividend growth predictive regressions similar to Cochrane (2008) on real dividend growth rates (Δd_t) and real earnings growth rates (Δe_t):

$$\begin{aligned} \frac{12}{h}\sum_{i=1}^{h}\Delta d_{t+i} &= \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h} \\ \frac{12}{h}\sum_{i=1}^{h}\Delta e_{t+i} &= \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h}. \end{aligned} \tag{4}$$

The results in Table 8 show no evidence of dividend growth predictability by GP.[14] For dividend growth, none of the estimated slopes from 1 year to 5 year horizons are statistically different from 0, and the R^2 are all nearly zero. For earnings growth, the R^2 are slightly higher but the t-statistics suggest the slopes are not significantly different from zero. This is evidence that the predictability I document arises because of variation in risk premia rather than dividend growth.

2.4 GP and the Cross-Section of Stock Returns

I examine the implications of GP risk for the cross-section of stock returns. As seen earlier, GP is countercylical and increases in times of economic distress. Stocks with high, positive covariation with GP innovations are therefore a good hedge against adverse states of high economic risk and low asset valuations, which suggests that GP should command a negative

[13] I discuss autocatalyst demand of platinum in Section 3.
[14] The earnings data is from Robert Shiller's website.

market price of risk in the cross-section. I estimate the risk exposures (betas) for each asset $i = 1, ..., N$ from time-series regressions

$$R^e_{i,t+1} = c_i + \beta_{i,\Delta gp} \Delta \log GP_{t+1} + \epsilon_{i,t+1} \tag{5}$$

where $R^e_{i,t+1}$ is the excess return for portfolio i and $\Delta \log GP_{t+1} = \log GP_{t+1} - \mathbb{E}_t [\log GP_t]$ is the innovation in GP from an AR(1) model.[15] The slope coefficient $\beta_{i,\Delta gp}$ represents the portfolio exposure of asset i to GP risk. In order to estimate the cross-sectional market price of risk associated with GP, I run a cross-sectional regression of time-series average excess returns on the risk exposures

$$\mathbb{E}\left[R^e_{i,t+1}\right] = \text{cons} + \beta_{i,\Delta gp} \lambda_{\Delta gp} + \upsilon_i \tag{6}$$

which yields estimates of the market price of risk $\lambda_{\Delta gp}$. I use the standard cross-section of ten portfolios sorted on the book-to-market ratio and ten portfolios sorted on size as my test assets. The data is monthly from 1975 - 2013. Recall that GP is constructed without any information from equity markets, which rules out any mechanical relationship between GP risk and the cross-section of stock returns. Furthermore, the parsimonious one-factor model avoids many statistical issues present in asset pricing tests that can mechanically produce high explanatory power. Panel A of the Table 9 shows that the market price of GP risk is significantly negative. Panel B of the Table further shows that the portfolio returns are all significantly and negatively exposed to GP risk; equity prices fall contemporaneously when GP increases. The one-factor model featuring only GP risk can explain over 60% of the cross-sectional variation in average returns. Figure 5 graphically depicts the strong negative relationship between average excess returns and risk exposures (Panel A), and correspondingly the fit between realized and model-predicted excess returns (Panel B). The cross-sectional results suggest that investors are willing to pay a premium for assets which hedge against increases in GP; in other words, the high-risk states which investors dislike are those associated with high GP.

2.5 GP and Tail Risk

The evidence so far suggests that 1) GP is countercyclical and increases in times of economic distress, 2) GP positively predicts future stock market excess returns, 3) GP risk is negatively priced in the cross-section, and 4) GP is high when the default spread is high, which is when firms with low credit ratings have higher probability of default. A plausible interpretation consistent with these results is that GP captures tail risk in the economy. This is broadly consistent with the findings of Manela and Moreira (2014), who use machine learning techniques to quantify tail risk (disaster concerns) from newspaper headlines: "gold" is one of the top words which explains variation in investors' tail risk concerns. GP is persistent, which is consistent with the evidence of persistent tail risk in Kelly and Jiang (2014). Options are an ideal way to measure tail risk because their convex payoff structures contain rich information about the tail distribution of returns. I extract tail risk measures from options markets and investigate the association between GP and tail risk.

[15] The results using first differences are nearly identical.

Out-of-the-money (OTM) index put options protect against stock market crashes. The slope of the implied volatility curve, defined as the implied volatility of an OTM put minus the implied volatility of an at-the-money (ATM) put with the same maturity, is a measure of tail risk in the economy (Pan (2002)). In the data, the implied volatility curve slopes upward to the left since OTM puts are relatively more expensive (Rubinstein (1994)). I take the implied volatility curve from OptionMetrics and define $SLOPE_t^{\Delta}$ as the implied volatility for an OTM put option between 20Δ to 40Δ, which I subtract from the implied volatility of an ATM put (50Δ).[16] The encompassing regression is:

$$\underbrace{SLOPE_t^{\Delta}}_{\sigma_{t,IV}^{OTM,\Delta}-\sigma_{t,IV}^{ATM}} = \beta_0 + \beta_1 \log GP_t + \beta_2 \sigma_{t,IV}^{ATM} + \epsilon_t. \tag{7}$$

To control for potential dependence of the slope on the level of implied volatility, I also control for $\sigma_{t,IV}^{ATM}$ on the right hand side of (7). Panel A of Table 10 shows the results. We see that GP is significant for all definitions of the implied volatility slope, both by itself and after controlling for the level of ATM implied volatility. The magnitude of the coefficients as well as the t-statistics and R^2 increase as the OTM put is further out-of-the-money (more tail risk). An alternative measure of tail risk is the Bakshi et al. (2003) model-free risk-neutral skewness. Bakshi and Kapadia (2003) and Jurek (2014) use this measure as a proxy for crash risk; more negative skewness is associated with more crash risk. The results in Panel B of Table 10 are similar to the results in Panel A: high GP is associated with more negative risk-neutral skewness and GP is significant even after controlling for the risk-neutral variance.

3 Gold and Platinum Markets

I examine key aspects of gold and platinum markets, including sources of demand for each metal, the leasing markets, and return dynamics. Understanding the leasing markets is important because no-arbitrage implies that investors are indifferent between buying gold (platinum) or leasing gold (platinum) in perpetuity. Understanding the variation in rental income will be important for the economic model and to compute gold and platinum returns.

3.1 Sources of Demand

Figure 6 shows the annual percentage demand for gold (top panel) and platinum (bottom panel) for each of its major uses.[17] From 1990 - 2013, approximately 70% of gold demand was for jewellery, 15% for uses in technology (semiconductors, electronics), and 15% for investments (coins, bars, ETF inventory building). Over the same period, approximately 40% of platinum demand was for jewellery, 15% for technology, while only a small fraction (less than 5%) was demanded for investment purposes. Quite conspicuously, the biggest difference between the two metals comes from the 40% of platinum demand used by the automotive industry as catalytic converters to reduce emissions in automobiles (autocat-

[16] Δ can be interpreted as the risk-neutral probability of expiring in-the-money. Lower Δ options are further out of the money.

[17] The data for gold is from Thomson-Reuters GFMS, and the data for platinum is from Johnson Matthey.

alysts).[18] Demand for platinum as autocatalysts was spurred by clean air legislation in the 1970s - securing sufficient supplies of platinum at stable prices became essential for car makers. Black (2000) (Chapter 6) describes the long-term arrangements made between platinum producers and car makers:

> *"With the introduction of autocatalysts...producers entered into long term supply contracts with the auto manufacturers. Prices were negotiated on contracts lasting up to five years".*

The private sale of platinum directly from producers to car makers means the amount of platinum used in auto production does not enter the market. Therefore, net autocatalyst demand (in excess of salvage) acts as a negative platinum supply shock since it reduces the supply of platinum available for other uses. Under this view, the major source of demand for gold and platinum comes from the jewellery industry.[19]

3.2 Lease Rates

Not surprisingly, jewellers are among the most active borrowers of gold and platinum. The LBMA and LPPM describe the leasing market:

> *"The inventory loan is the basic financial tool of the precious metals fabricating [industry]. For example, jewellery manufacturers can finance the raw material in their production process by leasing gold...The same kind of strategy would, for example, be adopted in platinum".*

Leasing is a convenient form of inventory financing widely practiced in both gold and platinum fabrication industries (LBMA and LPPM (2008)). Le and Zhu (2013) find that over the 1991 - 2007 sample, which purposely excludes the 2008 financial crisis to focus on normal time dynamics, gold lease rates are increasing in stock market returns. This is consistent with Aït-Sahalia, Parker, and Yogo (2004) who document strong positive covariation between stock returns and demand for luxury goods. In normal times, gold lease rates are procyclical: as stock returns go up, jewellers have increased need for raw materials to meet high demand for finished products and increase their gold borrowing, which drives up gold lease rates.

The picture is different in times of economic distress. Figure 7 plots annualized gold lease rates from 2007 - 2009. While gold lease rates are about 1% on average, the cost of borrowing gold during the financial crisis jumped up threefold and high gold lease rates persisted throughout the crisis. This is much greater than the observed decline in gold prices during this period, which implies the rental income (economic value of holding gold) must have been very high during the crisis. Several factors lead to countercyclical behavior of lease rates in bad times. In severe economic conditions, lenders fear default by borrowers and decrease the supply of loans, which increases the cost of leasing precious metals (LBMA (2009)). Furthermore, to the extent that there are greater benefits to service flows from holding

[18] While beyond the scope of this paper, Black (2000) (Chapter 5) provides an overview of the process by which platinum group metals catalyse the oxidation of hydrocarbons and carbon monoxide from internal combustion engines.

[19] I use the term "jeweller" to refer to gold and platinum fabricators.

gold relative to platinum in bad times (for historical and institutional reasons discussed in the introduction), gold lease rates will increase by more than platinum. In bad times, risk premia are high, which raises discount rates and lowers stock market valuations. Gold and platinum prices both fall due to strong discount rate effects, although gold prices fall by less than platinum prices since the fall in gold prices is cushioned by higher expected rental income. In fact, this reinforces why I use GP as a measure of risk in the economy, since it isolates the demand for gold as a safe asset, which is particularly important in times of high economic stress.

3.3 Gold and Platinum Returns

Previous studies of gold returns (see e.g., Erb and Harvey (2013), Barro and Misra (2013)) focus only on the price appreciation of gold and do not include the rental income over the ownership period. While Barro and Misra (2013) are correct in stating that gold dividends (rental income) are not directly observable given spot prices alone, we can compute and monetize rental income through the futures market.

Gold and platinum futures data comes from the commodities division (COMEX) of the CME (formerly NYMEX).[20] As is standard in the literature, I ignore mark-to-market of futures, and also the delivery options embedded in futures with physical settlement. I assume that futures contracts will roll in the first week of the expiration month; it is estimated that only 1% to 2% of commodities futures contracts are actually delivered, so this approach should not result in too much measurement error (Hirschey and Nofsinger (2008), Chapter 19). I examine the resulting contract maturities and verify similar to Schwartz (1997) that the maturities are relatively constant. The lease rate is given by:

$$\text{Lease rate} = \text{Libor rate - Futures premium.} \tag{8}$$

Table 11 provides a table of cash flows analysis of the above from the perspective of a jeweller, who is a typical borrower in the leasing market. For my analysis, I use futures contracts closest to 3 months to maturity, and match it with the 3 month Libor rate to calculate the lease rate, which I then annualize.[21] I choose 3 month maturities to get a contract with high liquidity, short time-to-maturity, yet not too short so that the physical delivery option does not affect prices too much. I use average daily futures prices as the monthly futures price, to be consistent with my measure of spot prices. Real gold and platinum returns (inclusive of rental income) are calculated in the standard way:

$$\begin{aligned} R_{t+1}^{real,\,gold} &= \frac{\left(\frac{P^g_{t+1}}{CPI_{t+1}} + \frac{D^g_{t+1}}{CPI_{t+1}}\right)}{\left(\frac{P^g_t}{CPI_t}\right)} \\ R_{t+1}^{real,\,platinum} &= \frac{\left(\frac{P^x_{t+1}}{CPI_{t+1}} + \frac{D^x_{t+1}}{CPI_{t+1}}\right)}{\left(\frac{P^x_t}{CPI_t}\right)} \end{aligned} \tag{9}$$

[20] The data is obtained from the Commodities Research Bureau (CRB) and are daily settlement prices direct from the exchange.

[21] Prior to 1986, I use Eurodollar deposit rates.

where P_t^g is the gold price, P_t^x is the platinum price, D_t^g is the gold rental income at time t, and D_t^x is the platinum rental income at time t, and CPI_t is the consumer price index.[22] The returns to gold and platinum can be interpreted from the perspective of an investor who owns gold or platinum, and continuously leases the metal out, earning the rental income and any price appreciation. The results are summarized in Table 12. Average gold excess returns are 2.40% per year, real gold return volatility is 16.76%, implying a Sharpe ratio of 0.14. Average platinum excess returns are 6.51% per year, real platinum return volatility is 22.18%, implying a Sharpe ratio of 0.29. For comparison, over the same period, the average excess return for U.S. equities is 7.53% per year, real equity return volatility is 15.11%. The gold risk premium is substantially lower than the equity risk premium. The risk premium for platinum is slightly lower than equities as well, although the volatility is higher. Gold lease rates are 1% per year on average. For comparison, Casassus and Collin-Dufresne (2005) estimate the gold lease rate to be 0.9% per year, while Le and Zhu (2013) find an average lease rate of about 1%. My estimate of the average platinum lease rate is 3.47% per year.

The economic model must match the low risk premium, high volatility, and low lease rate of gold. At the same time, the model must also capture the relatively high risk premium, high volatility, and high lease rate of platinum, while fitting the asset pricing dynamics of equity markets, the risk-free rate, and quantitatively accounting for the time variation and stock return predictability of GP observed in the data.

4 Economic Model

4.1 Economic Environment

I analyze whether a general equilibrium model featuring time-varying disaster risk (Wachter (2013)) and shocks to preferences for gold and platinum can jointly explain the empirical facts documented in the previous sections. I assume an endowment economy with complete markets and an infinitely-lived representative investor with Duffie and Epstein (1992) stochastic differential utility, which is the continuous-time analog of Kreps and Porteus (1978) and Epstein and Zin (1989) recursive preferences. Recursive preferences allow for a separation between risk aversion and the intertemporal elasticity of substitution (IES).[23] I focus on the case of unit IES, which is done both for tractability, and consistent with evidence in Vissing-Jørgensen (2002) and Hansen, Heaton, Lee, and Roussanov (2007).

Aggregate consumption growth is given by

$$d\log C_t = \bar{g}_c dt + \sigma_c dW_t^c + J_t^c dN_t^c \tag{10}$$

[22] I use superscript g to refer to gold, and superscript x to refer to platinum.

[23] A number of studies such as Bansal and Yaron (2004) argue that the IES should be greater than one. Others, such as Hall (1988) estimate IES to be significantly less than one, although time-varying consumption volatility can lead to large downward biases in the estimates of IES using the methodology employed in Hall (1988).

where W_t^c is a standard Brownian motion and N_t^c is a Poisson process whose intensity λ_t is given by a Cox, Ingersoll, and Ross (1985) square root process

$$d\lambda_t = \kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dW_t^\lambda + J_t^\lambda dN_t^\lambda \tag{11}$$

where W_t^λ a standard Brownian motion and N_t^λ is a Poisson process whose intensity is given by $\lambda_t^\lambda = \lambda_t$.[24] Drechsler and Yaron (2011) use a similar framework to model jumps in expected consumption growth and volatility.[25] Allowing λ_t to jump allows stock prices and volatility in the model to jump as well.[26] I solve for the stationary mean of λ_t in Appendix A3. λ_t can be approximately interpreted as the probability of a consumption disaster.[27] In the model, market volatility is endogenously determined, and evidence from the volatility estimation literature argues in favor of multiple time scales in volatility allowing for both long and short run components.[28] Also, as Seo and Wachter (2014) demonstrate, a one-factor model without time-variation in the long-run mean of λ_t generates the counterfactual prediction that the slope of the implied volatility curve decreases as the disaster intensity increases. This arises because stock return volatility is endogenously determined and is driven by λ_t itself. To relieve this tension, and consistent with the evidence from the volatility estimation literature, I follow Seo and Wachter and allow the long-run mean of λ_t to be a stochastic process ξ_t, which itself follows a square root process

$$d\xi_t = \kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi\sqrt{\xi_t}dW_t^\xi \tag{12}$$

where W_t^ξ is a standard Brownian motion. All Brownian motions and Poisson processes are assumed to be independent.

The size of the consumption jump, J_t^c is drawn from the multinomial disaster distribution of Barro and Ursua (2008), using data obtained from Robert Barro's website. The size of the jump in λ_t is given by J_t^λ, which follows an exponential distribution with mean μ_λ. Equity is modeled as a leveraged claim on aggregate consumption following Abel (1999). The aggregate dividend at time t is $D_t = C_t^\phi$, for leverage parameter ϕ, which implies that dividend growth dynamics are given by

$$d\log D_t = \phi\bar{g}_c dt + \phi\sigma_c dW_t^c + \phi J_t^c dN_t^c. \tag{13}$$

4.2 Gold and Platinum Supply

Gold and platinum do not depreciate, and consumption of the service flow from the stock of gold and platinum today does not render it less capable of providing the same service flow tomorrow. I model gold and platinum as non-depreciating durable goods. This means

[24]Nowotny (2011) considers the implications of self-exciting intensity processes to model persistent disaster states. My setup differs since realized jumps in consumption do not trigger increases in λ_t.

[25]For parsimony, I do not distinguish explicitly between X_{t-} and X_t in my notation, as it should be clear from the context.

[26]This is consistent with the evidence in Duffie et al. (2000), Broadie et al. (2007), Eraker and Shaliastovich (2008), and Tauchen and Todorov (2011).

[27]The probability of k jumps over an interval of time $\Delta t \approx e^{\lambda_t \Delta t}\frac{(\lambda_t\Delta t)^k}{k!}$.

[28]See e.g., Alizadeh, Brandt, and Diebold (2002) and Chernov, Gallant, Ghysels, and Tauchen (2003).

that the time t aggregate stock of gold and platinum increase one-to-one with the time t increment (accumulation) to the stock (Cuoco and Liu (2000)). In Appendix A4, I use data from world gold and platinum mine production to establish the properties of the gold and platinum endowment processes. The key stylized facts are: 1) the log growth rates of the aggregate per-capita gold and platinum stocks are smooth with no evidence of disasters, and 2) the aggregate per-capita gold and platinum stocks are cointegrated.[29] Given these facts, I model $\log G_t$ (the aggregate stock of gold) using a simple geometric Brownian motion which is not subject to disasters. Consistent with the empirical evidence, $\log G_t$ and $\log X_t$ (the aggregate stock of platinum) are modeled as cointegrated processes so that $\log X_t - \log G_t = \log Z_t$ is a stationary process which itself follows an Ornstein-Uhlenbeck process with long-run mean μ_z and reversion parameter θ_z:

$$\begin{aligned} d\log G_t &= \mu_g dt + \sigma_g dW_t^g \\ d\log Z_t &= \theta_z \left(\mu_z - \log Z_t\right) dt + \sigma_z dW_t^z \\ \log X_t &= \log G_t + \log Z_t. \end{aligned} \tag{14}$$

All parameters for the gold and platinum supply dynamics are directly estimated from the data.

4.3 Preferences

The representative investor's utility function is defined recursively as

$$\begin{aligned} V_t &= \mathbb{E}_t \left[\int_t^\infty f(\Omega_s, V_s) ds \right] \\ \text{for } f(\Omega, V) &= \delta(1-\gamma)V \left[\log \Omega - \frac{1}{1-\gamma} \log\left(1-\gamma\right)V \right] \\ \text{and } \Omega_t &= \left[C_t^{1-\frac{1}{\epsilon}} + \alpha_t G_t^{1-\frac{1}{\epsilon}} + \beta_t X_t^{1-\frac{1}{\epsilon}} \right]^{\frac{1}{1-\frac{1}{\epsilon}}} \end{aligned} \tag{15}$$

where $f(\Omega, V)$ describes the trade-off between current consumption Ω_t and the continuation utility V_t. The subjective time preference parameter is δ, and γ is commonly interpreted as the coefficient of relative risk aversion.

The consumption aggregator Ω_t is a constant elasticity of substitution (CES) aggregator over nondurable consumption C_t, the gold stock G_t, and the platinum stock X_t.[30] The intratemporal elasticity of substitution is ϵ.[31]

[29] Barro and Misra (2013) also find no evidence of disasters in the per-capita gold stock, using data since 1836.

[30] The agent derives utility from gold and platinum service flows in direct proportion to its stock. This is a standard way to model preference for multiple types of goods, which has been used in the durable goods literature by Ogaki and Reinhart (1998) and Yogo (2006).

[31] I use a CES aggregator with the same elasticity of substitution across all pairs of goods for parsimony and tractability relative to a specification with nested CES aggregators and separate elasticities.

The processes α_t and β_t capture in reduced-form time-varying preferences for gold and platinum:

$$\begin{aligned}\alpha_t &= \exp(a_1 + a_2\lambda_t) \\ \beta_t &= \exp(b_1 + b_2\lambda_t).\end{aligned} \tag{16}$$

Specifically, α_t and β_t represent the relative importance of gold and platinum service flows in the intratemporal consumption aggregator. Preference for precious metals responds to changes in λ_t but not directly to ξ_t, since λ_t is the probability of a consumption disaster. While the processes α_t and β_t gives me some additional flexibility, they depend completely on existing state variables and no new state variables are being added. The parameter a_2 (and b_2) cannot be arbitrarily set. We want a relatively high value of a_2 to generate enough countercyclical dynamics to match the low observed gold risk premium. However, when a_2 is too big, gold return volatility becomes too low, and gold lease rates will also be too low. Additionally, existence of solutions for gold and platinum price-dividend ratios places restrictions on the maximum a_2 and b_2 allowed, and this bound jointly depends on model parameters such as the volatility and persistence of state variables, the severity of jumps, and risk aversion. Changing these parameters to allow for high a_2 will affect equity market and risk-free rate dynamics as well.[32]

4.4 Asset Pricing

Duffie and Skiadas (1994) show that

$$\pi_t = \exp\left(\int_0^\infty f_V(\Omega_s, V_s)d\right)\ f_\Omega(\Omega_t, V_t)$$

can serve as the state-price density in this economy. Following Barro and Misra (2013), I assume that outlays on gold and platinum are negligible relative to nondurable consumption, which implies that $\Omega_t \approx C_t$. Under this assumption, the equilibrium relationship between V_t and the state variables is given by

$$V_t = \frac{C_t^{1-\gamma}}{1-\gamma} e^{a+b_\lambda\lambda_t+b_\xi\xi_t}.$$

where $\eta = -\delta(a+1)$ and $a,\ b_\lambda,\ b_\xi$ are the solutions to a system of equations given in Appendix A5. The state price density is then given by

$$\pi_t \approx \exp(\eta t - \delta b_\lambda \int_0^t \lambda_s ds - \delta b_\xi \int_0^t \xi_s ds)\ \delta\ C_t^{-\gamma} e^{a+b_\lambda\lambda_t+b_\xi\xi_t}. \tag{17}$$

I make the assumption of negligible outlays both for tractability and for economic reasons. The levels of α_t and β_t are small because per-capita expenditures on gold and platinum

[32] I have solved a version of the model based on the long-run risks (LRR) framework pioneered by Bansal and Yaron (2004) featuring jumps to uncertainty as in Eraker and Shaliastovich (2008) and Drechsler and Yaron (2011). The model is also able to deliver many similar results if the countercyclical components α_t and β_t are allowed to load on all state variables. The LRR framework delivers dividend growth predictability by GP, which is not seen in the data. I have opted for a stochastic disaster risk framework mostly for parsimony and to avoid these issues.

are small compared to expenditures on nondurable goods and services. When the CES aggregator is over multiple sources of consumption with large expenditure shares, such as nondurable and durable consumption or housing, this approximation will become wildly inaccurate; for example, Gomes, Kogan, and Yogo (2009) estimate the expenditure share of durable goods to be 50%, in which case this assumption would not be innocuous. In economic terms, the assumption implies that shocks to the supply of gold and platinum are unpriced. A mine shutdown in South Africa, for example, would affect gold and platinum prices, but would conceivably not affect aggregate stock market risk premia, which seems economically plausible. Going forward, I will assume that the approximation is accurate and describe dynamics of the stochastic discount factor in (17) with an equality sign.

The instantaneous risk-free rate is given by

$$r_t^f = \delta + (\bar{g_c} + \frac{1}{2}\sigma_c^2) - \gamma\sigma_c^2 + \lambda_t \mathbb{E}_v \left[e^{(1-\gamma)J_t^c} - e^{-\gamma J_t^c} \right].$$

I follow Barro (2006) and Wachter (2013) and suppose that if a disaster occurs, the government will default on debt obligations with probability q, leading to a loss in the same proportion as the consumption loss in the disaster.

The user costs (rental income) of gold and platinum are determined in equilibrium by the intratemporal optimality conditions:

$$\begin{aligned} Q_{g,t} = \frac{\Omega_G}{\Omega_C} &= \underbrace{\alpha_t \times \left(\frac{C_t}{G_t}\right)^{\frac{1}{\epsilon}}}_{\text{countercyclical} \times \text{procyclical}} \\ Q_{x,t} = \frac{\Omega_X}{\Omega_C} &= \underbrace{\beta_t \times \left(\frac{C_t}{X_t}\right)^{\frac{1}{\epsilon}}}_{\text{countercyclical} \times \text{procyclical}} \end{aligned} \tag{18}$$

where $Q_{g,t}$ is the user cost of gold and $Q_{x,t}$ is the user cost of platinum. Notice from equation (18) that the intratemporal elasticity of substitution ϵ behaves like the inverse of the leverage parameter ϕ, since shocks to gold and platinum supply are small and unpriced. When $\frac{1}{\epsilon} < \phi$, gold and platinum will be safer than levered equity and command a lower risk premium, while the opposite will be true if $\frac{1}{\epsilon} > \phi$. Lower values of ϵ lead to higher risk premia and volatility for gold and platinum returns, and also imply greater complementarity between nondurable consumption, gold, and platinum. Barro and Misra (2013) set $\epsilon > 1$, which makes gold less risky than unlevered equity. The authors use this mechanism to generate a low gold risk premium. However, as Wachter (2013) points out, under recursive preferences, when $\phi < 1$ (in this case, $\epsilon > 1$) the price-dividend ratio is increasing λ_t. The same result holds in my model since gold and platinum supply shocks are unpriced. This means that under the Barro and Misra (2013) assumption that $\epsilon > 1$, the model would predict that gold lease rates fall (gold prices rise) when the probability of a disaster increases, which is counterfactual in light of Figures 1 and 7. Intuition suggests $\epsilon < 1$ is more reasonable if we view gold and platinum as jewellery, since jewellery complements nondurable consumption but does not substitute for it. Furthermore, $\epsilon > 1$ results in gold return volatility being too low because in this case gold becomes a deleveraged consumption claim. In my calibration, I fix $\epsilon = \frac{1}{\phi}$ so that all the countercyclical properties of gold and platinum arise through α_t and β_t.

Let P_t be the price of a claim to the stream of dividends D_t, and $P_t^{t+\tau}$ be the price of the asset which pays the single risky dividend $D_{t+\tau}$ and nothing else. No arbitrage implies that $\pi_t P_t^{t+\tau}$ is a martingale, which implies that the equity price-dividend ratio is given by

$$\frac{P_t}{D_t} = \int_0^\infty e^{a_\phi(\tau)+b_\phi(\tau)\lambda_t+c_\phi(\tau)\xi_t} d\tau = G(\lambda_t, \xi_t). \tag{19}$$

Similar arguments hold for $P_{g,t}$ and $P_{x,t}$, which are the claims to gold and platinum, respectively:

$$\begin{aligned} \frac{P_{g,t}}{Q_{g,t}} &= \int_0^\infty e^{a_g(\tau)+b_g(\tau)\lambda_t+c_g(\tau)\xi_t} d\tau = G^g(\lambda_t, \xi_t) \\ \frac{P_{x,t}}{Q_{x,t}} &= \int_0^\infty e^{a_x(\tau)+b_x(\tau)\lambda_t+c_x(\tau)\xi_t+d_x(\tau)\log Z_t} d\tau = G^x(\lambda_t, \xi_t, \log Z_t). \end{aligned} \tag{20}$$

The equity functions $a_\phi(\tau)$, $b_\phi(\tau)$, $c_\phi(\tau)$, gold functions $a_g(\tau)$, $b_g(\tau)$, $c_g(\tau)$, and platinum functions $a_x(\tau)$, $b_x(\tau)$, $c_x(\tau)$, $d_x(\tau)$ are given by the solution to systems of ordinary differential equations described in Appendix A5.

4.5 GP in the Model

While I use the exact log GP ratio in my model simulations, a log-linearization conveys the economic intuition more clearly.[33]

In Appendix A6, I show that we can write log-linearized gold ($P_{g,t}$) and platinum ($P_{x,t}$) prices as

$$\begin{aligned} \log P_{g,t} &= A_g + \frac{1}{\epsilon}\log C_t - \frac{1}{\epsilon}\log G_t + \underbrace{(a_2 + b^*_{g,\lambda})}_{<0}\lambda_t + \underbrace{b^*_{g,\xi}}_{<0}\xi_t \\ \log P_{x,t} &= A_x + \frac{1}{\epsilon}\log C_t - \frac{1}{\epsilon}\log G_t + \underbrace{(b_2 + b^*_{x,\lambda})}_{<0}\lambda_t + \underbrace{b^*_{x,\xi}}_{<0}\xi_t + \underbrace{(b^*_{x,Z} - \frac{1}{\epsilon})}_{<0}\log Z_t \end{aligned} \tag{21}$$

where $A_g, A_x, b^*_{g,\lambda}, b^*_{g,\xi}, b^*_{x,\lambda}, b^*_{x,\xi}, b^*_{x,Z}$ are constants described in Appendix A6. Positive shocks to $\log C_t$ imply higher service flows and raise gold and platinum prices. The increase is greater than the increase in consumption itself because of complementarity between non-durable consumption and gold and platinum service flows ($\frac{1}{\epsilon} > 1$). High $\log G_t$ lowers gold prices since the quantity of gold becomes less scarce, and also lowers platinum prices due

[33] The exact log GP ratio is given by

$$\begin{aligned} \log GP_t &= \log \frac{P_{g,t}}{P_{x,t}} = \log \frac{G^g(\lambda_t, \xi_t)}{G^x(\lambda_t, \xi_t, \log Z_t)} \frac{Q_{g,t}}{Q_{x,t}} \\ &= (a_1 - b_1) + \log \frac{G^g(\lambda_t, \xi_t)}{G^x(\lambda_t, \xi_t, \log Z_t)} + (a_2 - b_2)\lambda_t + \frac{1}{\epsilon}\log Z_t. \end{aligned}$$

to cointegration. Higher $\log Z_t$ means that (all else equal) the quantity of platinum is less scarce, which also lowers platinum prices. Under my model calibration, strong discount rate effects imply that, despite $a_2, b_2 > 0$, the overall response of gold and platinum to increases in λ_t and ξ_t are negative, so that gold and platinum prices fall as disaster risks increase.

The log GP ratio is the difference between the log gold and platinum prices and is given by

$$\begin{aligned}\log GP_t &= \log \frac{P_{g,t}}{P_{x,t}} \\ &= \text{cons} + \underbrace{(\frac{1}{\epsilon} - b^*_{x,Z})}_{>0} \log Z_t + \underbrace{(a_2 - b_2 + b^*_{g,\lambda} - b^*_{x,\lambda})}_{>0} \lambda_t + \underbrace{(b^*_{g,\xi} - b^*_{x,\xi})}_{>0} \xi_t. \end{aligned} \tag{22}$$

Shocks to $\log C_t$ (which can be thought of as shocks to jewellery demand) affect gold and platinum prices equally, leaving GP insulated from consumption shocks. Likewise, shocks to $\log G_t$ alone also cancel out and only the relative difference in supply $\log Z_t$ matters for GP. Platinum is more expensive than gold on average because $\log X_t < \log G_t$ on average (platinum is more scarce). When $\log Z_t$ goes up, gold becomes scarce relative to platinum, which increases GP. In the model, GP is increasing in both λ_t and ξ_t.[34] High disaster probabilities imply high risk premia, which leads to high discount rates and low equity prices. Since a_2 and b_2 are positive, the service flows from gold and platinum increase when disaster probabilities increase, which partially offsets the higher discount rates and cushions the fall in prices. This works similar to a cash flow effect, where the cash flow represents gold and platinum rental income. Furthermore, $a_2 > b_2$ implies that the higher service flow is greater for gold relative to platinum, which not only affects the immediate service flow but also expected future service flow (rental income) through persistence in disaster probabilities. This means that gold and platinum prices both fall as disaster probabilities increase, but gold prices fall by less relative to platinum and GP is increasing in the disaster probabilities. The fact that GP increases in λ_t and ξ_t allows the model to generate the observed return predictability at both long and short horizons.

The $\log Z_t$ term is not priced by the stochastic discount factor but does affect the volatility and persistence of GP. Stationarity of GP in the model is assured because $\log Z_t$ is stationary, or in other words, because $\log G_t$ and $\log X_t$ are cointegrated. Interestingly, while shocks to $\log Z_t$ affect GP, they do not affect return predictability, which suggests that controlling for $\log Z_t$ in the data can potentially lead to even stronger return predictability by GP. I verify that this indeed holds in the data and discuss the results in the following section.

5 Calibration and Model Simulation Results

My parameter choices are given in Table 13. I have opted for smaller average jump sizes with an average disaster size of 15%. Barro (2006) uses the dataset of Madison (2003) and found the average disaster size to be 29%. Barro and Ursua (2008) update Madison (2003) and find that the average disaster size is about 22%; this disaster distribution is also used in Wachter

[34]That the GP ratio increases in both λ_t and ξ_t is dependent on the calibration. This holds under the model parameters I use for this model.

(2013).[35] I opt for smaller average disaster sizes in line with evidence from Nakamura, Steinsson, Barro, and Ursua (2013), who estimate the average permanent impact of disasters to be about 15%. While the actual probability of these smaller disasters is 5.85%, I opt for a more conservative calibration of 4%, which is achieved using a $\bar{\xi} = 0.0355$ as in Wachter (2013) along with an average jump size of $\mu_\lambda = 0.03$ in the event of a jump in λ_t. Figure 8 compares my multinomial jump size distribution with smaller average jump sizes to the distribution used in Barro and Ursua (2008) and Wachter (2013). An important challenge in calibrating representative investor models is to match the high observed volatility of the price-dividend ratio. The model places an upper bound on the amount of volatility in the state variables that can be allowed for solutions to exist (this is clearly seen in the equations for the Epstein-Zin discount factor in Appendix A5). I fix σ_ξ such that the discriminant in the solution to b_ξ is zero, which helps match the high volatility of the price-dividend ratio and also reduces the number of free parameters. The λ_t process is calibrated to be less persistent than ξ_t.

Table 14 describes the fit of the model to the data. State variables are simulated at a monthly frequency and aggregated to an annual frequency. The data moments are from 1975-2013. The model matches the low gold risk premium, relatively high gold return volatility, low Sharpe ratio, and low lease rate. The model-implied gold lease rate is 0.93%, which compares well to the 1% lease rate in the data. Lease rates in the model are the convenience yield, which corresponds to the dividend yield. For comparison, I also present the model 90% confidence intervals for simulation paths in which no disaster occurred. While these no-disaster intervals are more appropriate to compare against stock and bond moments (since no disasters have occurred in the recent U.S. data, on which the stock and bond returns are based), for gold and platinum returns it is more natural to compare against population moments, since there have been numerous economic disasters from 1975 - 2006 in international markets (using my disaster cutoff) based on the Barro and Ursua (2008) dataset (including several OECD countries), which can conceivably affect gold and platinum returns and volatilities. The model explains the expected returns, volatilities and lease rates for platinum as well, including the high lease rate and high volatility. The model also accounts for time variation in GP, with the volatility and persistence of GP falling inside the 90% confidence intervals. The median persistence for all simulations matches the data estimate nearly perfectly. Following this, I run the below return predictability regressions using model excess stock returns and GP:

$$\frac{1}{h}\sum_{i=1}^{h} \log(R^e_{t+i}) - \log(R^b_{t+i}) = \beta_0 + \beta_1 \log(GP_t) + \epsilon_{t+h}.$$

The left hand side is the annualized excess return for one year up through five years ahead, while the right hand side is the model GP. The results are shown in the top panel of Table 15. The data estimates fall right in the model confidence intervals, with the data R^2 estimates very close to the median values.[36] Thus, the model can explain the observed predictability

[35] The cutoff in Barro (2006) for a disaster was a 15% peak-to-trough decline in GDP per capita, while Barro and Ursua (2008) used a cut-off of 10%. To achieve an average disaster size of 15%, my cutoff is 6%.

[36] It is difficult to decide which, all simulations or no disasters, is most appropriate for the predictability exercise, since U.S. stock returns were not affected by domestic disasters, while the GP ratio is potentially affected by international disasters. For completeness, I include both sets of results.

of returns by GP. Similar to the data, the model delivers very low to negligible dividend growth predictability, similar to (Wachter (2013)). The model can also account for the observed relationship between GP and the slope of the implied volatility curve for index options, as detailed in Appendix A7.

How well have I captured the effect of gold and platinum supply dynamics on GP? Is there predictability coming from the supply effects (including autocatalyst demand)? The second panel of Table 15 investigates this issue. I regress GP on $\log Z_t$ inside the model, and we see that the data estimate falls right inside the 90% interval.[37] Since the leading coefficient on $\log Z_t$ in the model depends on $\frac{1}{\epsilon}$, this serves as a further check on the assumed complementarity ($\epsilon < 1$) between jewellery (gold and platinum) and nondurable consumption. Under a calibration where $\epsilon > 1$ as in Barro and Misra (2013), this regression in the data results in a coefficient smaller than 1. The second regression in this panel investigates return predictability by $\log Z_t$ in both the model and the data. In the model, $\log Z_t$ does not predict returns by construction, although in small samples it is occasionally possible to spuriously find weak evidence of predictability. Both the population and median values, however, show that there is no predictability coming from the supply channels. I run the same regression in the data and find no evidence of predictability through $\log Z_t$, which is evidence that the predictability does not come from gold and platinum supply dynamics.

These results for repeated samples of 39 years lead to an interesting finding. Time-variation in GP over finite samples is affected by $\log Z_t$, which is not a priced variable in this economy. The third panel shows return predictability regressions where I control for the effect of $\log Z_t$, which adds volatility and persistence to GP without adding predictive power. We see in this case that the point estimates increase at all horizons, and now the 90% interval for return predictability by GP does not contain 0. The R^2 increase over all horizons quite dramatically. In the data, we can separately identify $\log Z_t$ and $\log A_t$, the aggregate per-capita stock of platinum used as autocatalysts. Empirically, a regression of GP on $\log Z_t$ gives a significant, positive coefficient while a regression of GP on $\log A_t$ gives a significant, negative coefficient. When $\log Z_t$ is high, platinum is relatively more plentiful (gold more scarce) so gold is relatively more expensive than platinum. High values of $\log A_t$ correspond to high demand for platinum as autocatalysts, which is associated with higher platinum prices (lower GP). Neither $\log Z_t$ nor $\log A_t$ predict returns in the data. Controlling for persistent supply effects, the persistence of GP is lower; the monthy AR(1) coefficient is 0.96, which implies a half-life of about 1.5 years.

6 Conclusion

The risk and return tradeoff is one of the central tenets of asset pricing theory. However, empirically identifying a viable proxy for risk, as manifest through robust return predictability, cross-sectional pricing, and basic economic intuition, has been largely elusive in the literature. In this paper, I show that the ratio of gold to platinum prices (GP) proxies for an important aggregate source of risk in the economy. GP predicts future stock returns in the time-series, outperforms other predictors proposed in the literature, and GP risk is priced in the cross-section of stock returns. GP is persistent and significantly correlated with tail

[37] For the data $\log Z_t$, I interpolate annual values to monthly values.

risk measures implied by options markets. An equilibrium model with time-varying tail risk and shocks to preferences for gold and platinum can quantitatively account for the asset pricing dynamics of equity, gold, and platinum markets, as well as the time variation and return predictability of GP. In the model, higher aggregate risk lowers gold and platinum prices through strong discount rate effects, although gold prices fall by less due to higher expected rental income, which is consistent with the empirical evidence. I achieve these results by modeling the countercyclical component of gold and platinum service flows in reduced-form. The micro-foundations of this mechanism are an important open question, which I leave for fruitful future research.

Appendix

A1. Econometric Inference for Predictive Regressions

Stambaugh (1999) shows that predictive regressions using persistent predictors are biased in finite samples. The standard return predictability regression is:

$$r^e_{t+1} = \alpha + \beta x_t + \epsilon_{t+1}$$

where $r^e_{t+1} = \log\left(\frac{P_{t+1}+D_{t+1}}{P_t}\right) - r^f_t$ is the log excess return from time t to time $t+1$, and x_t is some predictor known at time t such as the log price-dividend ratio or the log GP ratio. If x_t is a persistent predictor, we can model it as an AR(1) process:

$$x_{t+1} = \mu + \rho x_t + u_{t+1}$$

For predictors such as the price-dividend ratio, $cov(\epsilon, u) \neq 0$, since a positive return shock typically means prices increased, which also increases the price-dividend ratio. Letting $\gamma = \frac{cov(\epsilon_{t+1}, u_{t+1})}{var(u_{t+1})}$, the bias in the estimate of the predictive beta can be written as:

$$\mathbb{E}[\hat{\beta} - \beta] = \gamma \underbrace{\mathbb{E}[\hat{\rho} - \rho]}_{\approx -\frac{(1+3\rho)}{T} < 0} \tag{23}$$

The degree of bias is proportional to γ, which can be estimated as the slope of the regression of residuals from the predictive regression on the residuals from the AR(1) regression of the predictor variable. For the price-dividend ratio, the correlation between ϵ and u in the data is 0.94, while it is only -0.17 for the GP ratio. Also note that there is no mechanical correlation between the residuals as is the case for the PD ratio. More formally, I project $\hat{\epsilon}_t$ on $\hat{u}_t$ and estimate $\hat{\gamma}$ to be 10.55 for the PD ratio, whereas for the GP ratio $\hat{\gamma}$ is only -1.78. Evaluating at the maximum bias ($\rho = 1$) estimates an upper bound of -0.090 for PD ratio bias, which is enough to change the sign, whereas it is only 0.015 for the GP ratio, which is small compared to the predictive beta of 0.237. The evidence suggests that the GP ratio predictability is not driven by finite sample bias.

Predictor persistence also potentially affects the size of tests (see e.g., Torous et al., 2004). For $\delta = corr(\epsilon, u)$, the test statistic for β has a non-standard limiting distribution:

$$t_\beta \Longrightarrow \delta \tau_\rho + \sqrt{(1-\delta^2)}\, z$$

where τ_ρ is non-normal and z is normal. I follow Elliot and Stock (1994) and use Monte Carlo simulations to asses the magnitude of these size distortions. I run 100,000 simulations of length equal to my sample size at a monthly frequency by simulating the above dynamics, evaluating all parameters using their sample values. When $\delta = 0.94$ (which is the case for the PD ratio), a 5% test has a true rejection rate of 17%. For the GP ratio, where $\delta = -0.167$, a 5% test has a true rejection rate of 6%, which is very close to the true size. Since the absolute value of δ is small in the case of the GP ratio, the significance of the predictability tests is not affected by potential size distortions due to predictor persistence.

A2. Realized Utility Gains

I calculate realized utility gains for an investor who maximizes mean-variance preferences given a risk aversion of γ.[38] Given forecasts of expected returns and stock market volatility, the investor optimally allocates between stocks which earn the market return and bonds which earn the risk-free rate. The allocation to stocks for period $t+1$ is formed in period t. Using the historical average as the estimate of expected return, the allocation is:

$$w_{1,t} = \left(\frac{1}{\gamma}\right)\left(\frac{\overline{r}_{t+1}}{\widehat{\sigma}^2_{t+1}}\right)$$

Using the GP ratio, the allocation to stocks is given by:

$$w_{2,t} = \left(\frac{1}{\gamma}\right)\left(\frac{\widehat{r}_{t+1}}{\widehat{\sigma}^2_{t+1}}\right)$$

For both portfolio choice problems, $\widehat{\sigma}^2_{t+1}$ is the forecasted variance of stock returns over the next month, which I estimate following Li et al. (2013) by using a ten-year trailing window of monthly stock returns. The average utility levels for the investor over the out-of-sample period are given by:

$$\begin{aligned} U_1 &= \mu_1 - \frac{1}{2}\gamma\widehat{\sigma}^2_1 \\ U_2 &= \mu_2 - \frac{1}{2}\gamma\widehat{\sigma}^2_2 \end{aligned} \tag{24}$$

where μ_i is the sample mean of the return for the portfolio formed based on strategy i, where strategy 1 uses the historical average and strategy 2 uses the GP ratio. We can view the utility level as a certainty equivalent return for an investor with these preferences (Kandel and Stambaugh (1996)).[39] The utility gain of using the GP ratio over the historical average in percentage terms is given by $1200 \times (U_2 - U_1)$, which can be thought of as the management fee that an investor with mean-variance preferences would be willing to pay to access the GP ratio to generate return forecasts. In the data, the GP ratio produces a large, positive utility gain of 4.53% while other popular forecasting variables offer much lower or weakly negative utility gains.[40]

A3. Stationary Mean of λ_t

I compute the stationary mean of the λ_t process. The process is given by:

$$d\lambda_t = \kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dW^\lambda_t + J^\lambda_t dN^\lambda_t$$

which implies that

$$\begin{aligned} \frac{\mathbb{E}\left[\lambda_t\right]}{dt} &= \kappa_\lambda(n(t) - m(t)) + \mu_\lambda\rho_0 + \mu_\lambda\rho_1 m(t) \\ \frac{\mathbb{E}\left[\xi_t\right]}{dt} &= \kappa_\xi\bar{\xi} - \kappa_\xi n(t) \end{aligned} \tag{25}$$

[38] Following the literature, I set $\gamma = 3$.

[39] Following Campbell and Thompson (2008), I constrain the allocation to stocks to be between 0% and 150%.

[40] The results for other forecasting variables are similar to the results in Li et al. (2013).

where $m(t) = \mathbb{E}[\lambda_t]$ and $n(t) = \mathbb{E}[\xi_t]$, with $n(t) \to \bar{\xi}$. Solving the ordinary differential equation for $m(t)$ implied by (25) results in the stationary mean of λ_t:

$$\mathbb{E}[\lambda_\infty] = \lim_{t\to\infty} m(t) = \frac{\kappa_\lambda \bar{\xi} + \mu_\lambda \rho_0}{\kappa_\lambda - \mu_\lambda \rho_1} \tag{26}$$

with necessary conditions $\kappa_\lambda > \mu_\lambda \rho_1$ and $\kappa_\lambda \bar{\xi} + \mu_\lambda \rho_0 > 0$, which are satisfied under the model calibration.

A4. Gold and Platinum Mine Production

The data for world gold mine production is from the U.S. Geological Survey (USGS) Annual Mineral Yearbook reports. The data for world platinum mine production is from Johnson Matthey. The data is annual from 1975 to 2013. The Johnson Matthey data details both annual platinum mine production as well as autocatalyst demand and salvage. I use as my measure of the increment to the platinum stock the total quantity mined in a given year minus the autocatalyst demand net of salvage. Thomas and Boyle (1986) estimate the initial world above-ground stock of gold at the end of 1974 to be 84,000 tonnes (2,700 million troy oz). There is not a consensus estimate of world above-ground platinum stock (net of autocatalysts) that I am aware of, although during the 1975 to 2013 period in the data, annual platinum production (net of autocatalyst demand) is consistently approximately 4.5% of gold production with very little variation each year. Annual gold production is approximately 2,000 tonnes and platinum production is approximately 90 tonnes. Therefore, I estimate the initial stock of platinum to be 3,780 tonnes.[41] Using market prices at the end of 2013, this puts the total dollar value of all gold in the world at $6.4 trillion, and the value of all platinum in the world (not found in autocatalysts) at just over $300 billion. I proxy for population growth using U.S. annual population growth data provided by the U.S. Census Bureau.[42]

Panel A of Table 16 describes the log growth rate of the aggregate per-capita stock of gold and platinum. G_t is the per-capita stock of gold, and X_t is the per-capita stock of platinum. The mean per-capita log growth rate of the aggregate gold stock is 0.72% per year, and the growth rate is very smooth: the standard deviation is only 0.21%. The platinum stock displays similar dynamics, with an average growth rate of 0.71% and a standard deviation of 0.29%. Furthermore, the means are close to the medians.

Given the stable relationship between gold and platinum production each year, I look for evidence of cointegration between the log per-capita stock of gold and platinum. Two processes $\log X_t$ and $\log G_t$ are cointegrated if there exists a vector β such that $\beta' \begin{bmatrix} \log X_t \\ \log G_t \end{bmatrix}$ is a stationary process. Panel B shows that $\log G_t$ and $\log X_t$ are unit root processes: an augmented Dickey and Fuller (1979) test fails to reject the null of a unit root at all

[41] Is this a reasonable estimate? While the discovery of platinum is often credited to Antonio de Ulloa in 1735, it was not until Hans Merensky identified large economic deposits of platinum in the Bushveld Igneous Complex of South Africa in 1924 that large scale platinum mining took place (Cawthorn (1999)). My results are robust to reasonable perturbations of the initial estimate.

[42] I use U.S. population growth as opposed to world population growth to be consistent with the consumption data in the calibration, which uses U.S. per-capita consumption data.

lags 1 through 5. However, the process $\log Z_t = \log X_t - \log G_t$ appears to be stationary. I estimate the cointegration vector using Dynamic Least Squares (DLS) as suggested by Stock and Watson (1993)in Panel C of Table 16:

$$\log X_t = \beta_0 + \beta_G \log G_t + \sum_{i=-k}^{k} \gamma_i \Delta \log G_{t-i} + \epsilon_t \tag{27}$$

for $k = 1, 2, 3$. The estimates of β_G are significant, ranging from 0.99 to 1.04, and in all cases a 95% confidence interval includes 1, which suggests that the cointegration vector is not statistically different from $[1, -1]$.

As further evidence of cointegration, I estimate the joint system $Y_t = \begin{bmatrix} \log X_t & \log G_t \end{bmatrix}'$ in a Engle and Granger (1987) Vector Error-Correction Model (VECM):

$$\Delta Y_t = \mu + \Pi Y_{t-1} + \sum_{j=1}^{p-1} \Gamma_j \Delta Y_{t-j} + \epsilon_t \tag{28}$$

and conduct Johansen (1988) rank tests for cointegration based on the rank of the matrix Π. The null hypothesis for the rank test is that there are no more than r cointegrating relationships, which implies that the remaining $K - r$ eigenvalues of Π must be zero where K is the dimension of Y_t. I follow Johansen (1995) and apply an iterative procedure which starts testing at $r = 0$ and accepts as $\hat{r}$ (number of cointegrating relationships) the first value of r for which the test fails to reject the null. Table 16 Panel D shows the results for the VECM with 1 through 4 lags. We see that for the estimation with 3 lags (which is the optimal lag length as chosen by the Akaike Information Criterion), we reject the null of zero cointegrating relationships, but fail to reject the null of 1 cointegrating relationship.[43] Figure 9 plots the demeaned $\log X_t$ and $\log G_t$ processes, where we can clearly see that gold and platinum supply seem to track each other over time.

A5. Model Solution

The equations for the Epstein-Zin discount factor coefficients are given by:

$$\begin{aligned} a &= \frac{(1-\gamma)(\frac{1}{2}(1-\gamma)\sigma_c^2)}{\delta} + \frac{b_\xi \kappa_\xi \bar{\xi} + \rho_0 \mathbb{E}_\eta \left[e^{b_\lambda J_t^\lambda} - 1 \right]}{\delta} \\ 0 &= \frac{1}{2}\sigma_\lambda^2 {b_\lambda}^2 - (\kappa_\lambda + \delta) b_\lambda + \rho_1 \mathbb{E}_\eta \left[e^{b_\lambda J_t^\lambda} - 1 \right] + \mathbb{E}_v \left[e^{(1-\gamma) J_t^c} - 1 \right] \\ b_\xi &= \frac{k_\xi + \delta}{{\sigma_\xi}^2} - \sqrt{\left(\frac{k_\xi + \delta}{{\sigma_\xi}^2} \right)^2 - \frac{2 b_\lambda \kappa_\lambda}{{\sigma_\xi}^2}}. \end{aligned}$$

In general, I can allow $\lambda_t^\lambda = \rho_0 + \rho_1 \lambda_t$. Equity price-dividend ratio is given by:

$$\frac{P_t}{D_t} = \int_0^\infty \exp(a_\phi(\tau) + b_\phi(\tau)\lambda_t + c_\phi(\tau)\xi_t) d\tau$$

[43] The estimated β_G in the VECM from $\Pi = \alpha\beta'$ has a 95% confidence interval of [-1.45,-1.04].

with $a_\phi(\tau), b_\phi(\tau), c_\phi(\tau)$ given by the ODEs:

$$\begin{aligned}
a'_\phi(\tau) &= c_\phi(\tau)\kappa_\xi\bar{\xi} + \phi\bar{g_c} + \frac{1}{2}\phi^2\sigma_c^2 - \delta - (\bar{g_c} + \frac{1}{2}\sigma_c^2) \\
&\quad + \gamma(1-\phi)\sigma_c^2 + \rho_0\mathbb{E}_\eta\left[e^{(b_\lambda + b_\phi(\tau))J_t^\lambda} - e^{b_\lambda J_t^\lambda}\right] \\
b'_\phi(\tau) &= \frac{1}{2}\sigma_\lambda^2 b_\phi(\tau)^2 + (b_\lambda\sigma_\lambda^2 - \kappa_\lambda)b_\phi(\tau) + \mathbb{E}_v\left[e^{(\phi-\gamma)J_t^c} - e^{(1-\gamma)J_t^c}\right] \\
&\quad + \rho_1\mathbb{E}_\eta\left[e^{(b_\lambda + b_\phi(\tau))J_t^\lambda} - e^{b_\lambda J_t^\lambda}\right] \\
c'_\phi(\tau) &= \frac{1}{2}\sigma_\xi^2 c_\phi(\tau)^2 + (b_\xi\sigma_\xi^2 - \kappa_\xi)c_\phi(\tau) + b_\phi(\tau)\kappa_\lambda
\end{aligned}$$

with initial conditions $a_\phi(0) = b_\phi(0) = c_\phi(0) = 0$.

Let $P_{g,t}^{t+\tau}$ be the price of zero-coupon gold which pays $Q_{g,t+\tau}$ and nothing else, and let $P_{x,t}^{t+\tau}$ be the analogous claim for platinum. Gold and platinum price-dividend ratios are solved by noting that $\pi_t P_{g,t}^{t+\tau}$ and $\pi_t P_{x,t}^{t+\tau}$ are martingales, so the sum of the drift and jump compensator must equal zero.

The gold price-dividend ratio is given by:

$$\frac{P_{g,t}}{Q_{g,t}} = \int_0^\infty \exp(a_g(\tau) + b_g(\tau)\lambda_t + c_g(\tau)\xi_t)d\tau$$

with $a_g(\tau), b_g(\tau), c_g(\tau)$ given by the ODEs:

$$\begin{aligned}
a'_g(\tau) &= c_g(\tau)\kappa_\xi\bar{\xi} + \frac{1}{\epsilon}\left[\bar{g_c} - \mu_g + \frac{1}{2\epsilon}(\sigma_c^2 + \sigma_g^2)\right] - \delta - (\bar{g_c} + \frac{1}{2}\sigma_c^2) \\
&\quad + \gamma(1 - \frac{1}{\epsilon})\sigma_c^2 + \rho_0\mathbb{E}_\eta\left[e^{(a_2 + b_\lambda + b_g(\tau))J_t^\lambda} - e^{b_\lambda J_t^\lambda}\right] \\
b'_g(\tau) &= \frac{1}{2}\sigma_\lambda^2 b_g(\tau)^2 + \left[(a_2 + b_\lambda)\sigma_\lambda^2 - \kappa_\lambda\right] b_g(\tau) \\
&\quad + \frac{1}{2}\sigma_\lambda^2 a_2^2 + a_2(b_\lambda\sigma_\lambda^2 - \kappa_\lambda) + \mathbb{E}_v\left[e^{(\frac{1}{\epsilon}-\gamma)J_t^c} - e^{(1-\gamma)J_t^c}\right] \\
&\quad + \rho_1\mathbb{E}_\eta\left[e^{(a_2 + b_\lambda + b_g(\tau))J_t^\lambda} - e^{b_\lambda J_t^\lambda}\right] \\
c'_g(\tau) &= \frac{1}{2}\sigma_\xi^2 c_g(\tau)^2 + (b_\xi\sigma_\xi^2 - \kappa_\xi)c_g(\tau) + (a_2 + b_g(\tau))\kappa_\lambda
\end{aligned}$$

with initial conditions $a_g(0) = b_g(0) = c_g(0) = 0$.

The platinum price-dividend ratio is given by:

$$\frac{P_{x,t}}{Q_{x,t}} = \int_0^\infty \exp(a_x(\tau) + b_x(\tau)\lambda_t + c_x(\tau)\xi_t + d_x(\tau)\log Z_t)d\tau$$

with $a_x(\tau), b_x(\tau), c_x(\tau), d_x(\tau)$ given by the ODEs:

$$
\begin{aligned}
a_x'(\tau) &= c_x(\tau)\kappa_\xi\bar{\xi} + \frac{1}{\epsilon}\left[\bar{g}_c - \mu_g + \frac{1}{2\epsilon}(\sigma_c^2+\sigma_g^2+\sigma_z^2)\right] - \delta - (\bar{g}_c + \frac{1}{2}\sigma_c^2) \\
&+ \gamma(1-\frac{1}{\epsilon})\sigma_c^2 + \frac{1}{2}\sigma_z^2 d_x(\tau)^2 + d_x(\tau)(\theta_z\mu_z - \frac{1}{\epsilon}\sigma_z^2) \\
&- \frac{1}{\epsilon}\theta_z\mu_z + \rho_0\mathbb{E}_\eta\left[e^{(b_2+b_\lambda+b_x(\tau))J_t^\lambda} - e^{b_\lambda J_t^\lambda}\right] \\
b_x'(\tau) &= \frac{1}{2}\sigma_\lambda^2 b_x(\tau)^2 + \left[(b_2+b_\lambda)\sigma_\lambda^2 - \kappa_\lambda\right] b_x(\tau) \\
&+ \frac{1}{2}\sigma_\lambda^2 b_2^2 + b_2(b_\lambda\sigma_\lambda^2 - \kappa_\lambda) + \mathbb{E}_v\left[e^{(\frac{1}{\epsilon}-\gamma)J_t^c} - e^{(1-\gamma)J_t^c}\right] \\
&+ \rho_1\mathbb{E}_\eta\left[e^{(b_2+b_\lambda+b_x(\tau))J_t^\lambda} - e^{b_\lambda J_t^\lambda}\right] \\
c_x'(\tau) &= \frac{1}{2}\sigma_\xi^2 c_x(\tau)^2 + (b_\xi\sigma_\xi^2 - \kappa_\xi)c_x(\tau) + (b_2 + b_x(\tau))\kappa_\lambda \\
d_x'(\tau) &= -\theta_z d_x(\tau) + \theta_z\frac{1}{\epsilon}
\end{aligned}
$$

with initial conditions $a_x(0) = b_x(0) = c_x(0) = d_x(0) = 0$.

A6. Log-Linearized Gold and Platinum Prices

Gold price-dividend ratios are given by

$$\frac{P_{g,t}}{Q_{g,t}} = G^g(\lambda_t, \xi_t) = \int_0^\infty e^{a_g(\tau)+b_g(\tau)\lambda_t+c_g(\tau)\xi_t}d\tau$$

Let $g(\lambda_t, \xi_t) = \log G^g(\lambda_t, \xi_t)$. Given fixed λ^*, ξ^*, Taylor expansion implies that

$$g(\lambda, \xi) \approx g(\lambda^*, \xi^*) + \left.\frac{\partial g}{\partial \lambda}\right|_{(\lambda^*,\xi^*)}(\lambda_t - \lambda^*) + \left.\frac{\partial g}{\partial \xi}\right|_{(\lambda^*,\xi^*)}(\xi_t - \xi^*)$$

where we have that

$$
\begin{aligned}
\left.\frac{\partial g}{\partial \lambda}\right|_{(\lambda^*,\xi^*)} &= \frac{1}{G(\lambda^*,\xi^*)}\int_0^\infty b_g(\tau)e^{a_g(\tau)+b_g(\tau)\lambda_t+c_g(\tau)\xi_t}d\tau = b_{g,\lambda}^* \\
\left.\frac{\partial g}{\partial \xi}\right|_{(\lambda^*,\xi^*)} &= \frac{1}{G(\lambda^*,\xi^*)}\int_0^\infty c_g(\tau)e^{a_g(\tau)+b_g(\tau)\lambda_t+c_g(\tau)\xi_t}d\tau = b_{g,\xi}^*.
\end{aligned}
$$

This implies that $G^g(\lambda_t,\xi_t) \approx G^g(\lambda^*,\xi^*)e^{b_{g,\lambda}^*(\lambda_t-\lambda^*)+b_{g,\xi}^*(\xi_t-\xi^*)}$, and I set λ^* and ξ^* equal to the stationary means of λ_t and ξ_t, respectively. Since $Q_{g,t} = e^{a_1+a_2\lambda_t}e^{\frac{1}{\epsilon}\log C_t - \frac{1}{\epsilon}\log G_t}$, This implies that log-linearized gold prices are given by

$$\log P_{g,t} = A_g + \frac{1}{\epsilon}\log C_t - \frac{1}{\epsilon}\log G_t + (a_2 + b_{g,\lambda}^*)\lambda_t + b_{g,\xi}^*\xi_t.$$

Similarly, log-linearized platinum prices are given by

$$\log P_{x,t} = A_x + \frac{1}{\epsilon}\log C_t - \frac{1}{\epsilon}\log G_t + (b_2 + b_{x,\lambda}^*)\lambda_t + b_{x,\xi}^*\xi_t + (b_{x,Z}^* - \frac{1}{\epsilon})\log Z_t.$$

The constants A_g and A_x only affect the level of GP and are mainly determined by the scaling term $a_1 - b_1$.

A7. Implied Volatilities in the Model

While the model implies that the GP ratio captures time-variation in the tail risk measures λ_t and ξ_t, can it formally reconcile the empirical evidence that the GP ratio can explain the slope of the implied volatility curve for equity index options? To investigate this, I price options in the model following Eraker and Shaliastovich (2008), Nowotny (2011), and Seo and Wachter (2014). Since the model is in the class of affine jump diffusion models studied by Duffie et al. (2000), the solution for the discounted characteristic function is known up to a system of differential equations.The risk-neutral dynamics of the state variables are also in the class of affine jump diffusions and we can derive the discounted characteristic function where X_t is a vector of the state variables:

$$\mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau} r_s^f ds} e^{uX_{t+\tau}}\right] = e^{\alpha(\tau)+\beta(\tau)X_t} \tag{29}$$

and $\alpha(\tau), \beta(\tau)$ are the solutions to a system of ordinary differential equations as given in Duffie et al. (2000). Nowotny (2011) describes the change of measure in detail.

I approximate the price-dividend ratio using a Taylor approximation following Seo and Wachter (2014), who also show that the approximation is very accurate when the intertemporal elasticity of substitution is equal to one. This allows me to express the price-dividend ratio as:

$$G(\lambda_t, \xi_t) \approx G(\lambda^*, \xi^*) e^{b^*_{\phi,\lambda}(\lambda_t - \lambda^*) + b^*_{\phi,\xi}(\xi_t - \xi^*)} \tag{30}$$

for constants λ^* and ξ^* which I set equal to the long-run mean of λ_t and ξ_t, respectively. Let $A_0 = \log G(\lambda^*, \xi^*) - b^*_{\phi,\lambda}\lambda^* - b^*_{\phi,\xi}\xi^*$. We can express the price of equity at time $t+\tau$ as:

$$P_{t+\tau} = e^{A_0 + d' X_{t+\tau}}, \text{ for } d = [\phi, b^*_{\phi,\lambda}, b^*_{\phi,\xi}]'$$

Following Carr and Madan (1999), Lewis (2000), and Eraker and Shaliastovich (2008), I price options using Fourier transform methods. The forward and inverse Fourier transforms are:

$$\hat{f}(z) = \int_{-\infty}^{\infty} f(x) e^{izx} dx$$

$$f(x) = \frac{1}{2\pi} \int_{iz_i-\infty}^{iz_i+\infty} \hat{f}(z) e^{-izx} dz$$

where $i^2 = -1$ and z_i denotes the imaginary part of the complex variable z. The payoff of a European put option is given by $f(x) = \max(K - e^x, 0)$, whose inverse transform is given by:

$$\hat{f}(z) = -K \frac{K^{iz}}{z^2 - iz}$$

The price of a European put option is a function of the state variables X_t, strike price K, and time to maturity τ, and can be computed as:

$$\begin{aligned} P(X_t, K, \tau) &= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau} r_s^f ds} \max(K - e^{A_0 + d' X_{t+\tau}}, 0)\right] \\ &= \frac{-K}{2\pi} \int_{iz_i - \infty}^{iz_i + \infty} e^{iz_i A_0} \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau} r_s^f ds} e^{-izdX_{t+\tau}}\right] \left(\frac{K^{iz}}{z^2 - iz}\right) dz \\ &= \frac{-K}{2\pi} \int_{iz_i - \infty}^{iz_i + \infty} e^{iz_i A_0} e^{\alpha(\tau) + \beta(\tau) X_t} \left(\frac{K^{iz}}{z^2 - iz}\right) dz \end{aligned} \tag{31}$$

with $\alpha(0) = 0$, $\beta(0) = -izd$, and $z_i < 0$. I normalize strike price and option price by P_t so strike prices can be interpreted as moneyness. Given put option prices, I back out Black and Scholes (1973) implied volatilities using the endogenous dividend yield $\frac{1}{G(\lambda_t, \xi_t)}$ and set the risk-free rate equal to the government bond rate in the model.

Figure 10 plots the difference between the implied volatility of a OTM put option with 0.95 moneyness and an ATM put option with moneyness equal to 1, as a function of the state variables λ_t and ξ_t. The options have 1 month to maturity as in the data. This figure is quite informative because it gives intuition about how λ_t and ξ_t affect return volatility (endogenously determined) and tail risk in the model.[44] In a one-factor model, where λ_t reverts to a constant, both volatility and disaster risk are controlled by the same variable. The combined result is that when λ_t increases, this increases the average level of volatility in addition to the likelihood of negative jumps. Option maturities are typically short (measured in months, for most liquid equity index options), and at this horizon the increase in volatility level is actually stronger than the increase in tail risk, resulting in the implied volatility slope being a decreasing function of λ_t in the one-factor model.

In the two-factor model, ξ_t controls the level of volatility more than it affects the likelihood of jumps, which is directly controlled by λ_t. Therefore, we see that the implied volatility slope is increasing in λ_t, and particularly fast when ξ_t is low. Similarly, the implied volatility slope decreases in ξ_t, and especially quickly when λ_t and ξ_t are low. As shown earlier, GP loads positively on both λ_t and ξ_t. For computational tractability, I first compute the implied volatility slope over a fine mesh of the state variables (λ_t, ξ_t). I then simulate repeated samples of length equal to the data counterpart (17 years). The model implied volatility slope is calculated by interpolating the mesh surface, and regressed on GP in the model. We see a coefficient as large as the data estimate in over 20% of sample paths. Staying within the class of rare disaster models, the two-factor model is necessary to allow the disaster intensity and the implied volatility slope to be positively related.

[44] Seo and Wachter (2014) offer a detailed discussion of this topic.

Tables

Table 1: Summary Statistics for Predictors

Table 1 gives descriptive statistics for the log GP ratio and other known stock return predictors. Monthly data from 1975 - 2013. $\log GP_t$ is the log GP ratio, computed as the log of the ratio of monthly gold to platinum fixing prices. Monthly prices are the average of daily prices. Prior to April 1990, I use monthly average dealer prices for platinum. Gold fixing prices are from the LBMA, and platinum fixing prices are from the LPPM. Platinum dealer prices are from the USGS. $\log PD_t$ is the log price-dividend ratio for the CRSP value-weighted index. $\log PE_t$ is the cyclically adjusted price-earnings ratio from Robert Shiller's website. $\log PNY_t$ is the net payout yield from Michael Roberts' website, available until December 2010. ICC_t is the implied cost of capital from Li et al. (2013), available from January 1977. $DFSP_t$ is the default spread, calculated as the difference between the yield of Baa and Aaa corporate bonds; the data is from FRED. $TMSP_t$ is the term spread, calculated as the difference in yield between a 10 year constant maturity U.S. government bond and a 3 month constant maturity U.S. treasury bill. The data is from FRED. $INFL_t$ is the growth rate of the consumer price index from the FRED. CAY_t is the consumption-wealth ratio from Lettau and Ludvigson (2001) and the data is from Martin Lettau's website, available until March 2013. The data is quarterly, interpolated to a monthly frequency. VRP_t is the variance premium, calculated as the difference between the squared VIX index and annualized realized volatility over the past month. The VIX data is from the CBOE website, and the high-frequency realized variance is from Hao Zhou's website, available from January 1990. ADF is the augmented Dickey and Fuller (1979) test statistic, and p-val is its p-value. The number of lags in the ADF test is selected based on the Ng and Perron (1995) sequential t-test.

Variable	Mean	Std. Dev.	AR(1)	ADF	p-val.	Min.	Max.	Corr. GP_t	Start	End
$\log GP_t$	-0.233	0.266	0.981	-2.872	0.049	-0.850	0.299	1.000	1975.1	2013.12
$\log PD_t$	3.608	0.447	0.994	-1.238	0.657	2.764	4.510	-0.588	1975.1	2013.12
$\log PE_t$	2.878	0.474	0.995	-1.227	0.662	1.893	3.789	-0.539	1975.1	2013.12
$\log PNY_t$	2.247	0.249	0.979	-2.276	0.180	1.700	3.235	-0.461	1975.1	2010.12
ICC_t	7.445	2.694	0.949	-2.759	0.064	-0.040	13.850	0.339	1977.1	2013.12
$DFSP_t$	1.126	0.474	0.961	-4.252	0.001	0.550	3.380	0.336	1975.1	2013.12
$TMSP_t$	1.819	1.261	0.952	-3.458	0.009	-2.650	4.420	0.232	1975.1	2013.12
$INFL_t$	0.322	0.323	0.643	-4.574	0.000	-1.787	1.420	0.027	1975.1	2013.12
CAY_t	0.003	0.018	0.995	-1.798	0.381	-0.035	0.034	0.258	1975.1	2013.03
VRP_t*	0.022	0.024	0.258	-5.700	0.000	-0.217	0.140	0.051	1990.1	2013.12

Table 2: U.S. Stock Return Predictability

Table 2 shows return predictability regressions for the U.S. equity market, January 1975 to December 2013, 468 monthly observations. The regression is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^{f}_{t+i} = \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h}$$

The left hand variable is the excess log return of the CRSP value-weighted index, annualized by the horizon h. The right hand predictor is $\log GP$. Returns are calculated from overlapping monthly data, and t-statistics are based on Newey and West (1987) HAC robust standard errors. The bottom panel estimates a VAR following Hodrick (1992).

VW Excess Returns	1m	3m	6m	1y	2y	3y	4y	5y
OLS Regression								
$\log GP_t$	0.237	0.246	0.260	0.243	0.202	0.161	0.145	0.129
t-stat.	(2.82)	(3.14)	(2.94)	(2.76)	(2.67)	(3.12)	(4.11)	(4.77)
${R_{adj}}^2$ (%)	1.21	4.23	9.55	16.57	23.60	23.57	27.80	31.66
VAR Estimation								
$\log GP_t$	0.236	0.234	0.228	0.215	0.192	0.172	0.155	0.140
t-stat.	(2.75)	(2.89)	(2.82)	(2.66)	(2.69)	(2.62)	(2.61)	(2.60)
${R_{VAR}}^2$ (%)	1.27	3.46	6.34	10.89	16.43	18.99	19.84	19.73

Table 3: Univariate Return Predictability

Table 3 shows univariate return predictability regressions for the U.S. equity market, controlling for other known predictors. January 1975 to December 2013, 468 monthly observations. The regression is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^{f}_{t+i} = \beta_0 + \beta_1 X_t + \epsilon_{t+h}$$

The left hand variable is the excess log return of the CRSP value-weighted index return, annualized by the horizon h. The right hand predictor is $\log GP$. Returns are calculated from overlapping monthly data, and t-statistics use Newey and West (1987) HAC robust standard errors. ${R_{adj}}^2$ is the adjusted R^2 statistic. X_t is a return predictor, including $\log GP$.

	1 month horizon			3 month horizon			1 year horizon			5 year horizon		
	Coef.	t-stat.	R^2_{adj}	Coef.	t-stat.	R^2_{adj}	Coef.	t-stat.	R^2_{adj}	Coef.	t-stat.	R^2_{adj}
$\log GP_t$	0.237	2.82	1.21	0.246	3.14	4.23	0.243	2.76	16.57	0.129	4.77	31.66
$\log PD_t$	-0.067	-1.19	0.10	-0.066	-1.39	0.68	-0.065	-1.30	3.15	-0.064	-2.94	22.12
$\log PE_t$	-0.045	-0.86	-0.05	-0.045	-1.00	0.24	-0.050	-1.04	2.01	-0.049	-1.99	14.70
$\log PNY_t$	-0.096	-0.84	-0.03	-0.084	-0.86	0.21	-0.122	-1.29	3.35	-0.095	-3.52	14.31
ICC_t	0.019	1.98	0.74	0.014	1.75	1.22	0.015	2.33	6.08	0.013	2.86	22.35
$DFSP_t$	0.013	0.18	-0.20	0.016	0.25	-0.16	0.035	0.88	0.90	0.032	1.07	5.23
$TMSP_t$	0.015	0.77	-0.08	0.013	0.77	0.04	0.025	1.80	3.77	0.017	2.15	12.41
$INFL_t$	-0.071	-0.87	-0.02	-0.008	-0.10	-0.21	-0.081	-1.93	2.48	-0.013	-0.68	0.18
CAY_t	0.566	0.46	-0.18	0.743	0.66	-0.03	1.158	1.10	1.49	1.970	3.43	23.03
VRP_t*	5.011	4.32	5.12	4.370	6.27	11.12	1.265	2.60	2.95	-0.487	0.99	1.90

Table 4: Bivariate Return Predictability: Short Horizon

Table 4 shows bivariate return predictability regressions for the U.S. equity market for 1 and 3 month horizons, controlling for other known predictors. January 1975 to December 2013, 468 monthly observations. The regression is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^{f}_{t+i} = \beta_0 + \beta_1 \log GP_t + \beta_2 X_t + \epsilon_{t+h}$$

For the monthly frequency, $h = 1$. The left hand variable is the excess logarithmic return of the CRSP value-weighted index return annualized by the horizon h. The right hand predictors are $\log GP$ and another return predictor X_t. Returns are calculated from overlapping monthly data, and t-statistics use Newey and West (1987) HAC robust standard errors. ${R_{adj}}^2$ is the adjusted R^2 statistic.

	1 month horizon					3 month horizon				
	GP Coef.	GP t-stat.	Coef.	t-stat.	R^2_{adj}	GP Coef.	GP t-stat.	Coef.	t-stat.	R^2_{adj}
$\log PD_t$	0.262	2.71	0.026	0.40	1.03	0.278	3.25	0.032	0.66	4.16
$\log PE_t$	0.274	2.94	0.038	0.68	1.09	0.288	3.45	0.043	0.98	4.33
$\log PNY_t$	0.232	2.44	0.019	0.16	0.80	0.250	2.65	0.040	0.39	3.54
ICC_t	0.180	1.91	0.013	1.24	1.25	0.220	2.37	0.007	0.72	4.12
$DFSP_t$	0.258	2.81	-0.036	-0.47	1.09	0.267	3.00	-0.035	-0.52	4.28
$TMSP_t$	0.233	2.60	0.004	0.18	1.01	0.245	2.82	0.001	0.03	4.02
$INFL_t$	0.239	2.88	-0.077	-0.91	1.22	0.246	3.17	-0.014	-0.17	4.04
CAY_t	0.241	2.64	-0.345	-0.26	0.99	0.250	3.03	-0.201	-0.18	4.03
VRP_t*	0.265	2.79	4.850	4.25	6.91	0.304	3.54	4.177	6.96	18.49

Table 5: Bivariate Return Predictability: Long Horizon

Table 4 shows bivariate return predictability regressions for the U.S. equity market for 1 and 5 year horizons, controlling for other known predictors. January 1975 to December 2013, 468 monthly observations. The regression is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_1 \log GP_t + \beta_2 X_t + \epsilon_{t+h}$$

For the monthly frequency, $h = 1$. The left hand variable is the excess logarithmic return of the CRSP value-weighted index return annualized by the horizon h. The right hand predictors are $\log GP$ and another return predictor X_t. Returns are calculated from overlapping monthly data, and t-statistics use Newey and West (1987) HAC robust standard errors. ${R_{adj}}^2$ is the adjusted R^2 statistic.

	1 year horizon					5 year horizon				
	GP Coef.	GP t-stat.	Coef.	t-stat.	R^2_{adj}	GP Coef.	GP t-stat.	Coef.	t-stat.	R^2_{adj}
$\log PD_t$	0.279	3.06	0.036	0.85	17.03	0.102	1.98	-0.024	-0.69	33.27
$\log PE_t$	0.280	3.05	0.037	0.95	17.26	0.118	2.26	-0.010	-0.29	31.92
$\log PNY_t$	0.231	2.11	-0.008	-0.08	14.76	0.113	2.92	-0.038	-0.91	33.35
ICC_t	0.219	2.15	0.008	1.02	17.79	0.116	5.04	0.009	2.02	44.49
$DFSP_t$	0.252	2.72	-0.014	-0.32	16.54	0.127	3.07	0.004	0.15	31.57
$TMSP_t$	0.228	2.47	0.014	0.84	17.57	0.116	4.96	0.011	1.96	36.51
$INFL_t$	0.248	2.92	-0.089	-1.94	19.65	0.131	4.97	-0.020	-1.40	32.57
CAY_t	0.241	2.71	0.131	0.16	16.40	0.100	4.11	1.172	1.89	38.09
VRP_t*	0.319	3.13	1.028	2.31	30.93	0.174	4.66	-0.579	-1.31	44.12

Table 6: Out-of-Sample Tests

Table 6 shows results for out-of-sample testing, using the out-of-sample R^2 statistic. Let T be the sample length, and m equal to the size of the initial training window (for expanding regressions) or the size of the training window (for rolling regressions). The Out-of-Sample R^2 is given by:

$$R^2_{OS} = 1 - \frac{\sum_{k=1}^{T-m} (r^e_{m+k} - \widehat{r^e}_{m+k})^2}{\sum_{k=1}^{T-m} (r^e_{m+k} - \overline{r^e}_{m+k})^2}$$

For expanding window regressions, the first out-of-sample forecast $\widehat{r^e}_{m+1}$ is based on parameters estimated using observations from 1 to m, the second out-of-sample forecast $\widehat{r^e}_{m+2}$ is based on parameters estimated using observations 1 to $m+1$, and so on. For expanding window regressions, the historical average excess return $\overline{r^e}_{t+1}$ is calculated as the average excess return from time 1 to time t. For rolling window regressions, the first out-of-sample forecast $\widehat{r^e}_{m+1}$ is based on parameters estimated using observations from 1 to m, the second out-of-sample forecast $\widehat{r^e}_{m+2}$ is based on parameters estimated using observations 2 to $m+1$, and so on. For rolling window regressions, the historical average excess return is calculated as the average excess return from over the last m periods, where m is the window length. I consider windows of length 120 months or 180 months to estimate betas, and predict the return in the next month. The In-Sample R^2 is the adjusted R^2. The p-values are calculated using the adjusted-MSPE statistic of Clark and West (2007) given by:

$$f_{t+1} = (r_{t+1} - \overline{r}_{t+1})^2 - \left[(r_{t+1} - \widehat{r}_{t+1})^2 - (\overline{r}_{t+1} - \widehat{r}_{t+1})^2\right]$$

which is regressed against a constant and the test is a one-sided test.

	In-Sample (%)	Out-of-Sample (%)							
		Rolling				Expanding			
Horizon		120m	p-val	180m	p-val	120m	p-val	180m	p-val
1m	1.21	-1.46	0.315	1.44	0.018	0.88	0.033	1.35	0.028
3m	4.23	3.95	0.010	8.28	0.002	4.59	0.009	6.18	0.008
6m	9.55	14.39	0.002	18.29	0.001	10.89	0.010	13.39	0.010
1y	16.57	23.02	0.009	29.37	0.007	19.31	0.021	21.35	0.024
2y	23.60	30.98	0.044	37.28	0.030	25.64	0.041	26.02	0.045
3y	23.57	22.08	0.069	29.30	0.027	23.56	0.033	24.51	0.029
4y	27.80	24.03	0.047	31.39	0.010	29.36	0.008	29.69	0.004
5y	31.66	21.09	0.057	31.88	0.010	33.60	0.010	34.37	0.003

Table 7: International Markets Return Predictability

Table 7 shows return predictability regressions for international equity markets, January 1975 to December 2013, 468 monthly observations. The regression is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h}$$

In Panel A, the left hand variable is the excess log capital gain on the MSCI World Index, annualized by the horizon h. The index is calculated in U.S. dollars and the risk-free rate is the U.S. treasury bill rate. Panel B presents the regression results for individual countries, using the respective MSCI country indices denominated in local currency. The risk-free rate for the U.K is the 3-month U.K Treasury rate from FRED. The risk-free rate for Switzerland is the 3-month Swiss franc interbank rate. The risk-free rate for Japan is the interest rate on Japanese Government Treasury bills from FRED. The risk-free rate for Sweden is the 3-month Swedish Treasury rate from FRED from 1982 onwards. Prior to 1982, I use the historical short-term Swedish interest rates from the Sveriges Riksbank website. Newey and West (1987) HAC robust standard errors.

World Excess Returns	1m	3m	6m	1y	3y	5y
Panel A - World Portfolio						
$\log GP_t$	0.183	0.200	0.214	0.202	0.140	0.111
t-stat.	(2.15)	(2.22)	(2.02)	(1.88)	(1.94)	(2.52)
${R_{adj}}^2$ (%)	0.67	2.62	5.88	9.97	14.91	21.50
Panel B - Individual Countries						
United Kingdom						
$\log GP_t$	0.266	0.239	0.232	0.213	0.156	0.105
t-stat.	(3.05)	(3.30)	(3.02)	(2.70)	(2.90)	(4.41)
${R_{adj}}^2$ (%)	1.19	3.46	7.34	14.15	26.00	25.52
Switzerland						
$\log GP_t$	0.196	0.216	0.240	0.236	0.166	0.129
t-stat.	(2.29)	(2.50)	(2.45)	(2.23)	(2.01)	(2.61)
${R_{adj}}^2$ (%)	0.71	2.56	6.24	11.28	16.61	19.73
Japan						
$\log GP_t$	0.186	0.207	0.221	0.212	0.163	0.148
t-stat.	(1.74)	(1.78)	(1.75)	(1.67)	(1.44)	(1.70)
${R_{adj}}^2$ (%)	0.39	1.72	3.74	6.57	10.34	15.36
Sweden						
$\log GP_t$	0.391	0.427	0.427	0.367	0.190	0.132
t-stat.	(2.72)	(3.00)	(2.69)	(2.34)	(1.82)	(2.70)
${R_{adj}}^2$ (%)	1.54	5.08	8.88	11.51	10.62	11.96

Table 8: Predicting Dividend Growth

Table 8 shows dividend growth predictability regressions, January 1975 to December 2013, 468 monthly observations. The regression in Panel A is:

$$\frac{12}{h}\sum_{i=1}^{h}\Delta d_{t+i} = \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h}$$

The regression in Panel B is:

$$\frac{12}{h}\sum_{i=1}^{h}\Delta e_{t+i} = \beta_0 + \beta_1 \log GP_t + \epsilon_{t+h}$$

Δd_t is the annualized log dividend growth rate calculated from the CRSP value-weighted index, annualized by the horizon h. Δe_t is the annualized log earnings growth rate calculated from the earnings data on Robert Shiller's website. The CPI used to deflate dividends is from FRED. Annual dividend growths are calculated from overlapping monthly data, and t-statistics use Newey and West (1987) HAC robust standard errors.

Panel A - Real Dividend Growth	1y	2y	3y	4y	5y
$\log GP_t$	0.018	0.027	0.028	0.023	0.014
t-stat.	(0.39)	(0.63)	(0.74)	(0.62)	(0.51)
R^2 (%)	0.08	0.41	0.77	0.89	0.57
Panel B - Real Earnings Growth	**1y**	**2y**	**3y**	**4y**	**5y**
$\log GP_t$	0.371	0.199	0.121	0.100	0.057
t-stat.	(1.31)	(1.06)	(0.93)	(1.12)	(1.04)
R^2 (%)	5.08	3.61	2.79	3.48	2.78

Table 9: Cross-Sectional Implications

Table 9 shows the implications of GP risk for the cross-section of stock returns. I first run time-series regressions to estimate betas:

$$R^e_{i,t+1} = c_i + \beta_{i,\Delta gp} \Delta \log GP_{t+1} + \epsilon_{i,t+1}$$

where $R^e_{i,t+1}$ is the excess return for portfolio i and $\Delta \log GP_t + 1 = \log GP_{t+1} - \mathbb{E}_t [\log GP_t]$ is the innovation in GP. The slope coefficient $\beta_{i,\Delta gp}$ represents the portfolio exposure of asset i to GP risk. To estimate the cross-sectional market price of risk associated with GP, I run a cross-sectional regression of time-series average excess returns on the risk exposures:

$$\mathbb{E} \left[R^e_{i,t+1} \right] = \text{cons} + \beta_{i,\Delta gp} \lambda_{\Delta gp} + \upsilon_i.$$

Panel A reports the market price of risk λ with Shanken (1992) t-statistic. Panel B shows the results for $\beta_{i,\Delta gp}$ with Newey and West (1987) HAC robust standard errors. The book-to-market and size portfolio returns are from Kenneth French's website. The data is monthly from 1975 - 2013.

Panel A: Price of Risk	$\lambda_{\Delta gp}$	t-stat
Cross-Section	-0.0217	-6.45
$R^2(\%)$	60.64	
Panel B: Risk Exposures	$\beta_{\Delta gp}$	t-stat
BM1	-0.163	-2.83
BM2	-0.150	-2.46
BM3	-0.137	-2.93
BM4	-0.165	-2.62
BM5	-0.128	-2.67
BM6	-0.170	-2.79
BM7	-0.149	-2.56
BM8	-0.150	-2.59
BM9	-0.177	-3.12
BM10	-0.262	-3.05
SIZE1	-0.277	-4.42
SIZE2	-0.254	-4.05
SIZE3	-0.247	-4.17
SIZE4	-0.233	-4.01
SIZE5	-0.234	-3.85
SIZE6	-0.201	-3.42
SIZE7	-0.213	-3.10
SIZE8	-0.198	-2.97
SIZE9	-0.181	-2.67
SIZE10	-0.132	-2.60

Table 10: GP and Tail Risk

In Panel A, $SLOPE_t^{n\Delta} = \sigma_t^{OTM,\Delta} - \sigma_t^{ATM}$, where $\sigma_t^{OTM,\Delta}$ is the implied volatility of an out-of-the-money put option with $n\Delta$, where $n = 40, 30, 20$, and σ_t^{ATM} is at-the-money implied volatility. Option prices and implied volatilities are from OptionMetrics, January 1996 to August 2013. The regression of the slope of the implied volatility curve for index options against the log GP ratio is

$$\sigma_t^{OTM,\Delta} - \sigma_t^{ATM} = \beta_0 + \beta_1 \log GP_t + \beta_2 \sigma_t^{ATM} + \epsilon_t.$$

The OTM option ranges from deep out-of-the-money (20Δ) to slightly out of the money (40Δ), and the ATM option is defined as a put option with 50Δ. The options have just under one month until expiration, and are taken on the last trading day of the month.

In Panel B, the encompassing regression is:

$$SKEW_t^{\mathbb{Q}} = \beta_0 + \beta_1 \log GP_t + \beta_2 \left(VAR_t^{\mathbb{Q}} \times 100 \right) + \epsilon_t$$

where $SKEW_t^{\mathbb{Q}}$ is as defined in Bakshi et al. (2003) and equal to:

$$SKEW_t^{\mathbb{Q}} = \frac{\mathbb{E}_t^{\mathbb{Q}} \left\{ \left(R_{t,t+\tau} - \mathbb{E}_t^{\mathbb{Q}} \left[R_{t,t+\tau} \right] \right)^3 \right\}}{\mathbb{E}_t^{\mathbb{Q}} \left\{ \left(R_{t,t+\tau} - \mathbb{E}_t^{\mathbb{Q}} \left[R_{t,t+\tau} \right] \right)^2 \right\}^{3/2}}.$$

For both panels, t-statistics use Newey and West (1987) HAC robust standard errors.

	$\log GP_t$		$\sigma_t^{IV,ATM}$		
Panel A: IV Slope	Coef.	t-stat.	Coef.	t-stat.	${R_{adj}}^2$
$SLOPE_t^{40\Delta}$	0.585	(2.74)			9.35
			0.056	(12.54)	59.26
	0.332	(2.97)	0.053	(12.91)	62.13
$SLOPE_t^{30\Delta}$	1.402	(3.10)			11.29
			0.123	(12.84)	60.24
	0.846	(3.85)	0.117	(13.06)	64.22
$SLOPE_t^{20\Delta}$	2.578	(3.29)			13.05
			0.209	(11.55)	59.21
	1.640	(4.12)	0.197	(11.66)	64.33

	$\log GP_t$		$VAR_t^{\mathbb{Q}} \times 100$		
Panel B: BKM	Coef.	t-stat.	Coef.	t-stat.	${R_{adj}}^2$
$SKEW_t^{\mathbb{Q}}$	-0.479	(-2.84)			5.65
					-0.06
			0.500	(3.09)	5.50
	-0.514	(-2.71)			5.36
	-0.587	(-3.58)	0.615	(3.62)	13.97
	-0.560	(-3.07)	0.633	(3.58)	13.67

Table 11: Inventory Financing for Jewellers

Table 11 shows the cash ledger for alternative raw materials inventory financing considerations for a hypothetical jeweller deciding between leasing gold or buying gold on credit. $P_{g,t}$ is the spot price of 1 troy oz of gold in period t. In both scenarios, the gold ledger is flat. $r_{f,t}$ is the borrowing rate for dollars, $r_{p,t}$ is the futures premium over the spot price, and δ_t is the gold lease rate. No-arbitrage and no institutional frictions implies $\delta_t = r_{f,t} - r_{p,t}$.

Gold Leasing	**Period t**	**Period $t+1$**
Lease gold at δ_t, repay in 1 period		$-P_{g,t+1} - \delta_t \times P_{g,t}$
Jewellery fabrication & sales		$P_{g,t+1}$ + Markup
Net Cash Flow	0	Markup $- \delta_t \times P_{g,t}$
Buying Gold on Credit	Period t	Period $t+1$
Borrow money at $r_{f,t}$, repay in 1 period	$+P_{g,t}$	$-(1+r_{f,t}) \times P_{g,t}$
Buy spot gold	$-P_{g,t}$	
Jewellery fabrication & sales		$P_{g,t+1}$ + Markup
Short gold futures		$(1+r_{p,t}) \times P_{g,t} - P_{g,t+1}$
Net Cash Flow	0	Markup $- (r_{f,t} - r_{p,t}) \times P_{g,t}$

Table 12: Gold and Platinum Returns

Table 12 estimates gold and platinum returns. All estimates are annualized real percentage terms, except for the Sharpe ratios. The data is monthly from 1975 - 2013. The spot prices are calculated from the LBMA fixing price (gold), LPPM fixing price (platinum). Before April 1990, I use USGS dealer prices for platinum. Futures prices are based on the futures contracts with closest to 3 months to maturity from the COMEX division of the CME (formery NYMEX). The interest rate is the U.S. dollar Libor rate (before 1986, the Eurodollar deposit rate). CPI data is from FRED, and the risk-free rate is the 1 month U.S. Treasury bill rate. δ is the dividend yield (for stocks) or lease rate (convenience yield, for gold and platinum).

Equity, Gold, and Platinum Returns

Equity		Gold		Platinum	
Variable	Data	Variable	Data	Variable	Data
$\mathbb{E}[R^m - R^b]$	7.53	$\mathbb{E}[R^g - R^b]$	2.40	$\mathbb{E}[R^x - R^b]$	6.51
$\sigma(R^m)$	15.11	$\sigma(R^g)$	16.76	$\sigma(R^x)$	22.18
Sharpe Ratio	0.50	Sharpe Ratio	0.14	Sharpe Ratio	0.29
$\mathbb{E}[\delta^m]$	2.71	$\mathbb{E}[\delta^g]$	1.00	$\mathbb{E}[\delta^x]$	3.47

Table 13: Model Parameters

Parameter values below are in annual terms. The third column gives the reference source for the parameter. "Standard" means that the parameter value is in a standard range of values commonly used in the literature. "Data" means that this parameter choice is disciplined by the data. Citations mean this parameter is from the cited paper. "Bounded Parameter" means this parameter is free to be set within bounds imposed by either existence of solutions (for a_2 and b_2) or by data estimates (for μ_λ). "Fixed" is for the intratemporal elasticity of substitution and volatility of ξ_t, as discussed in the main text. NSBU (2013) refers to Nakamura et al. (2013) and BU (2008) refers to Barro and Ursua (2008).

Parameter	Value	Reference
Relative risk aversion γ	3	Standard
Subject time preference δ	0.012	Standard
Mean log consumption growth (normal times) $\bar{g_c}$	0.025	Data
Volatility of log consumption growth (normal times) σ_c	0.020	Data
Leverage ϕ	2.6	Wachter (2013)
Mean-reverting target of ξ_t process $\bar{\xi}$	0.0355	Wachter (2013)
Rate of mean reversion κ_λ	0.25	Target AR1($\log GP$)
Rate of mean reversion κ_ξ	0.095	Target AR1($p-d$)
Volatility of λ_t process σ_λ	0.183	Target $\sigma(\log GP)$
Volatility of ξ_t process σ_ξ	0.0861	Fixed
Probability of default given disaster q	0.40	Barro (2006), Wachter (2013)
Average disaster size $\mathbb{E}_\nu[1-\exp J_t^c]$	0.15	BU (2008), Nakamura et al (2013)
Intratemporal elasticity of substitution ϵ	$1/\phi$	Fixed
Log growth rate in gold and platinum stock μ_g	0.0072	Data
Volatility of log gold stock growth σ_g	0.0021	Data
Mean-reverting target of $\log Z_t$ process μ_z	-3.114	Data
Rate of $\log Z_t$ mean reversion θ_z	0.022	Data
Volatility of $\log Z_t$ process σ_z	0.003	Data
Scaling term $a_1 - b_1$	6.34	Match $\mathbb{E}[\log GP]$
Intensity jump mean μ_λ	0.03	Bounded Parameter
Gold preference loading a_2	5.73	Bounded Parameter
Platinum preference loading b_2	1.25	Bounded Parameter

Table 14: Simulation Results: Asset Pricing Moments

Table 14 shows results from model simulations for stocks, bonds, gold, and platinum. State variables are simulated at a monthly frequency and aggregated to an annual frequency. The population moments are computed from a 1,000,000 year simulation. The model confidence intervals are computed from 100,000 simulations of length equal to the length of the data. Data moments are calculated using monthly observations, from 1975 to 2013 and annualized. Expected returns, yields, and volatilities are given in real percentage terms. δ represents the dividend yield (for stocks) or lease rate (for gold and platinum). $p-d$ is the log price-dividend ratio.

	Data	All Simulations				No Disaster Simulations		
	Est.	5%	50%	95%	Pop.	5%	50%	95%
Stocks and Bonds								
$\mathbb{E}[R_m - R_f]$	7.53	3.30	6.77	12.69	7.23	3.63	6.38	10.68
$\sigma(R_m)$	15.18	11.77	20.27	31.99	21.53	10.35	16.32	24.99
$\mathbb{E}[\delta^m]$	2.71	1.34	1.81	3.76	2.08	1.29	1.56	2.40
$\mathbb{E}[R_f]$	1.11	-1.06	2.46	3.50	2.00	1.93	3.16	3.57
$\sigma(R_f)$	1.07	0.29	1.71	6.96	3.51	0.17	0.73	2.06
$\sigma(p-d)$	0.45	0.12	0.28	0.58	0.44	0.10	0.21	0.44
$AR1(p-d)$	0.92	0.56	0.82	0.94	0.92	0.50	0.78	0.92
Gold								
$\mathbb{E}[R_g - R_f]$	2.40	0.23	2.62	5.74	2.78	0.88	2.67	4.96
$\sigma(R_g)$	16.76	6.19	10.38	20.64	12.48	5.67	7.60	11.15
$\mathbb{E}[\delta^g]$	1.00	0.59	0.77	1.79	0.93	0.58	0.66	0.96
Gold Sharpe	0.14	0.02	0.26	0.52	0.22	0.13	0.35	0.57
Platinum								
$\mathbb{E}[R_x - R_f]$	6.51	2.37	5.34	10.22	5.70	2.70	5.05	8.55
$\sigma(R_x)$	22.18	9.37	15.28	23.98	16.27	8.40	12.30	17.49
$\mathbb{E}[\delta^x]$	3.47	2.33	3.01	5.45	3.34	2.26	2.63	3.67
Platinum Sharpe	0.29	0.15	0.37	0.60	0.35	0.24	0.42	0.62
GP Ratio								
$\mathbb{E}[\log GP]$	-0.23	-0.37	-0.25	0.15	-0.23	-0.38	-0.31	-0.09
$\sigma(\log GP)$	0.26	0.05	0.14	0.42	0.28	0.04	0.10	0.29
$AR1(\log GP)$	0.79	0.50	0.78	0.92	0.91	0.46	0.74	0.90

Table 15: Simulation Results: Return Predictability

Table 15 shows results from model simulations for return predictability. State variables are simulated at a monthly frequency and aggregated to an annual frequency. The population moments are computed from a 1,000,000 year simulation. The model confidence intervals are computed from 100,000 simulations of length equal to the length of the data. I run the regression in the first panel is:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_{h,gp}\log GP_t + \epsilon_{t+h}$$

and the period is annual. Excess returns are log equity returns over the log return on the government bond, as in Barro (2006) and Wachter (2013). The regressions I run in the second panel are:

$$\log GP_t = \beta_0 + \beta_Z \log Z_t + \epsilon_t$$

and predicting returns using $\log Z_t$

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_{h,Z}\log Z_t + \epsilon_{t+h}.$$

In the third panel, I run the return predictability regression controlling for supply effects:

$$\frac{12}{h}\sum_{i=1}^{h}\log R_{t+i} - \log R^f_{t+i} = \beta_0 + \beta_{h,gp\perp Z}\log GP_t + \gamma_{h,Z}\log Z_t + \epsilon_{t+h}.$$

	Data	All Simulations				No Disaster Simulations		
	Est.	5%	50%	95%	Pop.	5%	50%	95%
Excess Returns								
$\beta_{1y,gp}$	0.243	-0.007	0.334	1.023	0.132	0.095	0.444	1.148
$\beta_{3y,gp}$	0.161	-0.006	0.303	0.753	0.127	0.100	0.390	0.818
$\beta_{5y,gp}$	0.129	-0.010	0.270	0.598	0.121	0.090	0.333	0.635
$R^2_{1y,gp}$	16.57	0.28	7.69	23.01	3.07	1.22	10.36	25.94
$R^2_{3y,gp}$	23.57	0.76	20.63	49.42	8.85	3.31	26.50	53.13
$R^2_{5y,gp}$	31.66	1.08	30.06	63.78	13.98	3.92	37.04	66.72
Supply Regressions								
β_Z	7.86	-23.20	1.60	26.80	1.69	-14.64	1.59	18.57
R^2_Z	11.35	0.09	9.10	49.09	0.77	0.07	8.41	46.15
$\beta_{1y,Z}$	-0.55	-10.19	0.01	10.26	0.00	-7.49	0.03	7.63
$R^2_{1y,Z}$	0.16	0.00	1.12	9.27	0.00	0.00	1.03	8.80
Excess Returns Controlling for Supply								
$\beta_{1y,gp\perp Z}$	0.283	0.004	0.392	1.154	0.132	0.111	0.512	1.286
$\beta_{3y,gp\perp Z}$	0.188	0.005	0.343	0.817	0.127	0.115	0.432	0.884
$\beta_{5y,gp\perp Z}$	0.145	0.002	0.293	0.634	0.121	0.105	0.358	0.668
$R^2_{1y,gp\perp Z}$	19.99	1.65	11.25	28.00	3.07	3.34	13.88	30.58
$R^2_{3y,gp\perp Z}$	28.84	4.58	28.55	56.53	8.85	9.28	33.81	59.51
$R^2_{5y,gp\perp Z}$	35.65	6.88	40.28	70.68	13.98	12.80	45.87	72.80

Table 16: Gold and Platinum Supply Dynamics

Table 16 shows the per-capita growth rate of the world gold and platinum stock, annual data from 1975 to 2013. The data is calculated from annual world production data. The data for platinum production is from Johnson Matthey. The data for gold production is from the U.S. Geological Survey Minerals Yearbook. The estimate of the initial gold stock is 84,000 tonnes (Thomas and Boyle (1986)), and the estimate of the initial platinum stock (net of autocatalysts) is set to be 4.5% of the initial gold stock. Population growth is proxied by the U.S. annual population growth from the U.S. census. G_t is the per-capita gold stock and X_t is the per-capita platinum stock. Panel A gives descriptive statistics for the log growth rates of per-capita gold and platinum stock. Panel B shows results from augmented Dickey-Fuller tests. Panel C reports the estimate of the cointegrating coefficient β_G from a Stock and Watson (1993) Dynamic Least Squares (DLS) regression:

$$\log X_t = \beta_0 + \beta_G \log G_t + \sum_{i=-k}^{k} \gamma_i \Delta \log G_{t-i} + \epsilon_t$$

using Newey and West (1987) HAC robust standard errors. Panel D results of the Johansen (1988) rank test for the VECM:

$$\Delta Y_t = \mu + \Pi Y_{t-1} + \sum_{j=1}^{p-1} \Gamma_j \Delta Y_{t-j} + \epsilon_t$$

where $Y_t = \begin{bmatrix} \log X_t & \log G_t \end{bmatrix}'$. The null hypothesis is $H_0 : rank(\Pi) = r$ against the alternative $H_a : rank(\Pi) > r$.

Panel A - Growth Rates	Mean (%)	Median (%)	Std. Dev. (%)	Min. (%)	Max. (%)
$\Delta \log G_t$	0.72	0.76	0.21	0.21	1.00
$\Delta \log X_t$	0.71	0.68	0.29	0.24	1.28
Panel B - ADF	1 lag	2 lags	3 lags	4 lags	5 lags
$\log G_t$	1.046	1.645	1.326	1.464	1.580
p-val	0.995	0.998	0.997	0.997	0.998
$\log X_t$	2.592	1.954	1.879	1.259	1.319
p-val	0.999	0.999	0.999	0.996	0.997
$\log Z_t$	-1.776	-3.085	-3.032	-2.781	-3.024
p-val	0.392	0.028	0.032	0.061	0.033
Panel C - DLS	Estimate	Std. Err.	t-stat	[95% confidence	interval]
β_G (k=1)	0.990	0.034	29.39	0.92	1.06
β_G (k=2)	1.006	0.035	29.02	0.94	1.08
β_G (k=3)	1.038	0.044	23.38	0.95	1.13
Panel D - Rank Test	Statistic	90% CV	95% CV	p-val	$H_0 = r$
1 lag in VAR	11.61	13.43	15.50	0.177	0
	0.99	2.71	3.84	0.448	1
2 lags in VAR	23.28	13.43	15.50	0.004	0
	3.10	2.71	3.84	0.078	1
3 lags in VAR	26.01	13.43	15.50	0.001	0
	2.53	2.71	3.84	0.112	1
4 lags in VAR	21.76	13.43	15.50	0.005	0
	0.84	2.71	3.84	0.504	1

Figures

Figure 1: Gold and Platinum Prices

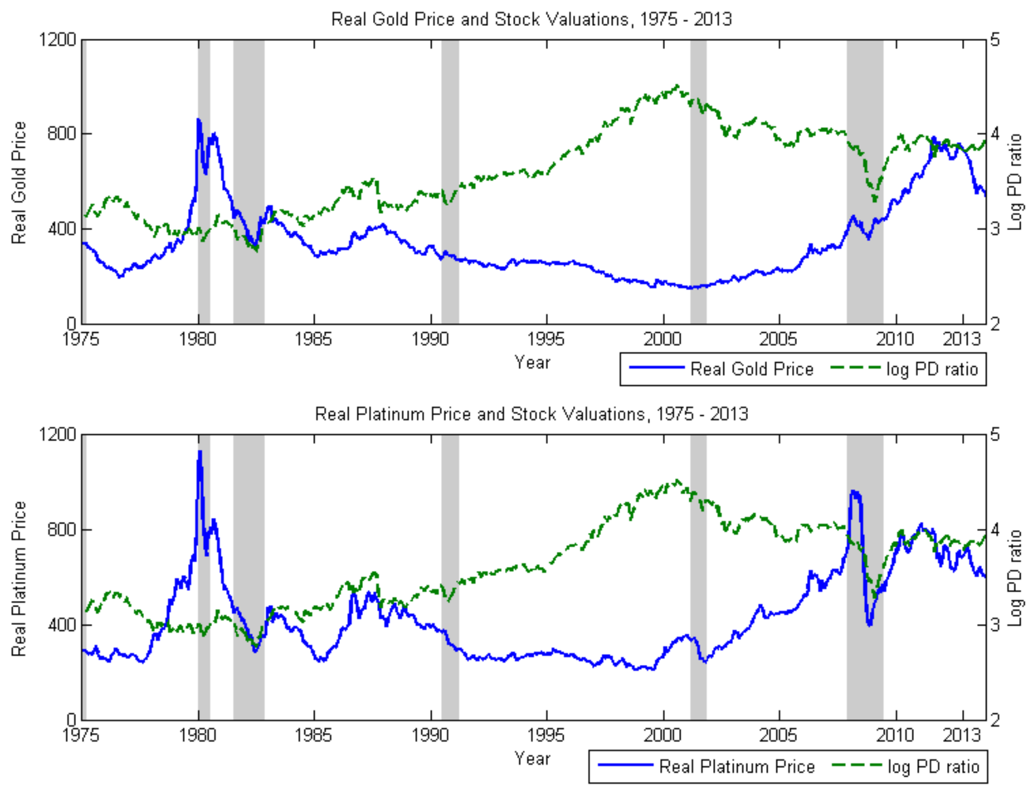

The top panel shows the behavior of real gold prices (solid line) and the log price-dividend ratio on the CRSP value-weighted portfolio (dashed line) from 1975 - 2013. The bottom panel shows real platinum prices (solid line) and the log price-dividend ratio on the CRSP value-weighted portfolio (dashed line) from 1975 - 2013. The shaded grey bars are NBER recessions.

Figure 2: Log GP Ratio

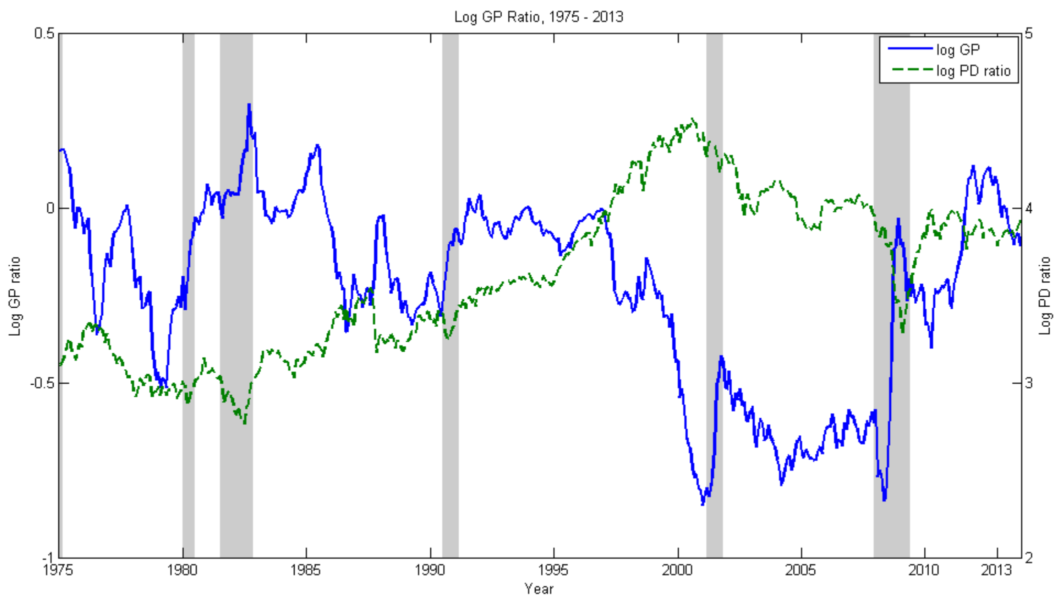

The figure above shows the natural logarithm of the ratio of gold to platinum prices (log GP ratio) from 1975 to 2013. The data is monthly frequency. Gold data is from LBMA, and platinum data is from LPPM and the U.S. Geological Survey. Shaded bars represent NBER recessions.

Figure 3: Rolling Regressions - GP ratio

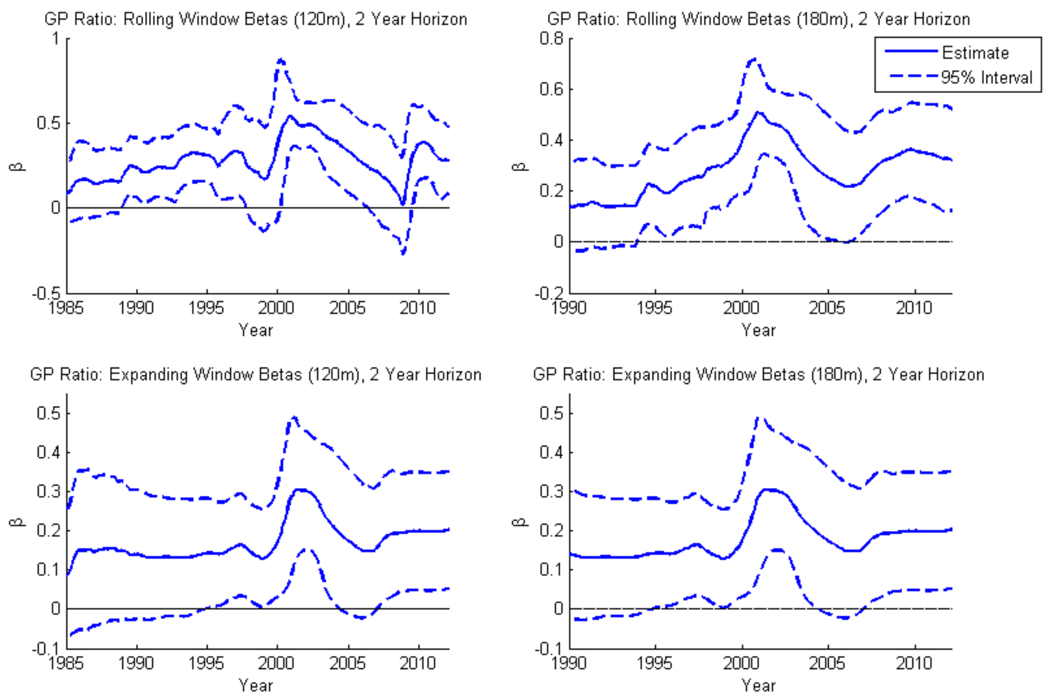

The figure above shows estimated betas and 95% confidence intervals for 2-year ahead predictive regressions of future U.S. stock market excess returns by the log GP ratio. The top left is for the rolling window method with a 120 month window. The top right is for the rolling window method with a 180 month window. The bottom left is for the expanding window method with a 120 month window. The bottom right is for the expanding window method with a 180 month window.

Figure 4: Rolling Regressions - PD ratio

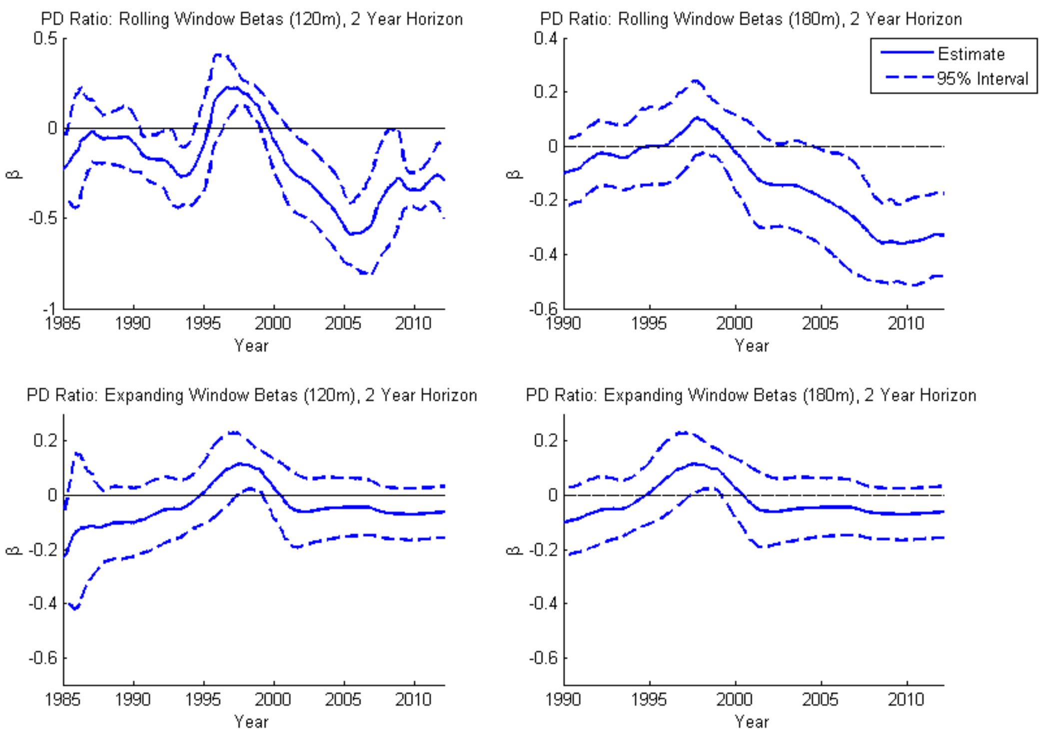

The figure above shows estimated betas and 95% confidence intervals for 2-year ahead predictive regressions of future U.S. stock market excess returns by the log price-dividend (PD) ratio. The top left is for the rolling window method with a 120 month window. The top right is for the rolling window method with a 180 month window. The bottom left is for the expanding window method with a 120 month window. The bottom right is for the expanding window method with a 180 month window.

Figure 5: Cross-Sectional Pricing

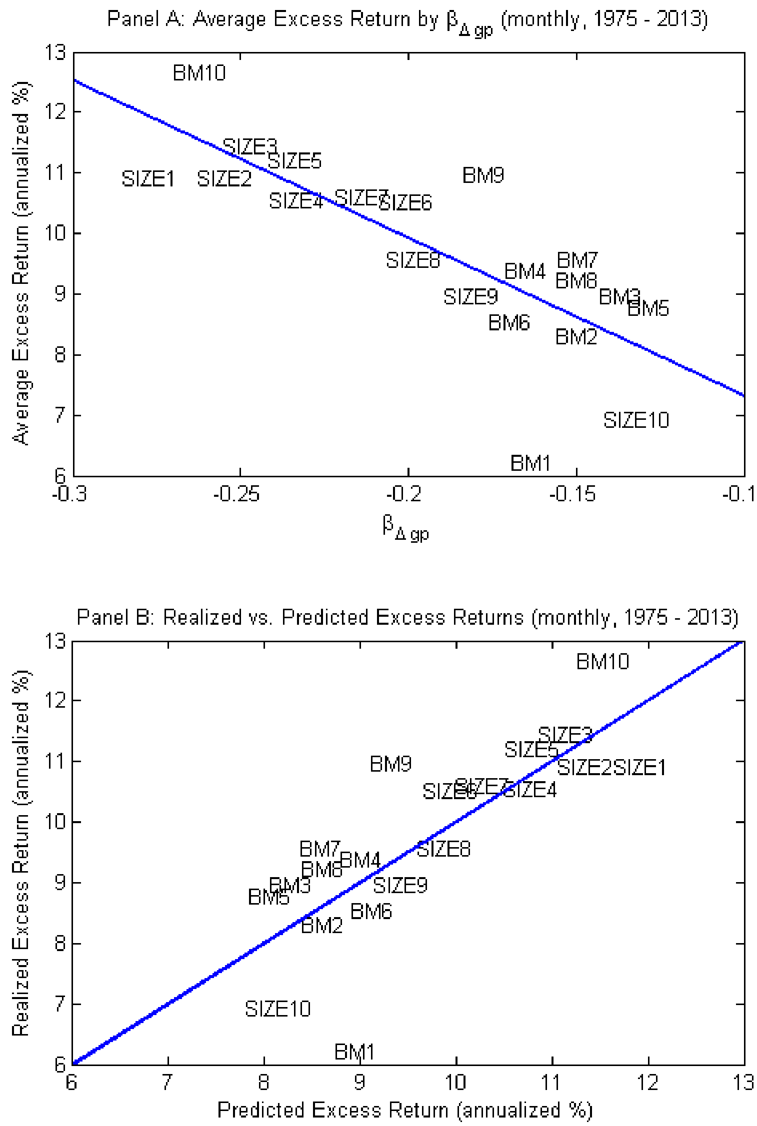

Panel A in the figure above shows the realized average excess returns for book-to-market and size portfolios against the risk exposures (betas). Panel B shows the realized average excess returns against the predicted excess returns. Results are based on the one-factor model with only GP risk. Monthly data 1975 - 2013, annualized percentage excess returns.

Figure 6: Gold and Platinum Demand

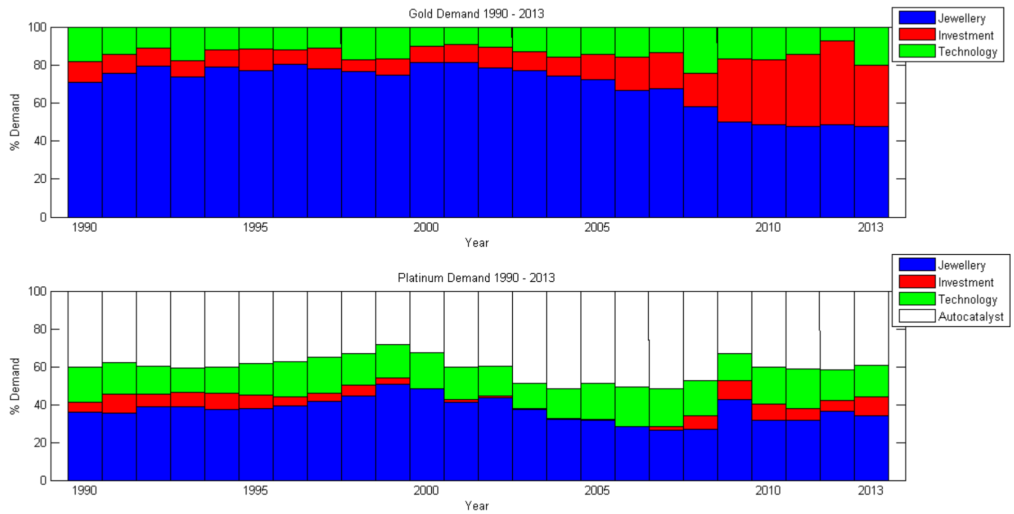

The figure above shows gold and platinum demand for jewellery and investment as a percentage of total demand in a given year, from 1990 - 2013. Gold data is from Thomson Reuters Gold Fields Mineral Services (GFMS) and platinum data is from Johnson Matthey.

Figure 7: Gold Lease Rates 2007 - 2009

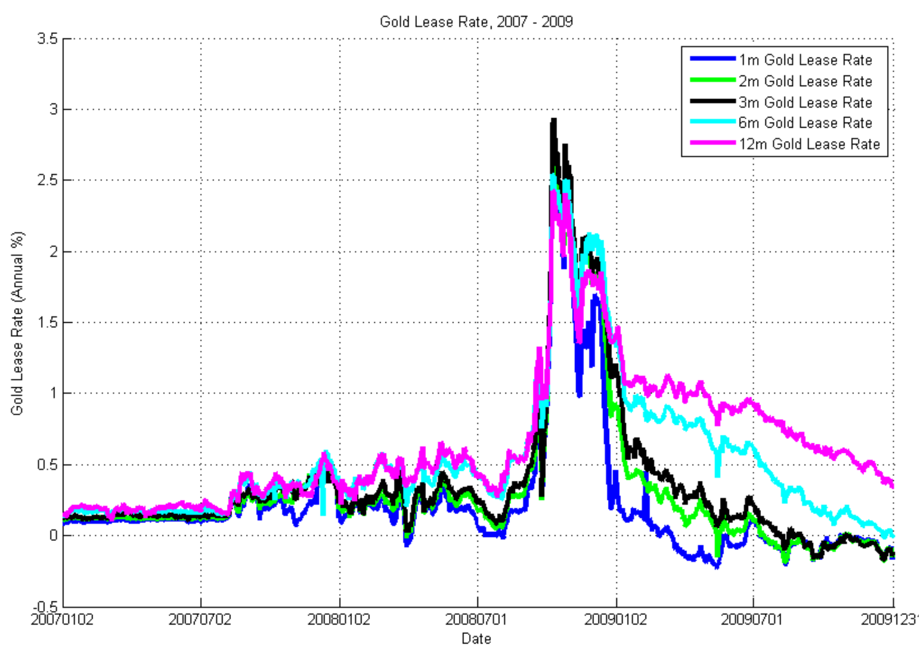

The figure plots the annualized gold lease rates in percentages from 2007 - 2009. The data is computed as the Libor rate minus the Gold Forward Offered Rate (GOFO) as published by the LBMA.

Figure 8: Multinomial Disaster Size Distribution

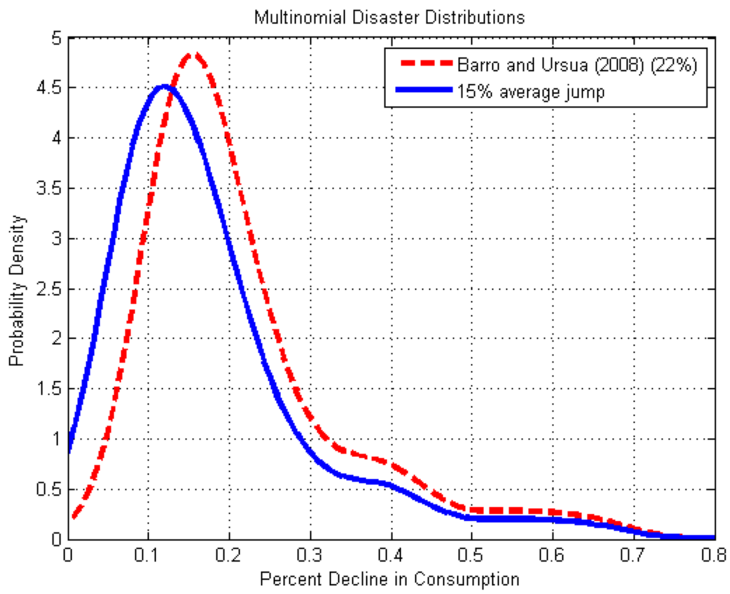

The figure shows the disaster size distribution for the quantity $1 - e^{J_t^c}$ in the model. The red dashed line is the distribution from the Barro and Ursua (2008) calibration, which is also used in Wachter (2013) with an average jump size of 22% using a 10% cutoff to identify disasters. The blue solid line is the distribution used in this paper with an average jump size of 15% using a 6% cutoff to identify disasters. The plots show distributions smoothed by a kernel estimator with a bandwith of 0.05. The model uses draws from the exact multinomial distribution.

Figure 9: Per-Capita Gold and Platinum Stock Growth

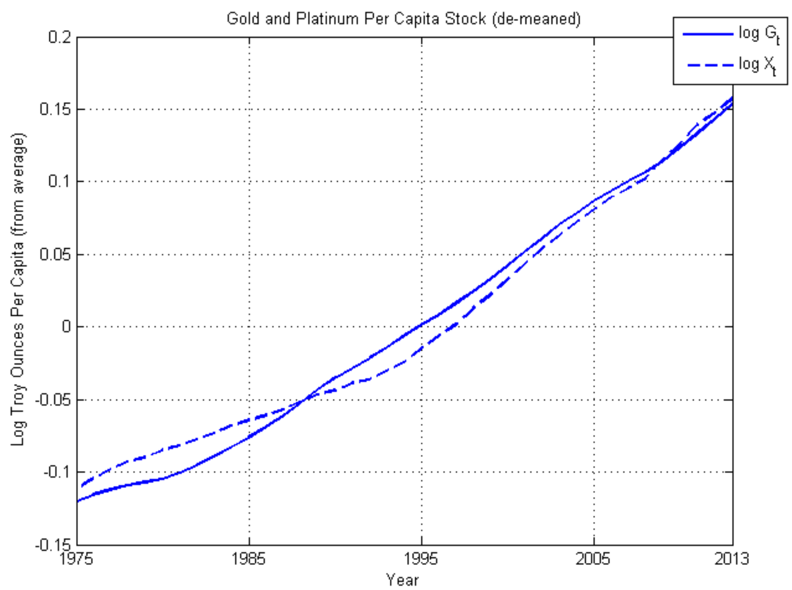

The figure shows the log per-capita gold and platinum stock (de-meaned) from 1975 to 2013. The data is annual frequency. Gold stock is calculated from world supply data from the U.S. Geological Survey, and platinum stock is calculated from world supply data from Johnson Matthey.

Figure 10: Implied Volatility Slope by Disaster Intensities

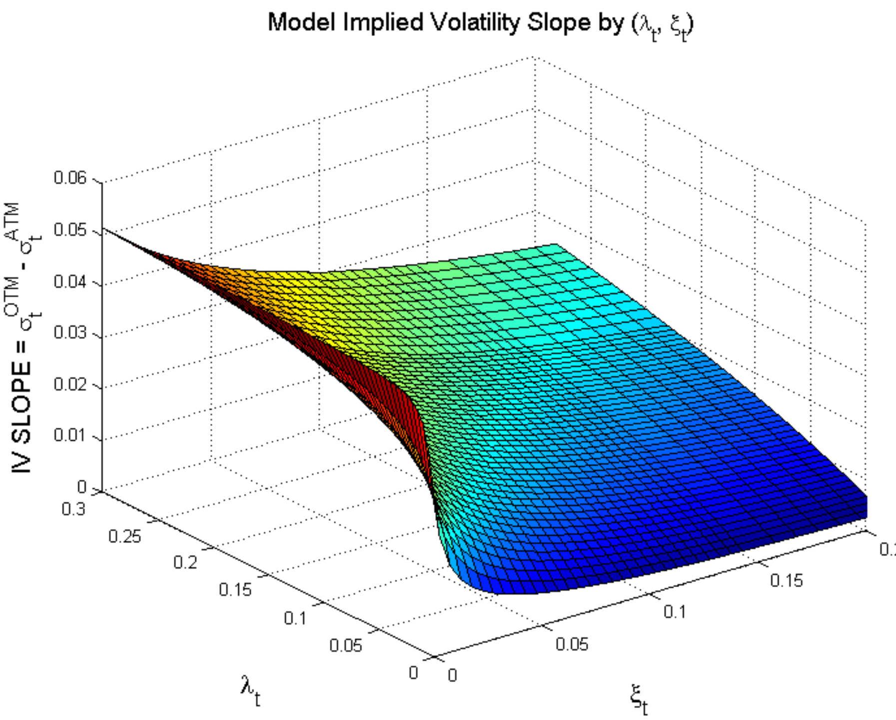

The figure plots the implied volatility slope, calculated as the difference between a out-of-the-money put option with $\frac{strike}{spot} = 0.95$ and an at-the-money put option with $\frac{strike}{spot} = 1$ with 1 month to maturity, as a function of the state variables λ_t and ξ_t. Implied volatility measures are calculated by inverting the Black-Scholes formula using the model implied price-dividend ratio and government bond rate.

References

Abel, Andrew, 1999, Risk premia and term premia in general equilibrium, *Journal of Monetary Economics* 43, 3–33.

Aït-Sahalia, Yacine, Jonathan Parker, and Motohiro Yogo, 2004, Luxury goods and the equity premium, *Journal of Finance* 59, 2959–3004.

Alizadeh, Sassan, Michael Brandt, and Francis X. Diebold, 2002, Range-based estimation of stochastic volatility models, *Journal of Finance* 57, 1047–1091.

Ang, Andrew, and Geert Bekaert, 2007, Stock Return Predictability: Is It There?, *Review of Financial Studies* 20, 651–707

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Bansal, Ravi, and Ivan Shaliastovich, 2011, Confidence risk and asset prices, *American Economic Review: Papers and Proceedings* 100, 537–541.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 54, 1481–1509.

Barro, Robert, 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

Barro, Robert, and Sanjay Misra, 2013, Gold returns, Working paper.

Barro, Robert, and Jose Ursua, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* 255–350.

Basel I, 1988, International convergence of capital measurement and capital standards. Basel Committee on Banking Supervision, Bank of International Settlements .

Basel II, 2004, Basel II: International convergence of capital measurement and capital standards: a revised framework. Basel Committee on Banking Supervision, Bank of International Settlements .

Basel III, 2012, Basel III counterparty credit risk and exposures to central counterparties - frequently asked questions. Basel Committee on Banking Supervision, Bank for International Settlements .

Bates, David, 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert Goldstein, 2011, Explaining asset pricing puzzles associated with the 1987 market crash, *Journal of Financial Economics* 101, 552–573.

Bernstein, Peter, 2012, *The Power of Gold: The History of an Obsession* (Wiley).

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Black, William, 2000, *The Platinum Group Metals Industry* (Woodhead Publishing Limited).

Bollerslev, Timothy, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael Roberts, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62, 877–915.

Brennan, Michael, 1958, The supply of storage, *American Economic Review* 48, 50–72.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.

Campbell, John, 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.

Campbell, John, and Robert Shiller, 1988, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

Campbell, John, and Samuel Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509–1531

Carr, Peter, and Dilip Madan, 1999, Option pricing and the fast fourier transform, *Journal of Computational Finance* 2, 61–73.

Casassus, Jaime, and Pierre Collin-Dufresne, 2005, Stochastic convenience yield implied from commodity futures and interest rates, *Journal of Finance* 60, 2283–2331.

Cawthorn, Grant, 1999, Seventy-fifth anniversary of the discovery of the platiniferous merensky reef, *Platinum Metals Review* 43, 146–148.

Chernov, Mikhail, Ronald Gallant, Eric Ghysels, and George Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225–257.

Clark, Todd, and Kenneth West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Cochrane, John, 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.

Cox, John, Jonathan Ingersoll, and Stephen Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.

Cuoco, Domenico, and Hong Liu, 2000, Optimal consumption of a divisible durable good, *Journal of Economic Dynamics and Control* 24, 561–613.

Dickey, David, and Wayne Fuller, 1979, Distribution of the estimators for autoregressive time series with a unit root, *Journal of the American Statistical Association* 74, 427–431.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it?, *Review of Financial Studies* 24, 1–45

Duffie, Darrel, and Costis Skiadas, 1994, Continuous-time asset pricing: A utility gradient approach, *Journal of Mathematical Economics* 23, 107–132.

Duffie, Darrell, and Larry Epstein, 1992, Asset pricing with stochastic differential utility, *Review of Financial Studies* 5, 411–436.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Elliot, Graham, and James Stock, 1994, Inference in time series regression when the order of integration of a regressor is unknown, *Econometric Theory* 10, 672–700.

Engle, Robert, and Clive Granger, 1987, Co-integration and error correction: Representation, estimation, and testing, *Econometrica* 55, 251–276.

Epstein, Larry, and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Eraker, Bjørn, and Ivan Shaliastovich, 2008, An equilibrium guide to designing affine pricing models, *Mathematical Finance* 18, 519–543.

Erb, Claude, and Campbell Harvey, 2013, The golden dilemma, *Financial Analysts Journal* 69, 10–42.

Fama, Eugene, and Kenneth French, 1988, Business cycles and the behavior of metals prices 43, 1075–1093.

Fama, Eugene, and Kenneth French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gomes, João, Leonid Kogan, and Motohiro Yogo, 2009, Durability of output and expected stock returns, *Journal of Political Economy* 117, 941–986.

Gourio, Francois, 2012, Disaster risk and business cycles, *American Economic Review* 102, 2734–2766.

Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Hall, Robert, 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–357.

Hansen, Lars P., John Heaton, Junghoon Lee, and Nikolai Roussanov, 2007, Chapter 61: Intertemporal substitution and risk aversion, in *Handbook of Econometrics, Volume 6, Part A*, 3967.

Hirschey, Mark, and John Nofsinger, 2008, *Investments: Analysis and Behavior* (McGraw-Hill).

Hodrick, Robert, 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.

Johansen, Soren, 1988, Statistical analsyis of cointegration vectors, *Journal of Economic Dynamics and Control* 12, 231–254.

Johansen, Soren, 1995, *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models* (Oxford: Oxford University Press).

Jurek, Jakub, 2014, Crash-neutral currency carry trades, *Journal of Financial Economics* 113, 325–347.

Kandel, Shmuel, and Robert Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance* 51, 385–424.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *Review of Financial Studies* 27, 2841–2871.

Kreps, David, and Evan Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46, 185–200.

Lannoye, Vincent, 2011, *The History of Money for Understanding Economics* (CreateSpace Independent Publishing).

LBMA, 2009, Gold lease rates: A rapidly changing market, *Alchemist* 53.

LBMA and LPPM, 2008, A guide to the london precious metals markets .

Le, Anh, and Haoxiang Zhu, 2013, Risk premia in gold lease rates, Working paper.

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.

Lewis, Alan, 2000, *Option Valuation Under Stochastic Volatility* (Finance Press).

Li, Yan, David Ng, and Bhaskaran Swaminathan, 2013, Predicting market returns using aggregate implied cost of capital, *Journal of Financial Economics* 110, 419–436.

Lintner, John, 1975, Inflation and security returns, *Journal of Finance* 30, 259–280.

Lucas, Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Madison, Angus, 2003, *The World Economy: Historical Statistics* (Paris, OECD).

Manela, Asaf, and Alan Moreira, 2014, News implied volatility and disaster concerns, Working paper.

Nakamura, Emi, Jon Steinsson, Robert Barro, and Jose Ursua, 2013, Crises and recoveries in an empirical model of consumption disasters, *American Economic Journal: Macroeconomics* 5, 35–74.

Newey, Whitney, and Kenneth West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 59, 817–858.

Ng, Serena, and Pierre Perron, 1995, Unit root tests in arma models with data-dependent methods fort the selection of the truncation lag, *Journal of the American Statistical Association* 90, 268–281.

Nowotny, Michael, 2011, Disaster begets crisis: The role of contagion in financial markets, Working paper.

Ogaki, Masao, and Carmen Reinhart, 1998, Measure intertemporal substitution: The role of durable goods, *Journal of Political Economy* 106, 1078–1098.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Pastor, Lubos, Meenakshi Sinha, and Bhaskaran Swaminathan, 2008, Estimating the intertemporal risk-return tradeoff using the implied cost of capital, *Journal of Finance* 63, 2859–2897.

Pastor, Lubos, and Robert Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Rietz, Thomas, 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.

Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.

Schwartz, Eduardo, 1997, The stochastic behavior of commodity prices: Implications for valuation and hedging, *Journal of Finance* 52, 923–973.

Seo, Sang Byung, and Jessica Wachter, 2014, Option prices in a model with stochastic disaster risk, Working paper.

Shanken, Jay, 1992, On the estimation of beta-pricing models, *Review of Financial Studies* 5, 1–55.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Stock, James, and Mark Watson, 1993, A simple estimator of cointegrating vectors in higher order integrated systems, *Econometrica* 61, 783–820.

Tauchen, George, and Victor Todorov, 2011, Volatility jumps, *Journal of Business and Economic Statistics* 29, 235–371.

Thomas, P.R., and E.H. Boyle, 1986, Gold availability - world, a minerals availability program appraisal, *U.S. Bureau of Mines, Washington D.C.* .

Torous, Walter, Rossen Valkanov, and Shu Yan, 2004, On predicting stock returns with nearly integrated explanatory variables, *Journal of Business* 77, 937–966.

Tufano, Peter, 1996, Who manages risk? An empirical examination of risk management practices in the gold mining industry, *Journal of Finance* 51, 1097–1137.

Vissing-Jørgensen, Anette, 2002, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.

Wachter, Jessica, 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987–1035

Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539–580.