Hedging Sanctions Risk: Cryptocurrency in Central Bank Reserves

Matthew Ferranti*

Department of Economics, Harvard University

November 17, 2022

Abstract

Central banks may shift their international reserve holdings in order to protect themselves ex-ante against the risk of financial sanctions by fiat reserve currency issuers. For example, from 2016 to 2021, countries facing a higher risk of US sanctions increased the gold share of their reserves more than countries facing a lower risk of US sanctions. This paper explores the potential for Bitcoin to serve as an alternative hedging asset. I describe a dynamic Bayesian copula model to simulate the joint returns of Bitcoin and other reserve assets under a wide range of plausible sanctions probabilities. Assuming mean-variance preferences, a modest risk of sanctions significantly increases optimal gold and Bitcoin allocations. If a central bank cannot acquire sufficient physical gold to hedge its sanctions risk, the optimal Bitcoin share rises further, suggesting that gold and Bitcoin are imperfect substitutes. I conclude that sanctions risk may diminish the appeal of US Treasuries, propel broader diversification in central bank reserves, and bolster the long-run fundamental value of both cryptocurrency and gold.

Keywords: Central Banks and Their Policies, Foreign Exchange, Bayesian Computation

*ferranti@g.harvard.edu. The author thanks Kenneth Rogoff, Neil Shephard, Jeffrey Frankel, Adrien Bilal, Marc Melitz, Pol Antras, and participants at the Harvard International Lunch for helpful comments. This paper is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1745303. The author declares no conflicts of interest.

1 Introduction

As cryptocurrencies have become increasingly mainstream vehicles for investing and transferring wealth, governments have begun to explore potential applications of the technology. Most prominently, El Salvador began adding Bitcoin to its reserves in September 2021, accumulating over 2,381 Bitcoin worth $57 million as of July 2022. In April 2022, the Central African Republic adopted Bitcoin as legal tender alongside the CFA franc. And in February 2022, Ukraine began accepting cryptocurrency donations to fund its military and purchase humanitarian aid during Ukraine's war with Russia, ultimately receiving $100 million from supporters across the globe.

The global financial sanctions enforced against Russia following its February 2022 invasion of Ukraine are unprecedented in their scope. Never before has an economy the size of Russia's – the 11th largest in the world – been subjected to such a comprehensive, coordinated sanctions effort. Russia's central bank found its assets frozen by the US, EU, UK, Switzerland, Japan, Canada, Australia, and South Korea. Ultimately, these major and minor reserve currency issuers froze approximately $300 billion of Russia's assets, roughly half of Russia's international reserves.

The ability of fiat reserve issuers to freeze transactions, which constitutes a form of *de facto* default on the underlying obligations, calls into question fiat reserve currencies' status as "safe haven" assets. Therefore, it is timely to explore the question of how, and to what extent, the risk of financial sanctions may motivate changes in central bank reserve composition.

I adopt a unique econometric approach to modeling cryptocurrency. Rather than assuming that the high realized returns of Bitcoin as of 2022 are likely to persist, I estimate a Bayesian model with an informative prior, whose parameters I choose to reduce the expected compound returns of Bitcoin. I use simulations from the model to compute optimal portfolios as a function of risk aversion. To my knowledge, this is the first paper to quantify the potential effect of sanctions risk on international reserve allocations. Unlike most of the sanctions literature, which estimates the effects of sanctions ex-post, this paper focuses on understanding the ex-ante effect of sanctions risk.

The paper is structured as follows. Section 2 describes financial sanctions and related literature. Section 3 discusses evidence that fear of sanctions may motivate central bank gold holdings. Section 4 outlines features of cryptocurrency and its resistance to sanctions. Section 5 details the time series model to simulate the returns of Bitcoin and reserve assets, and Section 6 benchmarks the performance of the model. Section 7 uses the simulations to demonstrate the effect of sanctions on reserve allocations under a range of plausible assumptions. Finally, Section 8 concludes, with implications for the renminbi and central bank digital currencies.

2 Overview of Financial Sanctions

The history of economic sanctions dates back to the blockades of World War I, following which the League of Nations began employing sanctions in support of foreign policy objectives as an alternative to war. Mulder (2022) describes the usage of economic sanctions as a coercive tool in the interwar period. Economic sanctions encompass both trade sanctions (tariffs and embargoes) as well as financial sanctions. In the digital commerce era, financial sanctions have assumed greater prominence because of the degree of centralization of the global financial system and the immediacy with which electronic banking services can be disabled. Zarate (2013) details the expansion of the US Treasury Department's financial sanctions programs to assist counterterrorism efforts following the 9/11 attacks. Hufbauer and Jung (2020) updates Hufbauer, Schott, et al. (2009) and describes more recent developments in economic sanctions, including the Iran nuclear agreement and Trump tariffs.

In the United States, financial sanctions can be implemented through a legislative or presidential procedure. In the legislative procedure, Congress passes a law specifying sanctions, and either the President signs the law or Congress overrides the President's veto. In the presidential procedure, the President issues an executive order declaring a state of emergency concerning a particular country, region, or topic, which empowers the US Treasury's Office of Foreign Assets Control (OFAC) to issue sanctions. All US persons must comply with OFAC sanctions, including all persons and entities within the United States, all US incorporated entities and their foreign branches. If a US person identifies property belonging to an OFAC-sanctioned entity, the US person must "block" (freeze) the property – prohibiting transfers or dealings of any kind with regard to the property – unless OFAC grants an exception or lifts the specific sanction.[1] Penalties for failing to comply with OFAC sanctions can be significant. Fines can reach millions of dollars, and individuals can face jail time. In April 2022, a researcher received a 63-month sentence, along with a $100,000 fine, for delivering a presentation about cryptocurrency technology in North Korea.[2] Foreign entities beyond the reach of US law enforcement may face "secondary sanctions" for conducting business with sanctioned entities.

Other governing bodies implement various procedures for issuing sanctions. The European Union Common Foreign and Security Policy Council may impose sanctions if all EU members consent to the proposal. The United Nations Security Council may approve sanctions if nine out of the fifteen members vote in favor, but any permanent member (China, Russia, France, the United Kingdom, and the United States) may veto the proposal. Perhaps because of the relative ease of implementing sanctions through unilateral executive action, the US has sanctioned far more entities than the UN or EU. As of September 2019, the OFAC list included 8,755 entities, compared with 2,136 for the EU and 1,057 for the UN.[3] Partly due to concerns about the overuse of sanctions and unintended effects on vulnerable groups, the Biden administration announced in October 2021 that it intended to limit its

[1] OFAC operates several types of sanctions programs. This paper studies the effect of full blocking sanctions under OFAC's Specially Designated Nationals And Blocked Persons List ("SDN List").

[2] https://www.wsj.com/articles/cryptocurrency-guru-sentenced-to-more-than-five-years-in-prison-over-north-korea-trip-11649789150

[3] https://www.tradefinanceglobal.com/wire/accuity-data-reveals-increased-complexity-of-sanctions-compliance-and-implications-for-global-trade/

usage of sanctions.[4]

Empirical evidence concerning the effectiveness of sanctions programs is mixed. Felbermayr et al. (2020) compile a global database and find that sanctions are increasingly used over time; the share of financial sanctions is rising; the main objectives of sanctions are increasingly related to democracy or human rights; and the success rate of sanctions has fallen since 1995, averaging 30% across policy objectives. In a firm-level comparison of sanctioned to unsanctioned Russian firms, Ahn and Ludema (2020) show that the 2014 sanctions caused significant losses in operating revenue, asset values, and employees, but the Russian government shielded some strategic firms from the full effect of the Western sanctions.

Several central banks currently face or have faced US sanctions. As of July 2022, the central banks of Russia, Iran, Syria, North Korea, and Venezuela are under US sanctions. Additionally, after the 2021 Taliban takeover, the Biden administration froze the New York Fed account belonging to the central bank of Afghanistan, ultimately expropriating the funds to divide them equally between a trust for the people of Afghanistan and victims of the 9/11 attacks.[5] Previously, the US froze the reserves of Iraq following its 1990 invasion of Kuwait (which President Bush subsequently expropriated in 2003[6]) and temporarily suspended Iraq's cash withdrawals in 2015 over concerns that cash was being transported to terrorist groups and sanctioned Iranian banks.[7] In 2008[8] and 2020,[9] the US threatened to freeze Iraq's reserves if Iraq expelled US troops from the country.

The countries described above face sanctions for a variety of reasons including launching external wars, sponsoring terrorism, developing nuclear weapons, repressing protests, refusing to accept the outcome of elections, and seizing power from a previous government.

[4] https://www.wsj.com/articles/biden-administration-to-trim-use-of-sanctions-in-a-foreign-policy-shift-11634600029

[5] https://www.nytimes.com/2022/02/11/us/politics/taliban-afghanistan-911-families-frozen-funds.html

[6] https://www.washingtonpost.com/archive/politics/2003/03/21/us-seizes-14-billion-in-frozen-iraqi-assets/98cbb395-ec84-422e-b825-7a864eea340d/

[7] https://www.wsj.com/articles/u-s-cut-cash-to-iraq-on-iran-isis-fears-1446526799

[8] https://www.independent.co.uk/news/world/middle-east/us-issues-threat-to-iraq-s-50bn-foreign-reserves-in-military-deal-841407.html

[9] https://www.wsj.com/articles/u-s-warns-iraq-it-risks-losing-access-to-key-bank-account-if-troops-told-to-leave-11578759629

Therefore, a central bank cannot preclude the possibility of facing US sanctions if the country in question simply avoids particular types of activity. Moreover, there is no expiration date to US financial sanctions. Some sanctions against Iran have been in place since 1979.

3 Sanctions and Central Bank Gold Allocations

Empirical evidence regarding the ex-ante effect of sanctions risk on central bank gold holdings is helpful to motivate the subsequent discussion of cryptocurrency as a potential reserve asset.

The primary non-fiat reserve asset is gold. Gold reserves under the physical control of a central bank are largely beyond the reach of financial sanctions by third parties. For example, despite facing US financial sanctions, the Central Bank of Venezuela chartered Russian aircraft to sell its gold reserves in Africa.[10] Therefore, a desire to hedge against financial sanctions risk by fiat reserve currency issuers is potentially one reason why central banks may accumulate gold reserves.

Since the Great Recession, central banks have steadily added gold to their reserves, as illustrated in Figure 1. In 2020, the gold share of international reserves reached a 20-year high of 14.4%.

The distribution of gold shares across central banks is heterogeneous. Figure 2 displays the distribution conditional on a non-zero gold share. About 20% of central banks do not own any gold; a few maintain gold shares in excess of 50%. The heterogeneity in gold allocations across central banks suggests that political and logistical considerations – such as the cost of transporting and securing physical gold – are as important as a central bank's risk tolerance in determining its portfolio composition. Indeed, Aizenman and Inoue (2012) find that central bank gold holdings are correlated with "global power" such as the history of being an empire, the geographic size of the country, and the country's centrality to the international financial system.

[10] https://www.wsj.com/articles/how-7-4-tons-of-venezuelas-gold-landed-in-africaand-vanished-11560867792

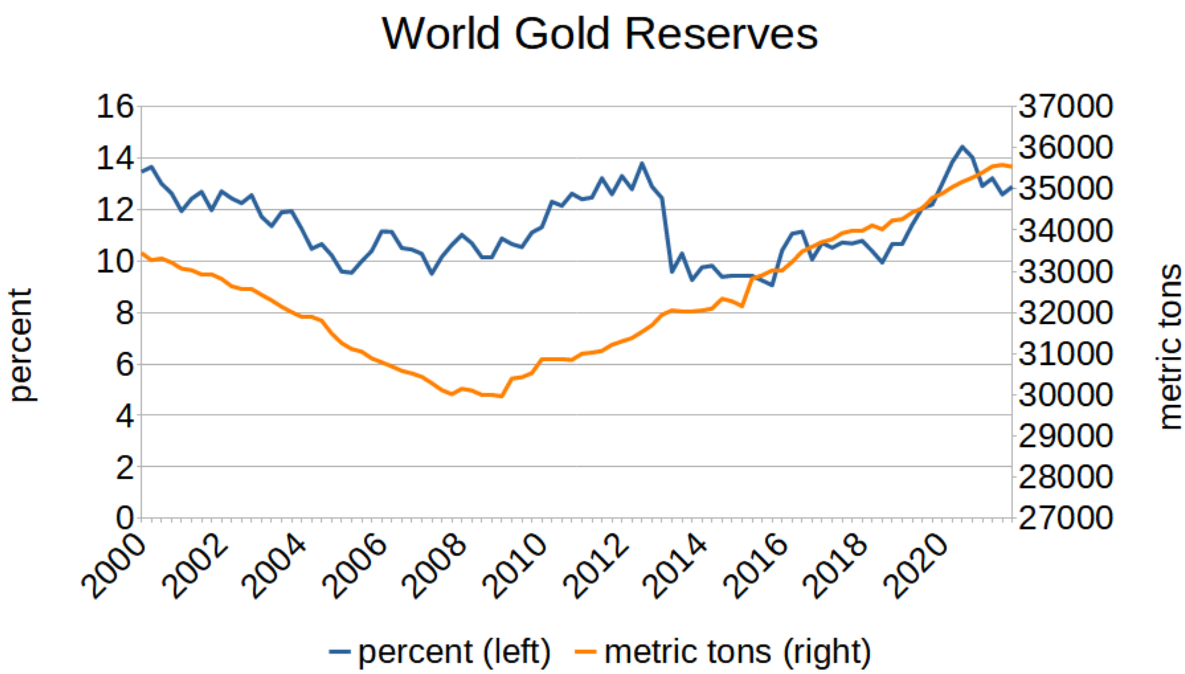

Source: World Gold Council

Figure 1: The quantity of gold in international reserves.

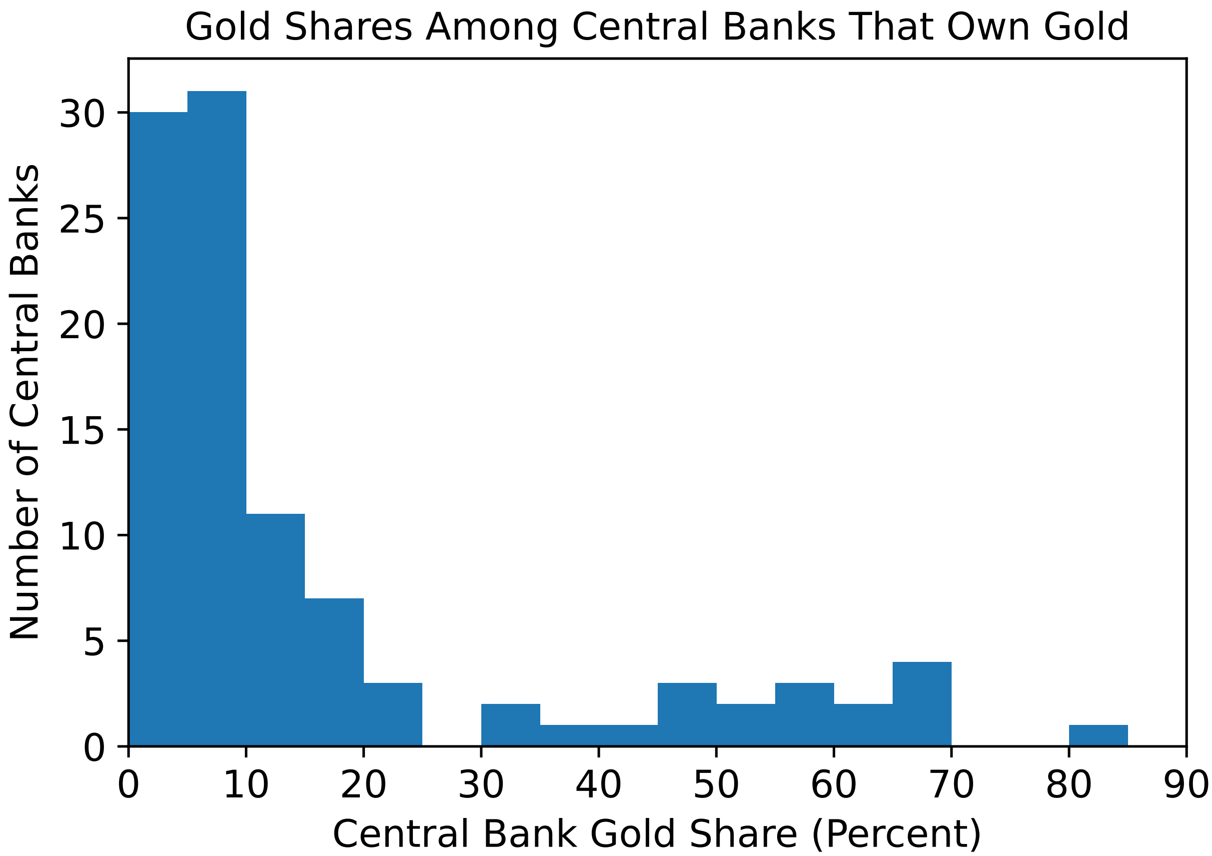

Source: World Gold Council

Figure 2: The distribution of gold shares across central banks that own gold.

It is also informative to study the changes in gold shares across central banks, and relate those changes to a variable that could proxy for financial sanctions risk. Military import deals evince political alignment and require a real commitment of financial and political capital.[11] I obtain annual data from 2017 to 2021 on military imports and exports from the Stockholm International Peace Research Institute (SIPRI). The financial terms of arms trade agreements are often undisclosed, so SIPRI computes a Trend Indicator Value (TIV) based on unit production costs and weapon characteristics that represents the military value of the items being traded.

Countries that import valuable military equipment from geopolitical rivals of the United States, particularly China and Russia, plausibly face a heightened risk of U.S. sanctions. In 2017, then-President Trump signed the Countering America's Adversaries Through Sanctions Act, providing for sanctions on entities that transact with the Russian defense sector. In 2020, President Trump issued Executive Order 13959, establishing financial sanctions against certain Chinese military companies. As previously discussed, entities that transact with sanctioned entities face the risk of secondary sanctions, so importing military goods from China or Russia raises the importer's sanctions risk.

I obtain a sample of central bank gold shares from the World Gold Council. I filter the sample by discarding countries that do not own any gold, or whose TIV military imports from the US, China, and Russia were zero from 2017 to 2021 inclusive, resulting in 81 countries. For each country, I define a measure of political alignment, mil_import_diff, in Equation (1). The result is a metric ranging from -100 to 100 that indicates the extent of a country's military-political alignment with the US (100) or rivals of the US (-100). Averaging over a 5-year period in Equation (2) helps account for the fact that some military trade deals require multiple years to plan and execute.

[11] I eschew using United Nations voting records, which are commonly used to measure political alignment among countries, because they are not obviously connected to sanctions risk and often constitute cheap talk.

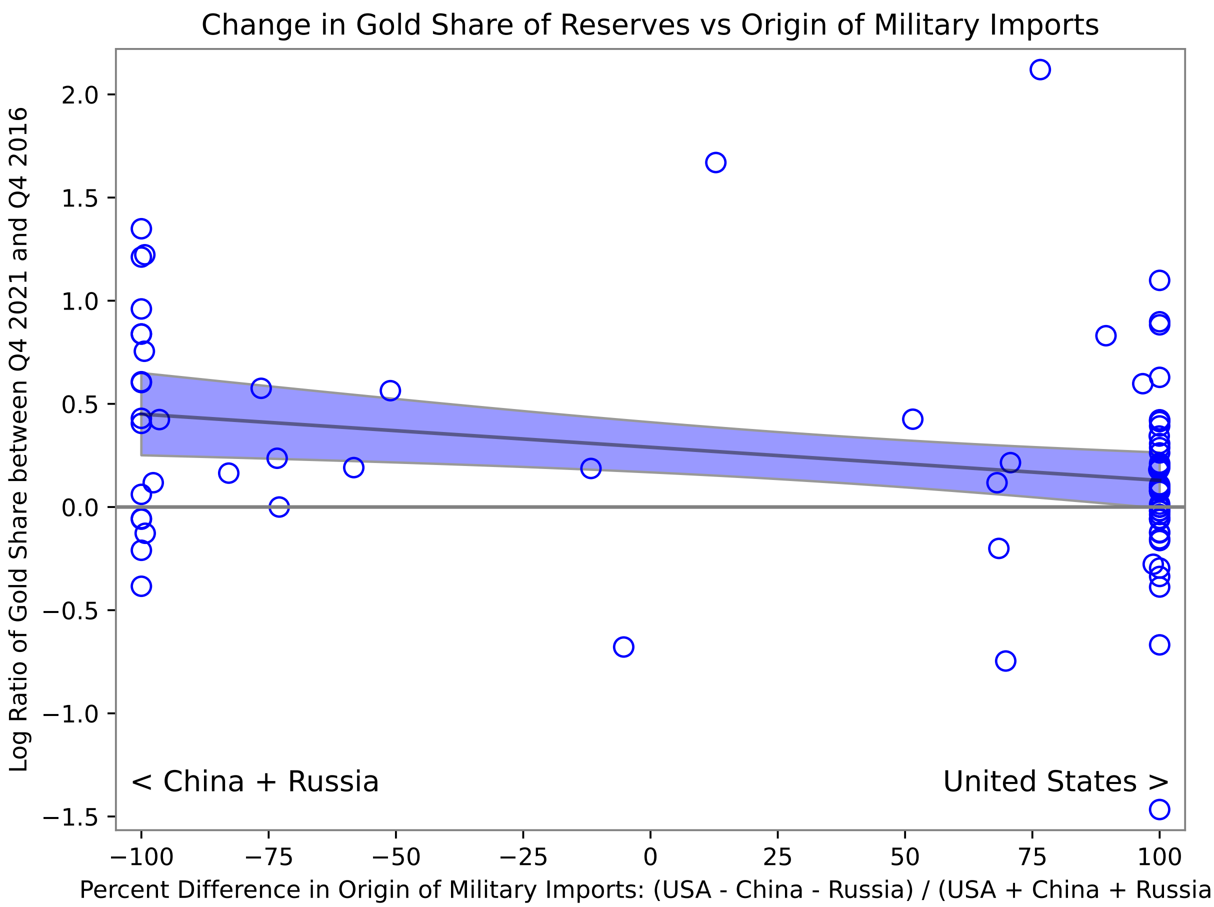

Source: SIPRI, World Gold Council, author's calculations

Figure 3: The origin of military imports explains changes in central bank gold shares from Q4 2016 to Q4 2021.

$$\text{mil\_import\_diff}_{i,t} = \frac{100(\text{TIV}_{\text{USA}\to\text{i;t}} - \text{TIV}_{\text{China + Russia}\to\text{i;t}})}{\text{TIV}_{\text{USA + China + Russia}\to\text{i;t}}} \tag{1}$$

$$\text{mil\_import\_diff}_{i} = \sum_{t=\text{Q4 2016}}^{T=\text{Q4 2021}} \text{mil\_import\_diff}_{i,t}/5 \tag{2}$$

Then, I run the following regression. I choose the log ratio of the gold share as the outcome variable, rather than the percentage point difference, because of the heterogeneity in the level of gold shares illustrated in Figure 2.

$$ln(\frac{\text{gold\_share}_{i,\text{Q4 2021}}}{\text{gold\_share}_{i,\text{Q4 2016}}}) = \beta_0 + \beta_1(\text{mil\_import\_diff}_i) + \gamma(\text{controls}_i) \tag{3}$$

Without including any controls, Figure 3 displays a scatterplot illustrating the regression line and its confidence interval.

More detailed regression results are available in Table 1. I compute standard errors using the heteroskedasticity robust Eicker-Huber-White estimator. Control variables include groupings based on geographic location and GDP per capita. Regardless of the choice of controls, the origin of a country's military imports retains statistically significant explanatory power over changes in its central bank gold shares from Q4 2016 to Q4 2021.

Table 1: Gold Share Regression Results

Economic Controls	No	Yes	No	Yes
Geographic Controls	No	No	Yes	Yes
R^2	0.07	0.13	0.22	0.26
mil_import_diff Coefficient (p-value)	-0.0016 (0.009)**	-0.0019 (0.008)**	-0.0021 (0.004)**	-0.0024 (0.006)**

Of course, these regressions do not establish a causal link between sanctions risk and gold allocations. But this analysis of central bank gold shares does demonstrate robust demand for a non-fiat reserve asset, especially among countries that may be less trusting of the major fiat reserve issuers. Cryptocurrencies may be of particular interest for countries with political or economic disagreements with fiat reserve issuers.

4 Characteristics of Cryptocurrency

Cryptocurrencies are fungible digital tokens, the history of which is stored on a digital ledger secured by cryptography. The largest and oldest cryptocurrency, Bitcoin, began use in 2009; since then, people have created thousands of different cryptocurrencies employing different digital architectures. As of July 2022, Bitcoin and Ether (the second largest cryptocurrency) collectively comprise about 60% of the approximately $1.0 trillion market capitalization of all cryptocurrencies. Since 2018, the prices of Bitcoin and Ether have been highly correlated, as illustrated in Figure 4. A detailed history and technical description of cryptocurrency is beyond the scope of this paper, but Hardle, Harvey, and Reule (2020) and Halaburda et al. (2022) provide an overview.

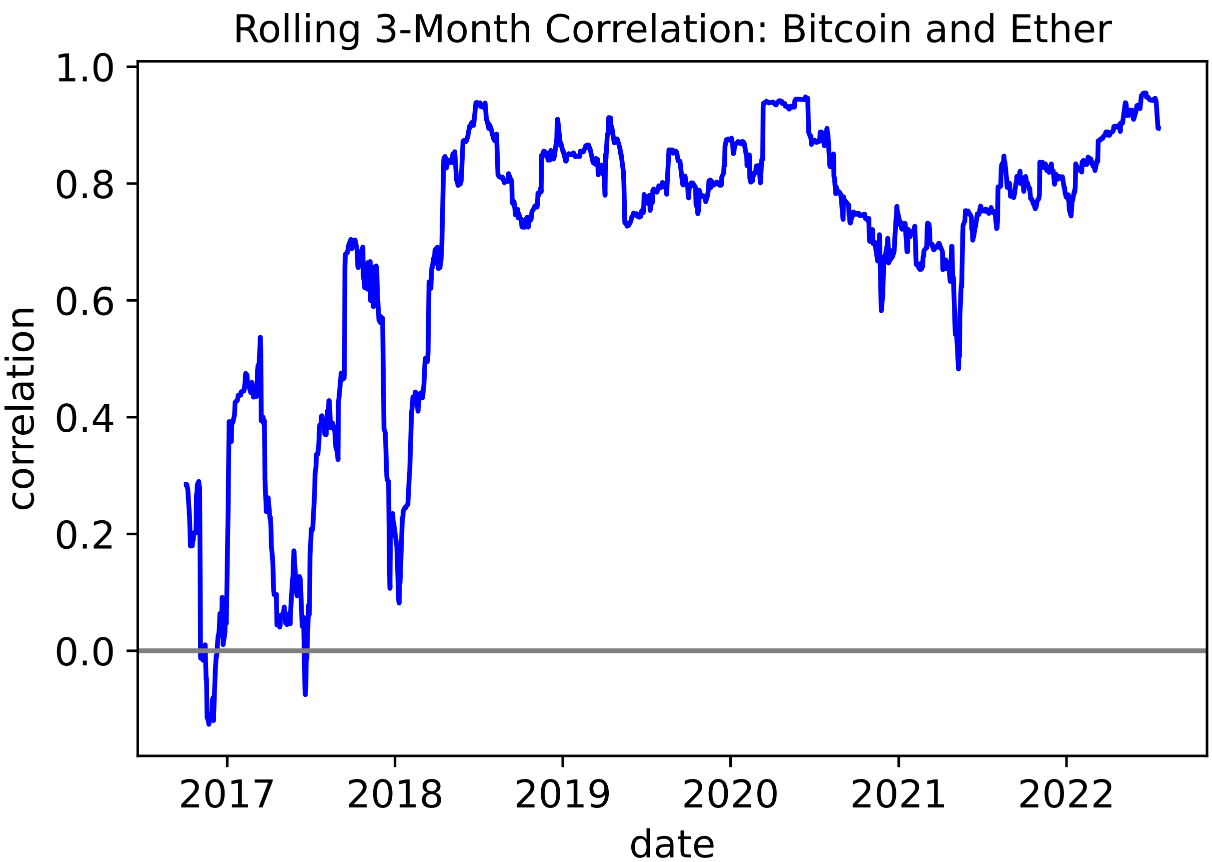

Source: investing.com

Figure 4: The 3-month rolling correlation between Bitcoin and Ether.

A central tenet of Bitcoin and Ethereum, the network that supports the cryptocurrency Ether, is the immutable public decentralized system, called a blockchain, through which tokens are created and transactions are authenticated. The blockchain is pseudo-anonymous in that transactions between wallets are public knowledge,[12] while the identities of the wallets' owners are generally not revealed (but sometimes can be inferred based on transaction patterns and other external information). Bitcoin currently operates a "proof-of-work" process for authenticating transactions, in which groups of miners ("mining pools") operating specialized computing equipment compete to confirm new blocks of transactions. The odds of winning the competition are proportional to the computing power expended on the task, and the winner receives both transaction fees and a quantity of newly minted cryptocurrency. Ethereum operates a "proof-of-stake" system in which holders of Ether receive transaction fees and newly minted Ether proportional to the quantity of Ether that they pledge to authenticate new transactions.

Although proponents of cryptocurrency cite its "trustlessness" as an advantage over fiat currencies, Bratspies (2018) points out that cryptocurrencies require different forms of trust. Specifically, users must trust that the cryptocurrency software itself is secure, that miners will not collude to attack the integrity of the blockchain, and that the governance process will not approve of a "hard fork" that fundamentally alters the blockchain itself or other parameters of the cryptocurrency.

Decentralized cryptocurrencies are resistant to governmental financial sanctions. A fiat currency issuer can issue sanctions against particular cryptocurrency wallets, rendering it illegal for holders of fiat currency to assist the owners of the sanctioned cryptocurrency wallets with converting their cryptocurrency into fiat currency. Sanctioned individuals may not be able to use large cryptocurrency exchanges, who are required to comply with sanctions programs if the exchanges want to continue converting cryptocurrency into fiat currency. But as long as the issuers of fiat currency do not control the blockchain itself, sanctioned individuals can continue to send cryptocurrency from one wallet to another. Additionally, sanctioned individuals could participate in "off-chain" transactions by providing the private keys to their cryptocurrency wallet in exchange for goods, services, or other forms of currency,

[12] Some cryptocurrencies, such as Monero, employ additional measures that obscure transactions.

as discussed in Luckner, Reinhart, and Rogoff (2021). For example, in April 2022, North Korean hackers continued to launder $600 million of Ether stolen during a hack of the video game Axie Infinity, eight days after the US Treasury sanctioned the digital wallet used in the hack.[13] In 2018, Iran began issuing licenses to cryptocurrency miners, who are required to sell their tokens to the central bank to facilitate sanctions evasion.[14] A discussion regarding stablecoins, which are inherently more centralized than Bitcoin or Ether and therefore not suitable for evading sanctions, can be found in Appendix A.

Under a proof-of-work system, the ability to censor transactions on the blockchain requires achieving "majority hash power," meaning that the censor must control at least 51% of the computing power employed by all miners. Achieving such a status is not feasible due to the sheer quantity of computing power dedicated to Bitcoin mining, as well as the amount of electricity required to power the mining chips. Furthermore, the structure of the Bitcoin network incentivizes Bitcoin owners to oppose any individual's acquisition of majority hash power by purchasing or producing their own mining chips, because majority hash power also enables a "double spending attack" that results in the duplication of Bitcoin, which would likely destroy confidence in the cryptocurrency. In 2014, the Bitcoin mining pool Ghash.io briefly acquired majority hash power,[15] and faced a combination of public criticism, cyber-attacks, and abandonment by miners that quickly reduced its market share below 50%. No mining pool has ever acquired majority hash power in Bitcoin since then. The impossibility of implementing a successful attack against Bitcoin contributes to Bitcoin's status as the most valuable cryptocurrency. Moreover, the proof-of-work mechanism establishes a barrier to the large-scale adoption of an alternative to Bitcoin. Rival proof-of-work currencies tend to face difficulty attracting miners, since prospective miners could mine Bitcoin and other established currencies more profitably. This creates a feedback cycle wherein small proof-of-work currencies are not valuable enough to attract many miners, and therefore remain

[13] https://www.washingtonpost.com/business/2022/04/23/north-korea-hack-crypto-access/

[14] https://techcrunch.com/2022/06/19/iran-to-cut-electricity-to-authorized-crypto-miners-report/

[15] https://www.extremetech.com/extreme/184427-one-bitcoin-group-now-controls-51-of-total-mining-power-threatening-entire-currencys-safety

vulnerable to majority hash attacks, a risk which prevents the small proof-of-work currencies from becoming more valuable. Indeed, many smaller cryptocurrencies that operate based on the proof-of-work mechanism have faced majority hash attacks, as described in Shanaev et al. (2020).

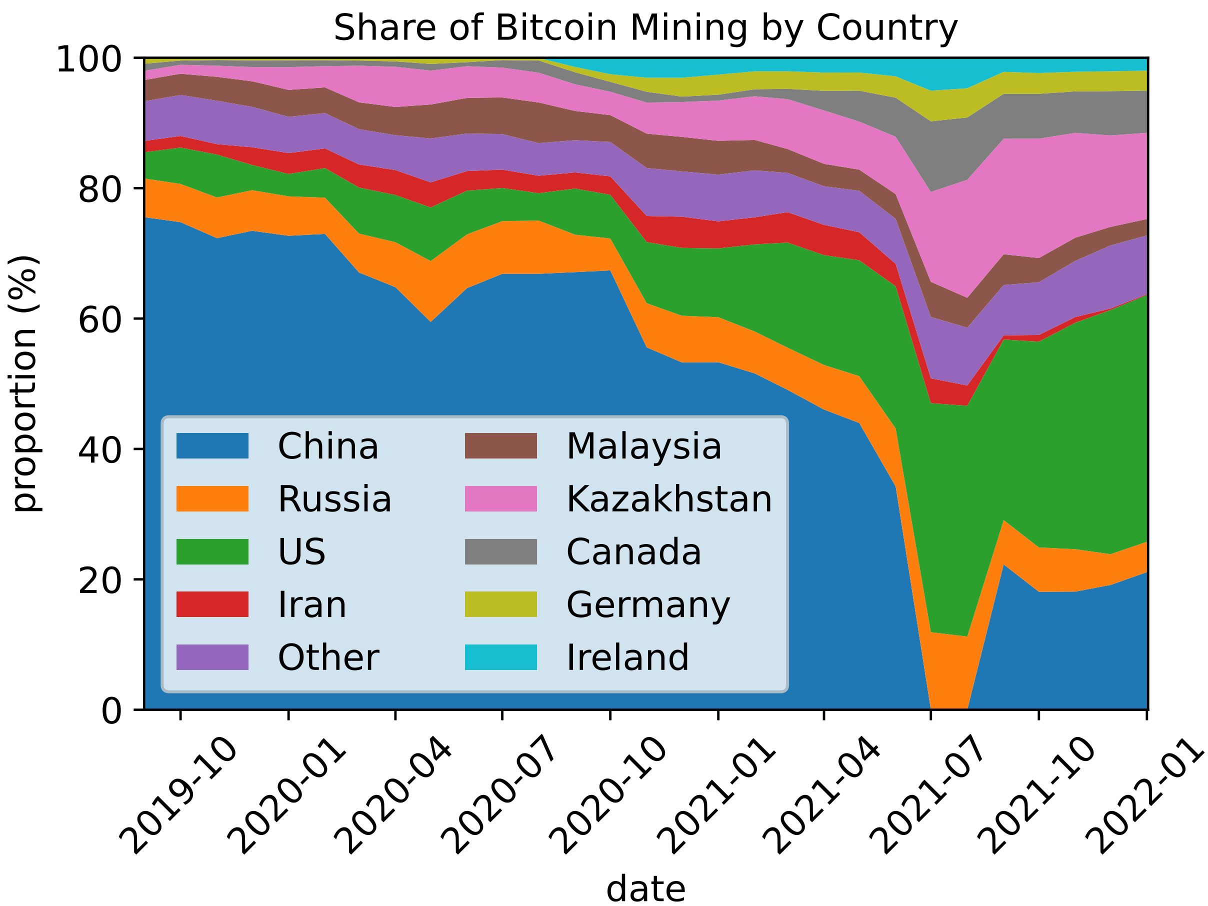

Source: Cambridge Centre for Alternative Finance

Figure 5: The proportion of Bitcoin mining, by country. China officially banned Bitcoin mining in June 2021, but mining operations continued, often via VPN's that mask the miners' locations. The shares for Germany and Ireland are likely inflated due to VPN services that route traffic through those countries.

Bitcoin miners generally do not comply with sanctions regarding the wallets whose transactions the miners are validating. One Bitcoin mining pool, Marathon, announced in May 2021 that it would not validate transactions involving wallets that appeared on the OFAC sanctions list.[16] Facing criticism, Marathon reversed itself one month later, noting that its mining income was far lower than that of its peers. Theoretically, a sanctioned individual could offer higher transaction fees if necessary to motivate miners to process the individual's transactions. The distribution of Bitcoin mining is well diversified and mobile across countries, as illustrated in Figure 5, complicating efforts by any individual country to censor the blockchain by regulating Bitcoin mining.

Under Ethereum's proof-of-stake system, the ability to censor transactions on the blockchain requires holding a majority of staked cryptocurrency. A sufficient condition to enforce censorship on the network is the acquisition of a majority of all digital tokens in circulation. Half of Ether's market capitalization is approximately $92 billion, as of July 2022. Of course, the price of Ether would rise if an individual began purchasing large quantities of Ether, so an enormous expenditure of resources would be required to enforce censorship on the Ethereum network. In order to defeat the censor, users would need to acquire some of the censor's digital tokens, or else implement a "hard fork" by migrating towards a new copy of the blockchain with a different distribution of tokens across users, effectively abandoning the original Ethereum as a dead project. Although network effects will likely support the value of Ether in the short run, staking Ether does not require a meaningful real-world expenditure of resources, so Ethereum may be more likely to face serious competition from alternative networks in the long run than Bitcoin.

The amount of electricity consumed by Bitcoin mining results in a significant negative environmental externality. According to the International Energy Agency and the Cambridge Centre for Alternative Finance, Bitcoin mining alone consumes approximately 0.5% of world energy production as of July 2022. A central bank that purchases significant quantities of Bitcoin will promote additional Bitcoin mining by increasing the price of Bitcoin, resulting in additional environmental harm. The environmental externalities of Bitcoin can be thought of as the cost of decentralization. Countries that are worried about the possibility of US or

[16]https://www.theblock.co/linked/106865/marathon-ofac-bitcoin-mining-pool-taproot

EU sanctions will likely find the environmental costs of Bitcoin mining to be an acceptable tradeoff in return for the benefit of hedging their reserves against sanctions risk.

Appendix B discusses Bitcoin's liquidity and role as a store of value, noting a "flight to safety" effect that appeared in February 2022 immediately following the global sanctions against the Central Bank of Russia. Reaching the opposite conclusion as Smales (2019), I argue that Bitcoin meets the minimum requirements to be considered a store of value.

5 Reserve Assets Model

In order to solve portfolio optimization problems including Bitcoin and other reserve assets, I require a means of generating samples from a plausible joint distribution of future returns of those assets. A major challenge is that the historical returns of Bitcoin are likely to severely overstate Bitcoin's forward-looking expected returns. Between July 1, 2012 and July 1, 2022 the compound annual return of Bitcoin exceeded 100% per year. It is simply not realistic to expect such high returns to continue indefinitely, because they are not based upon the economic fundamentals of cryptocurrency as it exists today.[17] Bitcoin's high returns were realized by early adopters, who made a highly risky investment (consisting of computing equipment and electricity) into a brand new digital payment system whose usefulness and longevity were both unclear.[18] Today, cryptocurrencies are far more mainstream; there is no reason to expect Bitcoin's future returns to beat optimistic forecasts of the stock market by an entire order of magnitude.[19]

Any statistical technique based on bootstrapping, maximum likelihood, or moment-matching will produce draws from a distribution of Bitcoin returns that resemble Bit-

[17] The concept that expected returns fall as asset prices rise is well documented in the context of the stock market. For example, after a stock joins the list of the top 10 largest US stocks, its 10-year expected return is 1.5% below the market return. See: https://www.dimensional.com/us-en/insights/large-and-in-charge-giant-firms-atop-market-is-nothing-new

[18] To emphasize the riskiness of an early-stage investment into Bitcoin, consider the thousands of cryptocurrencies that are nearly worthless today.

[19] A combination of high realized returns and declining expected returns could be explained by a declining rare disaster risk associated with Bitcoin, due to Bitcoin's increasing rate of adoption and increasing hashrate, both of which reduce the likelihood of a sudden collapse or a successful attack.

coin's historical returns, and likely overstate Bitcoin's future returns. Therefore, I employ a Bayesian approach, in which I use an informative prior to adjust Bitcoin's expected returns to a more reasonable level. I do not claim to have particular knowledge about the expected returns of any asset, including Bitcoin. I provide a framework for portfolio optimization in the context of sanctions where the end-user can explore optimal allocations by encoding his or her own beliefs. Such beliefs may be derived from principles of economics or other background knowledge available to the investor.

Although DCC-MGARCH models are almost always solved via the maximum likelihood technique, a small literature takes a Bayesian approach similar to this paper. Fioruci, Ehlers, and Louzada (2014) develop an R software package to estimate Bayesian DCC-MGARCH models using a handful of skewed and heavy-tailed error distributions. Shiferaw (2019) uses a Bayesian DCC-MGARCH model to study the correlation between energy and agricultural commodity prices. Tang and Aruga (2022) use a Bayesian DCC-MGARCH model to investigate relationships among the fossil fuel, clean energy stock, gold, and Bitcoin markets.

5.1 Time Series Model

This model combines the AGARCH specification of Engle and Ng (1993) with the DCC framework of Engle (2002), thereby capturing several important features of financial time series. The model specification is outlined below.

5.1.1 Univariate Equations

In equation (4), I model the log returns of each asset. I assume the mean return is not time varying; accordingly, I do not capture low-frequency features of financial data such as long-horizon mean reversion. Such effects are difficult to discern from noise at the daily frequency.

$$r_{i,t} = ln(p_{i,t}/p_{i,t-1}) = \mu_i + \epsilon_{i,t}, \quad i = 1, 2, \ldots N. \text{ (returns of } i\text{th asset)} \tag{4}$$

In equation (5), I model the volatility of returns as having a conditional Student-t distribution, allowing for heavy-tailed returns. When estimating the model, I further constrain $\nu_i > 3$ so that the first three moments of the error distribution are finite. While finite moments are not a prerequisite for Bayesian inference, they tend to produce more reasonable volatility estimates.

$$\epsilon_{i,t}|\vec{r_1}, \vec{r_2}, \ldots \vec{r_{t-1}} \sim \text{Student-t}\ (0, \sqrt{\sigma_{i,t}^2(\nu_i - 2)/\nu_i}, \nu_i)\ \text{ (innovations for } i\text{th asset)} \tag{5}$$

In equation (6), I employ a standard GARCH(1,1) specification for modeling the time-varying volatility. The GARCH component produces heavy-tailed returns and volatility clustering. The AGARCH δ_i term allows for asymmetric news impact effects: negative shocks increase volatility moreso than positive shocks. Several different potential specifications capture this effect (EGARCH, GJR-GARCH, APARCH) but only the AGARCH specification is differentiable over its entire domain, which is helpful for the Bayesian sampling procedure.

$$\sigma_{i,t}^2 = \omega_i + \alpha_i(\epsilon_{i,t-1} - \delta_i)^2 + \beta_i \sigma_{i,t-1}^2\ \text{ (variance of } i\text{th asset)} \tag{6}$$

5.1.2 Multivariate Equations

In equations (7), (8), and (9), I compute the standardized residuals, the dynamic conditional covariance, and the dynamic conditional correlation respectively. The DCC component of the model allows the correlations to vary over time in a mean-reverting fashion, with clustered periods of high and low correlation.

$$\vec{\eta_t} = diag(\vec{\sigma_t})^{-1}(\vec{r_t} - \vec{\mu})\ \text{ (standardized residuals)} \tag{7}$$

$$\boldsymbol{Q}_t = \boldsymbol{S} + a(\vec{\eta_{t-1}}\vec{\eta_{t-1}}^T - \boldsymbol{S}) + b(\boldsymbol{Q}_{t-1} - \boldsymbol{S})\ \text{ (covariance)} \tag{8}$$

$$\boldsymbol{R}_t = diag(\boldsymbol{Q}_t)^{-0.5}\boldsymbol{Q}_t diag(\boldsymbol{Q}_t)^{-0.5} \quad \text{(dynamic correlation)} \tag{9}$$

Rather than modeling the GARCH error term using a multivariate distribution, I use a copula to correlate the univariate marginal distributions of each asset. The theory of copulas was originally developed by Sklar (1959) and Sklar (1973), who proved that any continuous joint distribution can be modeled by linking its marginal distributions with a unique copula. Copula models are commonly used for risk management in finance; Patton (2012) provides a recent overview. The copula approach provides flexibility to model the marginal distributions and the manner of their correlation separately. There are many copulas that produce a wide variety of dependence structures. Two popular copulas are the Gaussian copula, which does not feature tail dependence, and the Student-t copula, in which the joint probability of extreme values is higher than the Gaussian copula. Nguyen et al. (2020) find that the Student-t copula provides the best fit when applying a GJR-GARCH model to a set of reserve currencies and gold. Accordingly, I use the Student-t copula, whose density is given in Equation (10).

$$\vec{\epsilon_t} \sim \frac{t_{\boldsymbol{R_t},\lambda}\left(T_\lambda^{-1}(F(\epsilon_{1,t})), T_\lambda^{-1}(F(\epsilon_{2,t})), \ldots, T_\lambda^{-1}(F(\epsilon_{n,t})\right)}{\prod_{i=1}^{N} t_\lambda\left(T_\lambda^{-1}(F(\epsilon_{i,t}))\right)} \quad \text{(Student-t copula density)} \tag{10}$$

Specifically, $t_{\boldsymbol{R_t},\lambda}$ refers to the multivariate Student-t density with correlation matrix[20] $\boldsymbol{R_t}$ and degrees of freedom λ, t_λ refers to the univariate Student-t density, T_λ^{-1} refers to the inverse Student-t distribution function, and $F(\cdot)$ refers to the marginal cumulative distribution for each asset (given in Equation (5)).

When estimating this model, I include five assets: Bitcoin, gold, 2-year US Treasury bonds, 2-year Euro bonds (measured in USD), and a global market capitalization-weighted stock index. I do not separately model Ether because of the high correlation between Bitcoin and Ether. However, the subsequent discussion regarding Bitcoin could be extrapolated to

[20]Formally, the correlation matrix is given by $\frac{\lambda}{\lambda-2}\boldsymbol{R_t}$, but as described in Demarta and McNeil (2005), the Student-t copula is invariant to a strictly increasing transformation of its components. So the Student-t dispersion matrix may be interpreted as a correlation matrix, provided that $\lambda > 2$.

a market capitalization-weighted portfolio of both Bitcoin, Ether, and other smaller non-stablecoin cryptocurrencies.

In adopting a dynamic Bayesian copula approach, this paper is most closely related to So and Yeung (2014), who also adopt a dynamic conditional correlation approach for a variety of copula functions, using a more granular vine copula structure that estimates the parameters of a separate copula function for each pair of assets. However, So and Yeung (2014) estimate their models using maximum likelihood, not a Bayesian approach. In this setting, vine copulas would introduce significantly more computational complexity without much additional benefit compared to a multivariate copula, because the estimated correlations across asset classes are moderate-to-low and the estimated tail-dependence coefficients[21] are all well below 0.1. Bayesian approaches that utilize vine copulas are more appropriate when modeling a collection of similar assets, and often involve approximations to the posterior distribution, as in Kreuzer and Czado (2019).

Because this model features a small number of assets, I avoid the curse of dimensionality, described by Pakel et al. (2021), that would hinder a Bayesian approach to modeling a large set of assets. Specifically, the number of parameters in this model grows proportional to the square of the number of assets. Although central banks have been diversifying their reserves away from US dollars and euro in the last two decades, as documented in Arslanalp, Eichengreen, and Simpson-Bell (2022), high-quality US dollar bonds, euro bonds, public equities, and gold still comprise over 80% of global reserves. Moreover, one could approximate the performance of many other assets, such as corporate bonds, by forming linear combinations of Treasuries and stock. Therefore, this model captures the set of investment options available to central banks.

[21] A measure of the probability that the return of one asset exceeds a particular quantile, conditional on another asset exceeding the same quantile, as the quantile approaches 1 or 0. Tail dependence coefficients vary between 0 and 1; for a Gaussian copula, tail dependence is equal to 0.

5.2 Priors

I use uninformative priors for the volatility and correlation processes, operating upon the assumption that historical volatility relationships among these assets are representative of the future.

However, I set informative priors for the mean returns of the assets. I considered several factors in selecting the Bitcoin prior. As described in Section 4, Bitcoin is a real asset in finite supply, so its long-run return should be related to the rate of inflation. But unlike gold, the effective supply of Bitcoin can only shrink in the long run, since Bitcoin can be rendered permanently inaccessible when owners lose, forget, or discard their private keys. The blockchain analysis firm Chainalysis found that about 20% of Bitcoin had not been touched in over 5 years, as of June 2020.[22] In the long run, perhaps 1% of Bitcoin will be ”lost” each year, boosting Bitcoin's expected return to 1% above inflation. From another perspective, Makarov and Schoar (2020) show that daily Bitcoin exchange volume explains up to 85% of variation in the Bitcoin price, suggesting that the long-run return of Bitcoin may be tied to the gross rate of wealth creation. Over the long run, wealth tends to remain a constant share of GDP, although wealth grew slightly faster than GDP in the 21st century as of 2020, according to Woetzel et al. (2021). If long-run real GDP growth averages 2%, then the long-run return of Bitcoin could be similar. Averaging these two approaches, I set the mean return prior for Bitcoin at 1.5% above that of gold (i.e., 1.5% above 2-year expected inflation).[23]

I set the expected returns of Treasuries and Euro bonds equal to their current yield as of September 16, 2022 (assuming uncovered interest parity holds in expectation for the Euro bonds), and I set the expected return of gold equal to the 2-year expected inflation rate. Based on Asness (2021), who finds that the expected return of both US and developed

[22] https://blog.chainalysis.com/reports/bitcoin-market-data-exchanges-trading/

[23] Setting the expected return of Bitcoin above that of gold suggests that Bitcoin commands a risk premium over gold, which seems reasonable since Bitcoin is far more volatile than gold, and Bitcoin's volatility is not clearly diversifiable. Indeed, Liu and Tsyvinski (2018) find that the performance of cryptocurrency is not explained by common macroeconomic risk factors.

international equities is about 5% above cash after adjusting for valuation changes, I set the expected return of world stock equal to 5% above that of Treasuries. Table 2 contains a list of priors, and an explanation of each.

Table 2: Priors on Model Parameters

Parameter	Prior	Notes
$\vec{\mu}$	Normal(m, σ), where $m = (4.68, 3.18, 3.85, 3.85, 8.85)'/253$ $\sigma = (1.0, 0.5, 0.02, 0.02, 0.50)'/253$	See above. Annual returns are divided by 253 trading days in the year.
$\vec{\omega}$	Cauchy(0, 0.1) constrained within the interval [0,1]	Gelman (2006) suggests the half-Cauchy distribution as a prior for variance parameters.
$\vec{\alpha}$	Uniform(0, 1)	Constraints that the variance process is mean-reverting.
$\vec{\beta}$	Uniform(0, 1-$\vec{\alpha}$)	Constraints ensure that the variance process is mean-reverting.
$\vec{\delta}$	Cauchy(0, 0.1), constrained within the interval [0,1]	Ensures that negative shocks have a stronger impact on the variance at time t than positive shocks.
$\boldsymbol{S}$	Lewandowski-Kurowicka-Joe(1)	Uniform prior over the set of correlation matrices.
a	Uniform(0,1)	Constraints ensure that the correlation process is mean-reverting.
b	Uniform(0,1-a)	Constraints ensure that the correlation process is mean-reverting.
$\vec{\nu}$	$\vec{\nu} - 3 \sim$ Gamma(2,0.1)	Juarez and Steel (2010) suggest this Gamma prior for Student-t degrees of freedom. $\vec{\nu} > 3$ ensures that the first three moments of the marginal distributions are finite.
λ	$\lambda - 2 \sim$ Gamma(2,0.1)	Juarez and Steel (2010) suggest this Gamma prior for Student-t degrees of freedom. $\lambda > 2$ ensures that the covariance of the multivariate copula is finite.

5.3 Data

I obtain daily historical Bitcoin and gold prices at the close of each business day from investing.com, 2-year US Treasury yields from the St. Louis Federal Reserve, and 2-year Euro bond yields from the European Central Bank. For world stock, I obtain the daily returns of the Vanguard Total World Stock Index ETF, adjusted for dividends, which tracks the FTSE Global All Cap Index.

Bitcoin's early history–when it was regarded moreso as a science experiment rather than a legitimate asset class–is likely to be less informative of Bitcoin's future returns compared to Bitcoin's modern history. Additionally, the correlation between Bitcoin and the stock market sharply rose starting in March 2020, as illustrated in Figure 6. Because I want to capture this higher correlation, I begin my sample on March 1, 2020, continuing until July 22, 2022.

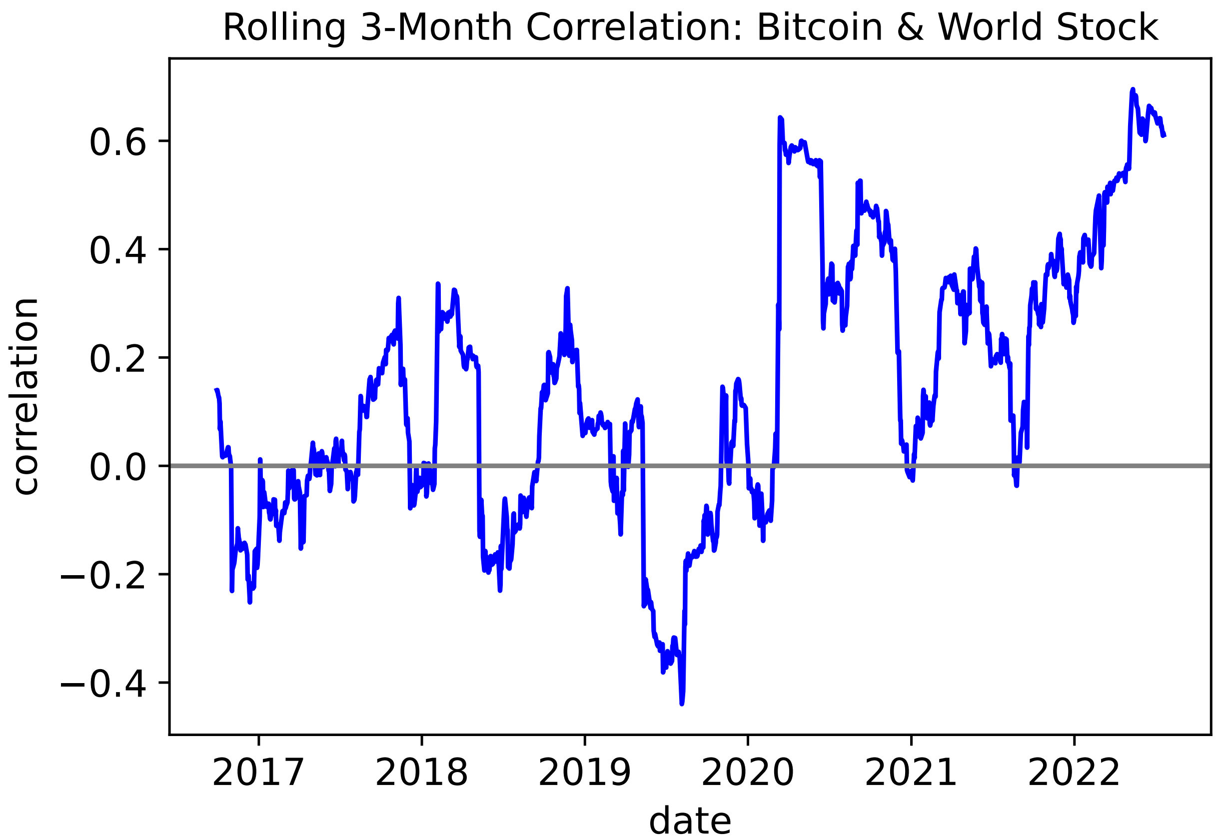

Figure 6: The estimated rolling correlation between Bitcoin and the FTSE Global All Cap Stock Index.

5.4 Computation

I sample from the posterior distribution over the parameter set $\{\vec{\mu}, \vec{\omega}, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \boldsymbol{S}, \vec{a}, \vec{b}, \vec{\nu}, \lambda\}$ using the Bayesian modeling software Stan developed by Carpenter et al. (2017). Stan implements a No-U-Turn sampler, a variant of Hamiltonian Monte Carlo. When running the sampler, I use 6 chains, with 1084 samples per chain, including 250 warmup iterations per chain that I discard. Figure 7 presents heatmaps of the mean return parameters for Bitcoin and gold.

In order to speed up the computation of this model, I implement an approximation to the inverse Student-t distribution, which must be computed many times as part of the multivariate Student-t copula in Equation (10). Computing the inverse Student-t distribution using root-finding procedures, such as using the inverse regularized beta function, is about 2x - 3x slower. The approximation, which combines two power series, achieves a maximum error of 0.15% for all degrees of freedom and all quantiles between 0.00001 and 0.99999. A description of the approximation can be found in Appendix C.

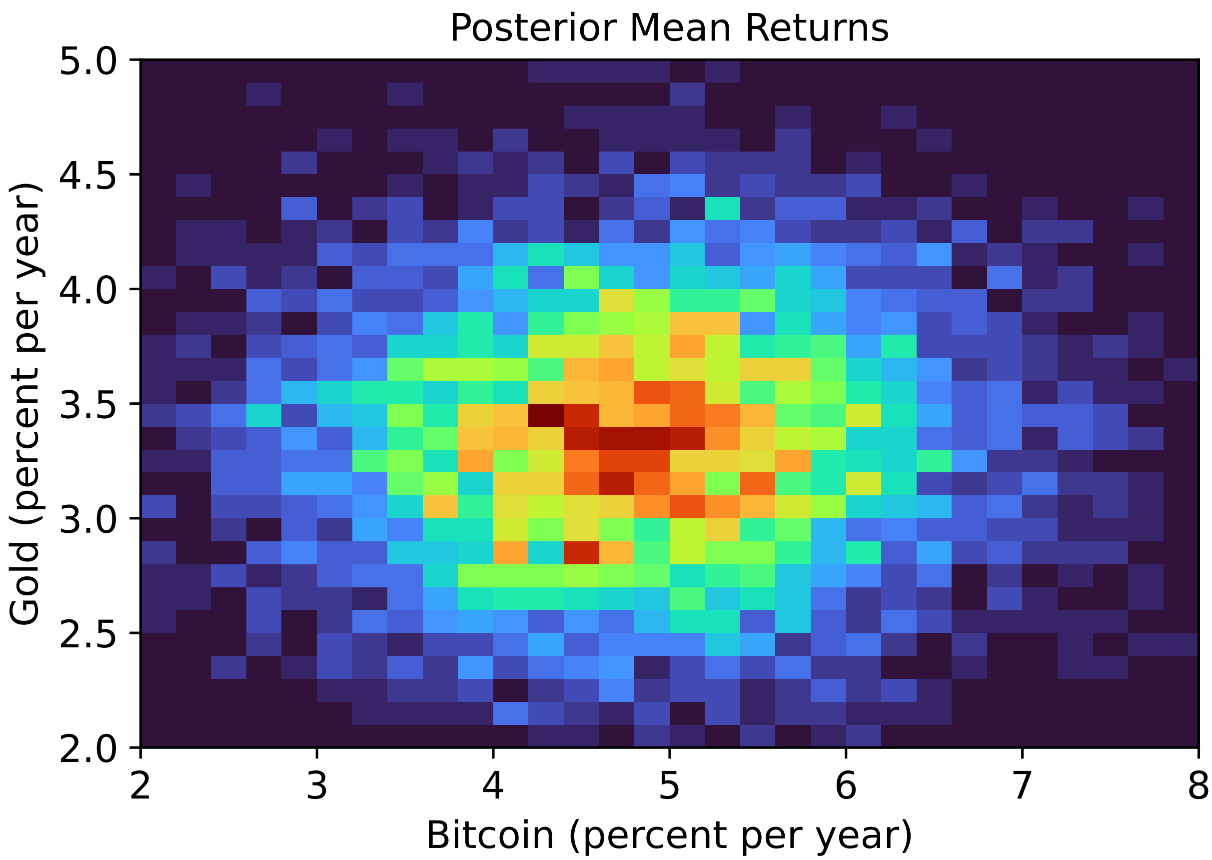

Figure 7: A heatmap of the posterior mean return parameters for Bitcoin and gold.

The unconditional standard deviation of asset i is given by:

$$\text{Std}(r_i) = \sqrt{\frac{\omega_i + \alpha_i \gamma_i^2}{1 - \alpha_i - \beta_i}} \text{ (unconditional standard deviation)} \tag{11}$$

For each draw from the model posterior, I compute the unconditional standard deviation, and compare the resulting distribution against the sample standard deviation. Figure 8 displays the result for Bitcoin and gold. The posterior unconditional standard deviation is heavy-tailed, especially for Bitcoin.

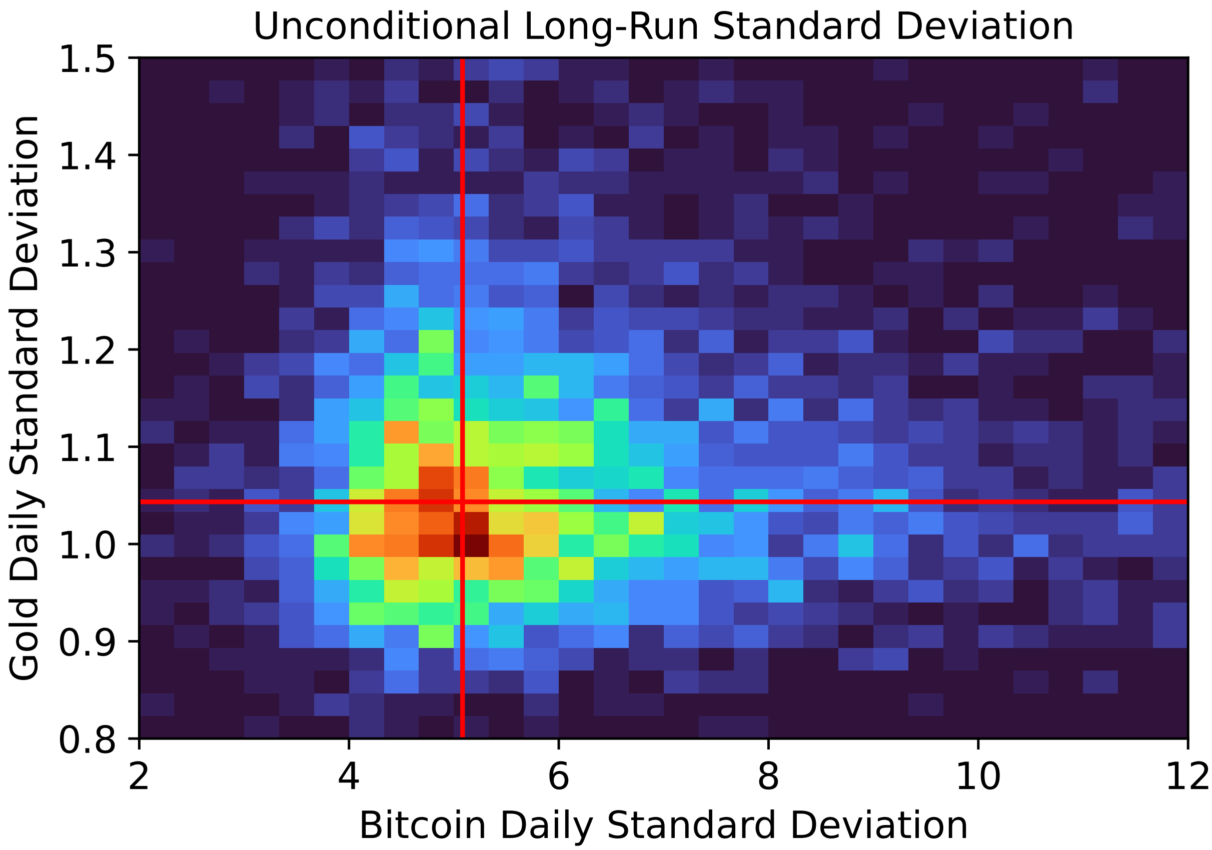

Figure 8: A heatmap of the posterior long-run standard deviation for Bitcoin and gold, compared with the sample standard deviation (red lines).

The relationship between Bitcoin's degrees of freedom parameter – which controls the extent of heavy-tailed behavior of its marginal distribution – and the long-run volatility of Bitcoin is displayed in Figure 9. The crescent-like shape of the posterior distribution indicates a relationship between samples in which Bitcoin is especially high risk (low degrees of freedom, high standard deviation), and samples in which Bitcoin is lower risk (high degrees of freedom, low standard deviation).

Lastly, it is informative to examine the distribution of $\boldsymbol{S}$, the long-run correlations across assets. Table 3 shows that Bitcoin is mostly uncorrelated with all reserve assets, except the stock market, with which Bitcoin has a moderate positive correlation.

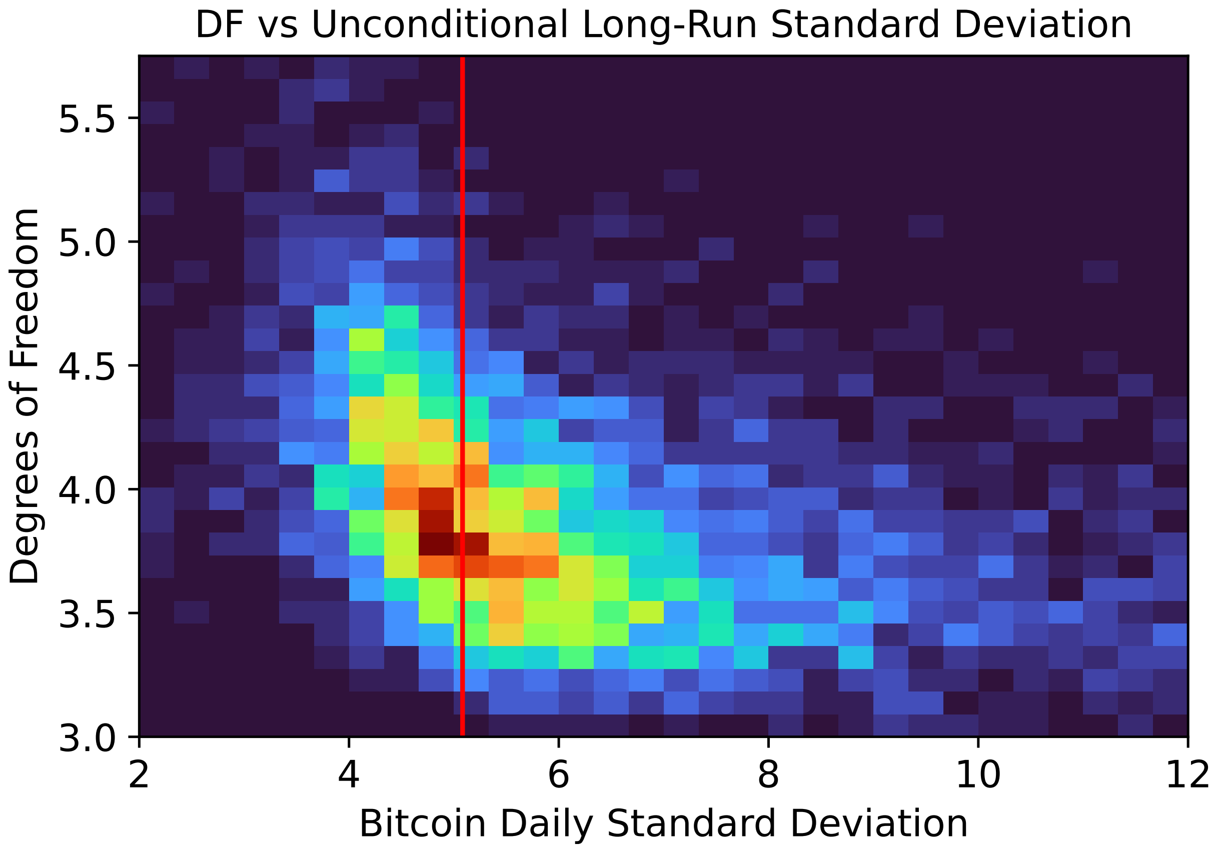

Figure 9: A heatmap of the posterior degrees of freedom of Bitcoin, compared with Bitcoin's posterior long-run standard deviation. The sample standard deviation is plotted as a red line.

Table 3: Median Posterior Long-Run Correlations

	Bitcoin	Gold	Treasuries	Euro bonds	World stock
Bitcoin	1				
Gold	0.07	1			
Treasuries	-0.01	0.21	1		
Euro bonds	0.16	0.41	0.18	1	
World stock	0.34	0.16	-0.08	0.36	1

6 Model Validation

Gelman et al. (2013) describe posterior predictive checks as a means of validating Bayesian models by simulating draws from the posterior distribution, and comparing those draws to the observed data. Because I conduct portfolio choice optimization over simulated 3-month periods, I simulate 3 months of returns (63 trading days) from the posterior distribution and compare the simulations to rolling 3-month periods from the data in various ways.

The subsequent sections illustrate several posterior predictive checks for Bitcoin. Checks for other assets in the model will be available in an online appendix. Overall, the model captures the characteristics of Bitcoin over 3-month periods.

6.1 Gross Return

Figure 10 demonstrates that the simulated returns are significantly more pessimistic than rolling 3-month periods in the sample. This is an intended consequence of the informative prior.

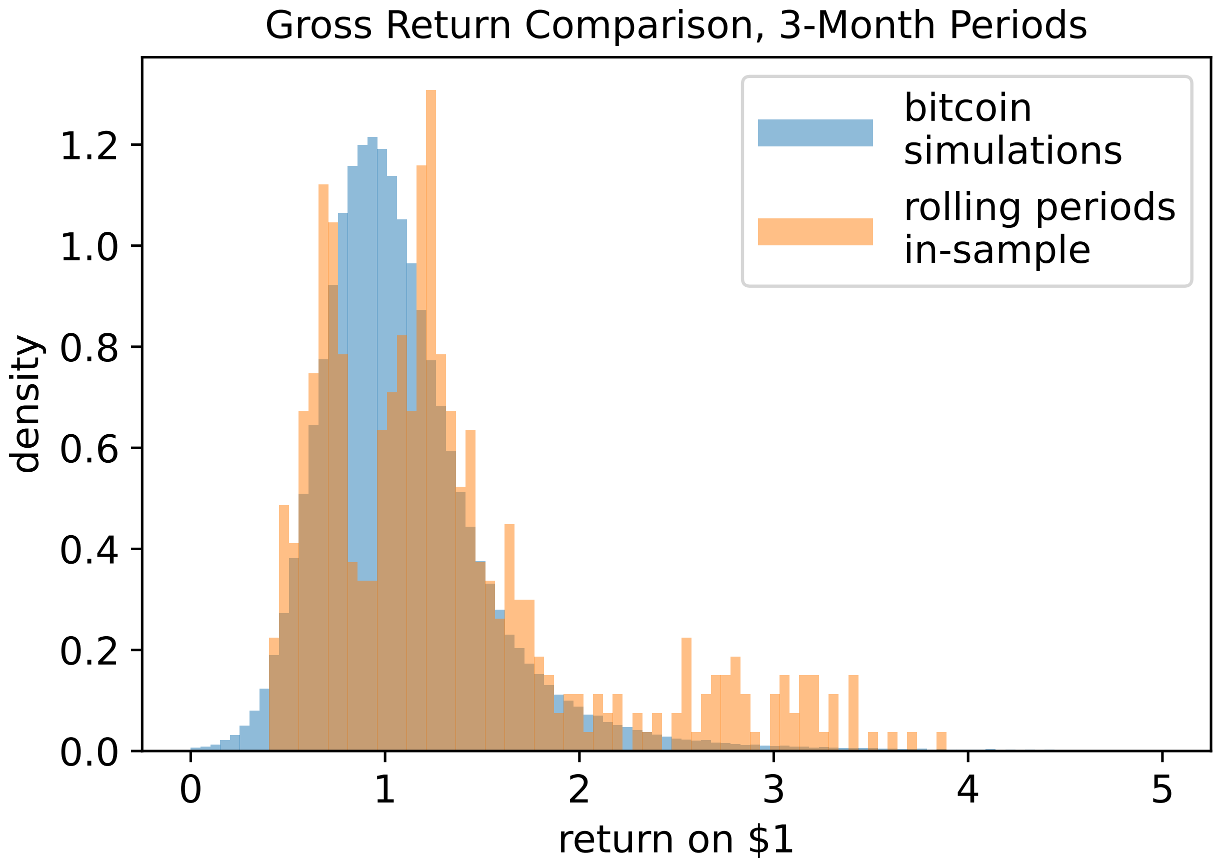

Figure 10: A comparison of gross returns from the model simulations, and gross returns from rolling 3-month periods in the data.

6.2 Correlation

Figures 11 and 12 compare simulated correlations between Bitcoin and each other asset to 3-month rolling periods in the data. In all cases, the simulations are in strong agreement with the data. Other comparisons can be found in the appendix.

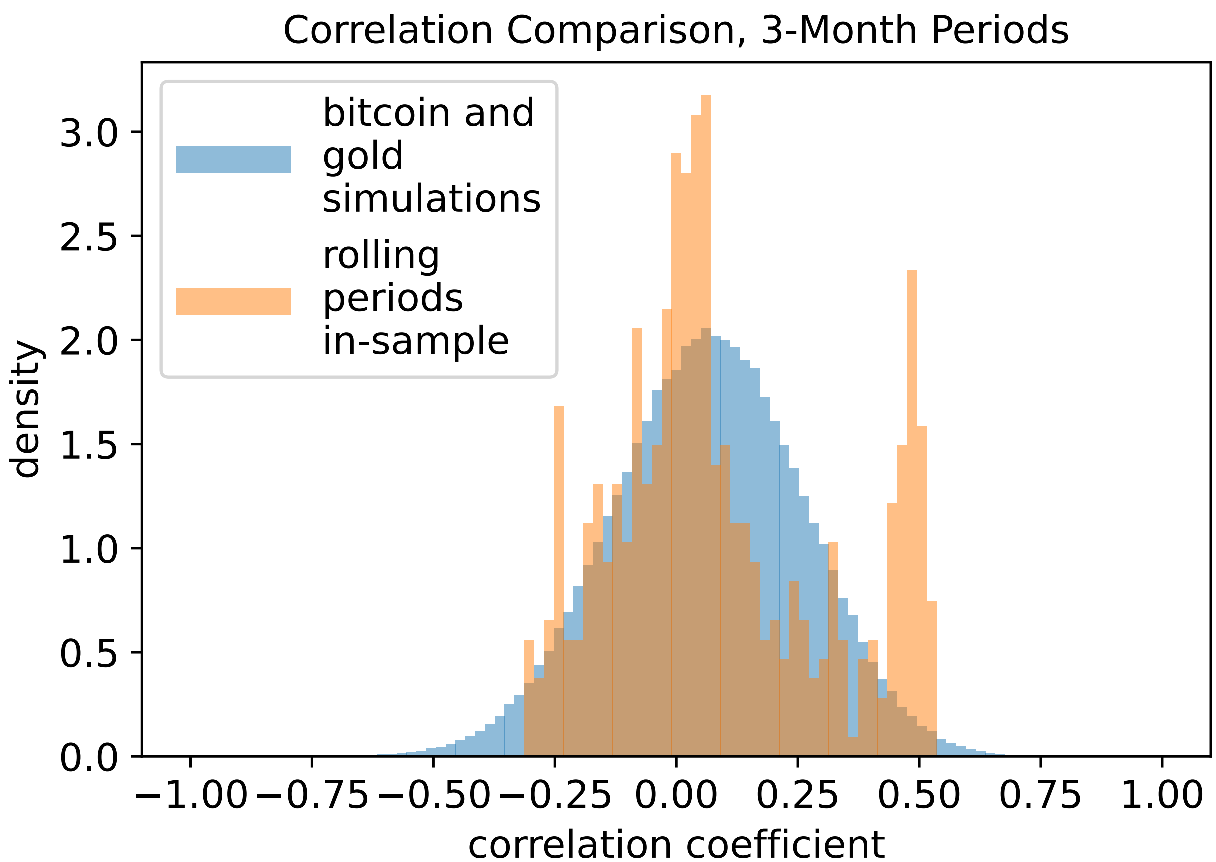

Figure 11: A comparison of the correlation between Bitcoin and gold from the model simulations, and the same correlation from rolling 3-month periods in the data.

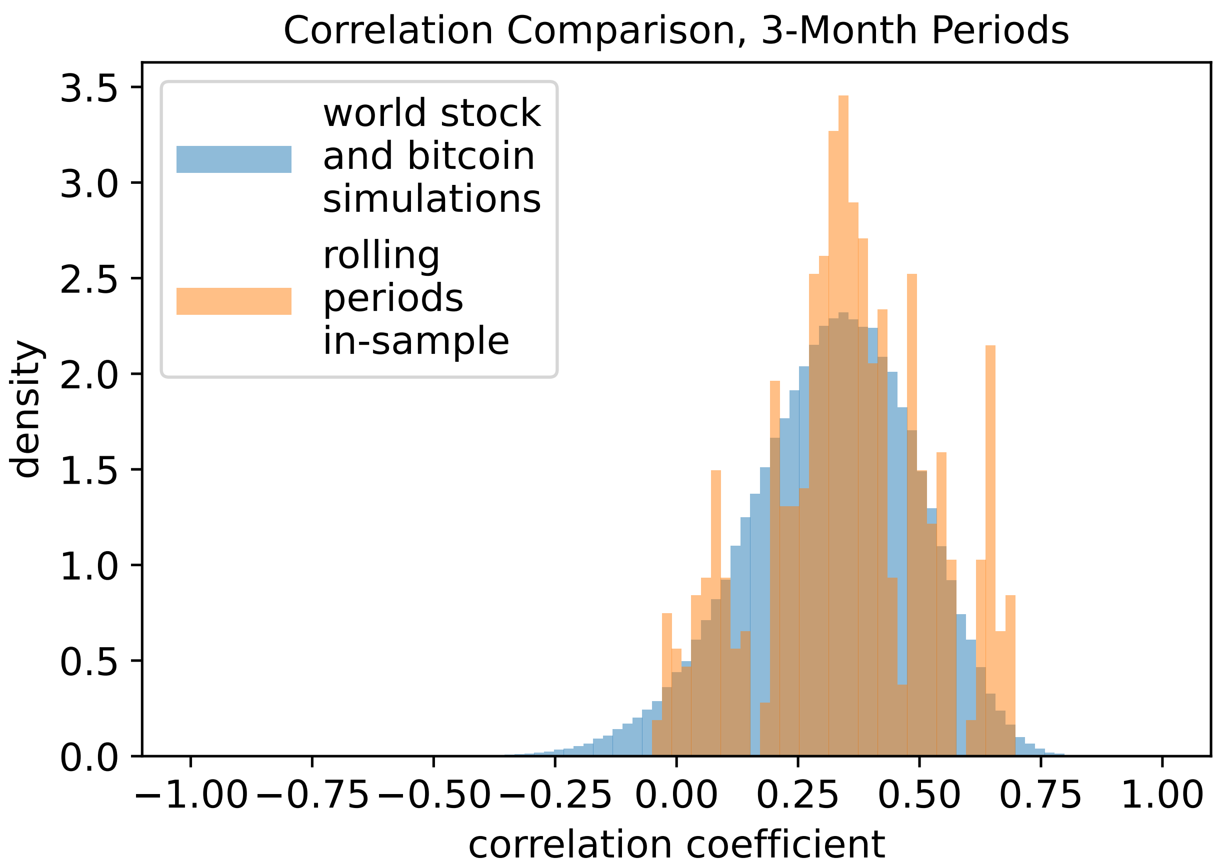

Figure 12: A comparison of the correlation between Bitcoin and world stock from the model simulations, and the same correlation from rolling 3-month periods in the data.

6.3 Standard Deviation

Figure 13 shows that the standard deviation of Bitcoin's log returns in the simulations aligns well with rolling periods in-sample. The standard deviation of the simulations has a long right tail.

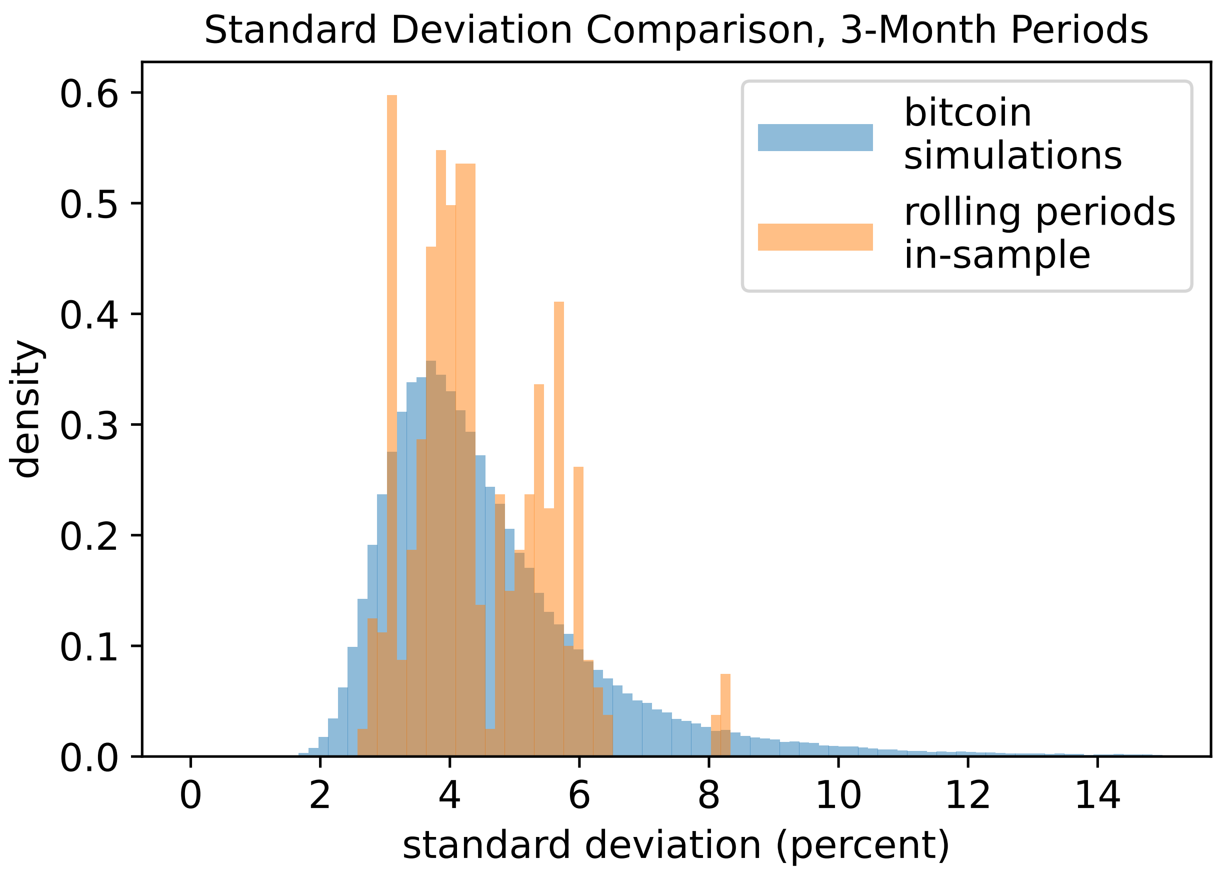

Figure 13: A comparison of the standard deviation of Bitcoin from the model simulations, and the same standard deviation from rolling 3-month periods in the data.

6.4 Skewness

Figure 14 demonstrates that the skewness in the Bitcoin simulations is centered at zero, because of the symmetric Student-t distribution of the error term. The simulations assign less mass to the left tail of the skewness distribution, compared to the sample.

While GARCH models are often estimated using skewed distributions for the error term, doing so is not practical with an informative prior on the mean return. If the mean return is time-invariant, then the location parameter of the skewed distribution must shift every time period as a function of the skewness and variance, which is not computationally feasible. In my model, although the log returns are unskewed, the gross returns are positively skewed (as a result of the exponential transformation) consistent with the findings of Farago and Hjalmarsson (2019).

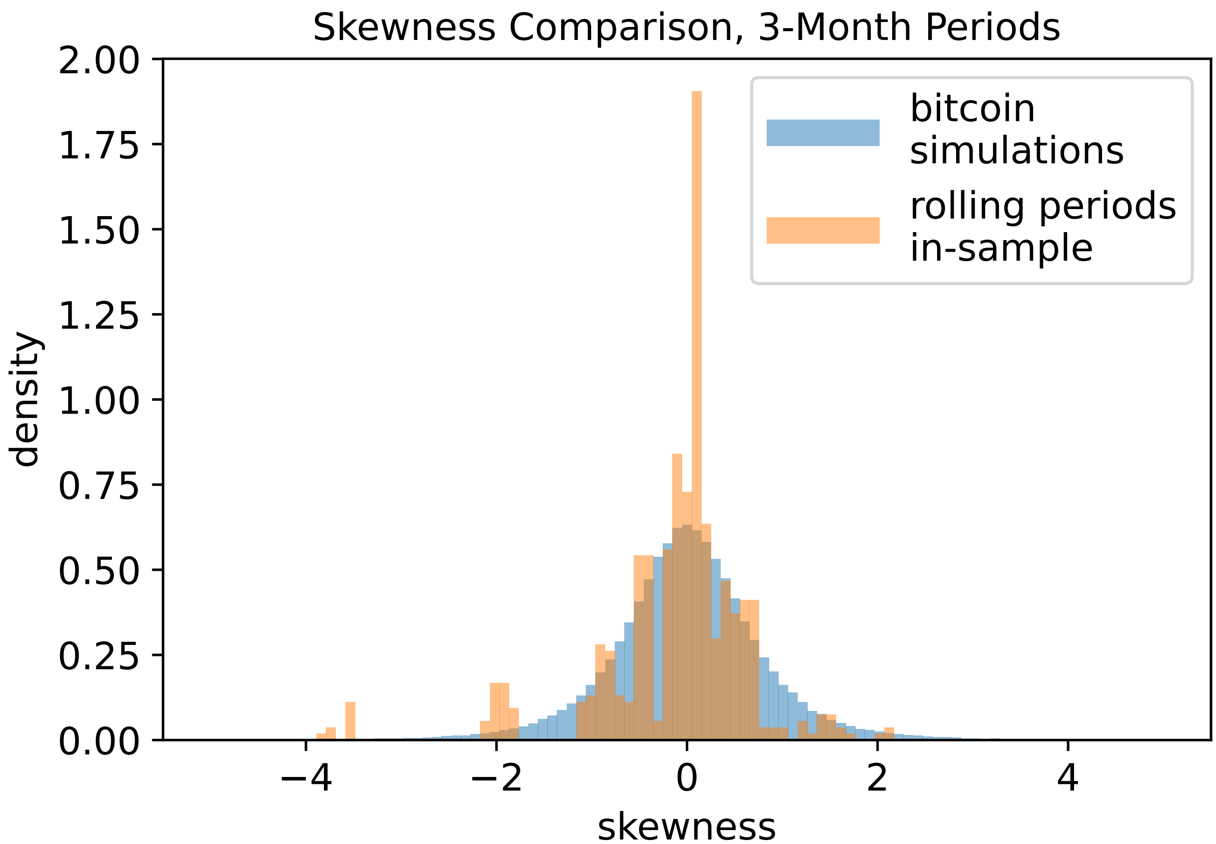

Figure 14: A comparison of the skewness of Bitcoin from the model simulations, and the same skewness from rolling 3-month periods in the data.

6.5 Excess Kurtosis

Figure 15 shows that the simulations slightly understate the kurtosis of Bitcoin (by contrast, the simulations slightly overstate the standard deviation, in Figure 13). It is plausible that the extreme tail risk posed by Bitcoin gradually declined over time as Bitcoin becomes more widely adopted, so that the kurtosis in-sample overstates Bitcoin's tail risk measured today.

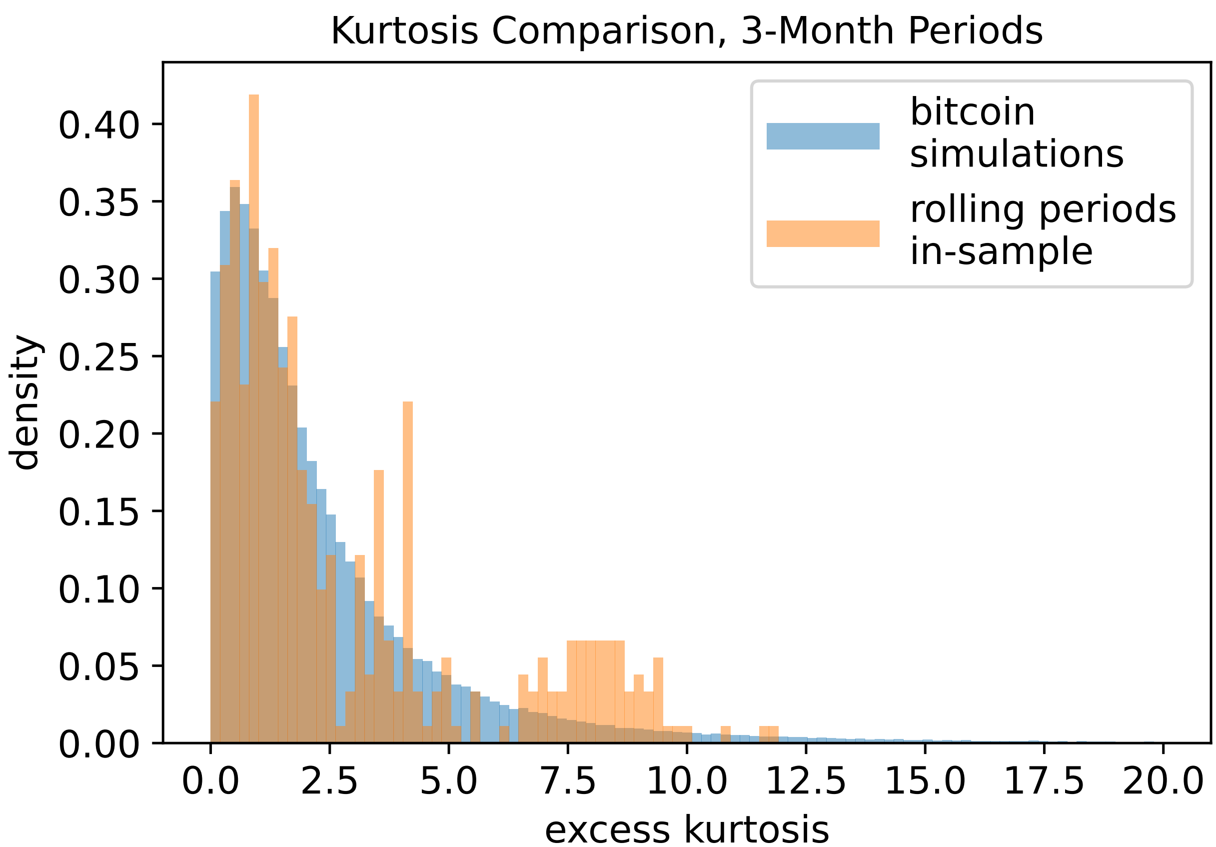

Figure 15: A comparison of the excess kurtosis of Bitcoin from the model simulations, and the same excess kurtosis from rolling 3-month periods in the data.

7 Portfolio Choice Optimization

According to the International Monetary Fund *Guidelines for Foreign Exchange Reserve Management* (2005), central banks tend to review the performance of their reserves quarterly. Therefore, I simulate quarters of returns (63 trading days) from the statistical model. I assume no leverage and no shorting.

After drawing 250,000 simulations from the dynamic Bayesian model, I employ a rejection sampling procedure designed to ensure that the samples are plausible and the portfolio optimization process is numerically stable. It is likely that the heavy-tailed distributions used in my GARCH model, combined with the uncertainty in the long-run variance of the assets, overstate the amount of probability mass in the far tails of the distribution. Additionally, my model does not capture lower-frequency mean-reversion that might reduce the variance. To correct this misspecification, I compute the gross return of each asset for each quarterly sample, and I reject the sample if, for any asset, the corresponding cumulative log return is outside the interval $(63 * \mathbb{E}\mu_i - log(k), 63 * \mathbb{E}\mu_i + log(k))$. I compute the interval width k by calculating the square of the maximum gain or loss across all 3-month rolling periods in the data. For example, if Bitcoin's best performance across rolling periods in the sample was a 4x return, and its worst performance was a 0.5x return, I set $k = max(4, 1/0.5)^2 = 16$. This rejection sampling results in discarding less than 0.5% of the simulated data, but ensures that all accepted samples are economically plausible and the sample variances are numerically stable. Moreover, this procedure preserves the dynamic structure of returns and average compound return within accepted samples.

The portfolio optimization process only considers financial characteristics of the reserve assets. Of course, central banks face many other considerations when allocating their portfolios. In particular, central banks may prefer to align their reserve currency composition with the currency composition of their imports, external debts, and the currency peg (if any). Matching the currency composition of reserve assets with the currency composition of consumption minimizes variability measured in units of consumption. I address these factors (which are exogenous to my model) by implementing constraints, detailed below, for some optimizations that include the risk of financial sanctions. Those constraints correspond to global average reserve shares. Holding the traditional reserve assets in fixed proportions can also be interpreted as enforcing a desire to minimize transaction costs, especially regarding costlier transactions in Bitcoin and physical gold. Insofar as the inputs for the optimization procedure are derived from a Bayesian statistical analysis, this approach is analogous to that of Black and Litterman (1992), who adapt the classic mean-variance framework of Markowitz (1952) in a Bayesian setting. Moreover, my mean-variance approach is similar to that of Papaioannou, Portes, and Siourounis (2006), who use a mean-variance framework to estimate the optimal euro share of reserves.

When modeling sanctions, I treat the sanctions probabilities as exogenous relative to the central bank's decisionmaking. I am unaware of any instance where a central bank's actions provided the primary impetus for a third party to freeze the central bank's reserves. Rather, sanctions result from political decisions made by leaders external to the central bank, of which the central bank probably would not have advance notice. I assume that the US and EU may separately choose to apply sanctions to the central bank, resulting in the total loss of the central bank's US Treasuries or Euro bonds, but leaving the central bank's gold and cryptocurrency untouched.[24] I assume that US sanctions result in a 2/3 loss in the value of the global stock holdings, while EU sanctions result in a 1/3 loss (a total loss occurs if both US and EU sanctions are applied). In my base case, the probability of US sanctions is 1/100, the probability of EU sanctions is 1/200, and the correlation coefficient between the two is 0.4. I implement the correlation across sanctions probabilities with a Gaussian copula. I conduct sensitivity analysis by varying these parameters, as shown in Table 4.

[24]In assuming that gold is not affected by sanctions, I implicitly assume that the central bank retains physical custody of its gold, rather than storing its gold within third-party custodians such as the vaults of the New York Fed. I do not account for the cost of transporting or storing gold.

Table 4: Scenarios for Portfolio Optimization with Sanctions

Scenario	US Sanction Probability	EU Sanction Probability	US & EU Correlation
Base Case	1/100	1/200	0.4
High Risk	1/10	1/20	0.4
Low Risk	1/1000	1/2000	0.4
High Correlation	1/100	1/200	0.8
Low Correlation	1/100	1/200	0

To implement mean-variance preferences, I solve the following optimization problem, where ψ is the degree of risk aversion:

$$\begin{aligned} \underset{\vec{w}}{\text{maximize}} \quad & \sum_{i=1}^{N_{sim}} \frac{x_i}{N_{sim}} - \psi \sum_{j=1}^{N_{sim}} \frac{(x_j - 1/N_{sim} \sum_{k=1}^{N_{sim}} x_k)^2}{N_{sim}} \\ \text{subject to} \quad & w_i \geq 0, \quad i = 0, \ldots, N_{assets}, \\ & \sum_{i=1}^{N_{assets}} w_i = 1. \end{aligned} \tag{12}$$

x_i is the geometric return resulting from the ith simulation, and $0 \leq \Omega_k \leq 1$ is an indicator for the extent of sanctions applied to asset k:

$$x_i(\vec{w}) = \left(\sum_{k=1}^{k=N_{assets}} \Omega_k w_k \prod_{t=1}^{T=63} exp(r_{i,k,t}) \right)^{1/63} - 1 \tag{13}$$

Two aspects of this optimization process are noteworthy. First, I assume that the central bank does not rebalance cross assets within the quarter. Chinn, Ito, and McCauley (2021) find mixed evidence regarding whether central banks rebalance. By assuming no rebalancing, I avoid the necessity of accounting for transaction costs regarding physical gold and Bitcoin, which are likely to be much more costly than transactions in fiat assets. Additionally, without rebalancing, the timing of sanctions within the quarter does not affect the final portfolio value, which simplifies the optimization procedure.

Second, I operate using the geometric return rather than the expected return. The geometric return is the best measure of performance over longer periods of time (such as 63 trading days), because the geometric return accounts for volatility decay. As Hughson, Stutzer, and Yung (2006) point out, the expected return of highly volatile assets is much higher than their compound return, because an asset which experiences a sequence of equal percentage gains and losses does not return to its starting value. Accounting for volatility decay over the appropriate time horizon is particularly important when modeling Bitcoin, which will appear more attractive over shorter holding periods where the effect of volatility decay is less pronounced. Computing the geometric return over 63 trading days produces results that are similar to those that I would obtain from optimization using the continuously-compounded log returns, a procedure that Zhang (2021) explores.

7.1 Without Sanctions

Figure 16 displays optimal shares without sanctions. The results show that an investor will blend an increasing share of a safe asset (Treasuries) with a set of risky assets (stocks and cryptocurrency) as the investor becomes more risk averse. The investor does not hold Euro bonds because the exchange rate risk is not compensated, and similarly, gold is not particularly attractive as long as Treasuries pay a positive real rate of return. Even the most risk-averse investor holds 2-3% in Bitcoin.[25]

[25] As a point of comparison, El Salvador's Bitcoin represents about 1.4% of its international reserves, as of September 2022.

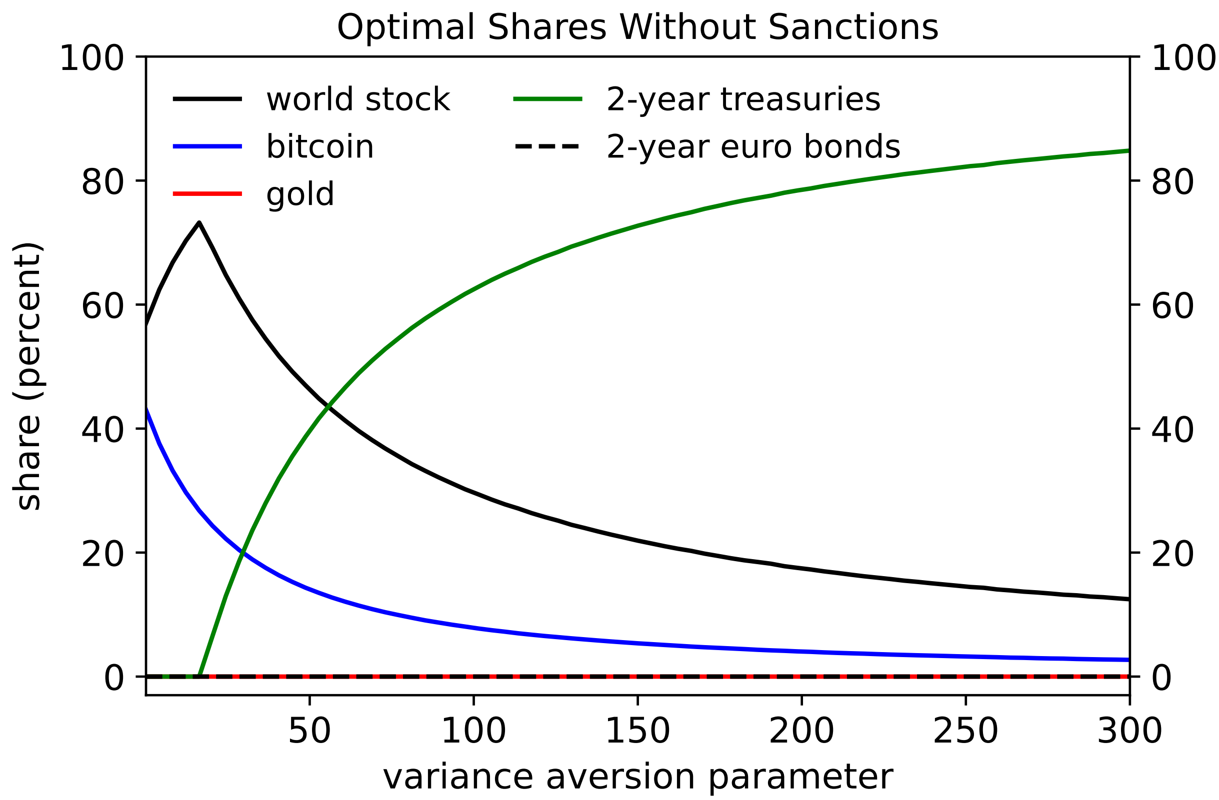

Figure 16: Optimal asset shares with mean-variance preferences.

7.2 With Sanctions, Varying All Shares

Figure 17 displays optimal Bitcoin shares in the mean-variance framework and the base sanctions case. Sanctions motivate major changes in the portfolio shares. Notably, the most risk-averse investor will diversify across all five assets, gold becomes far more attractive, Euro bonds are more appealing than US Treasuries due to their lower sanctions risk, and the optimal Bitcoin share rises to about 5%.

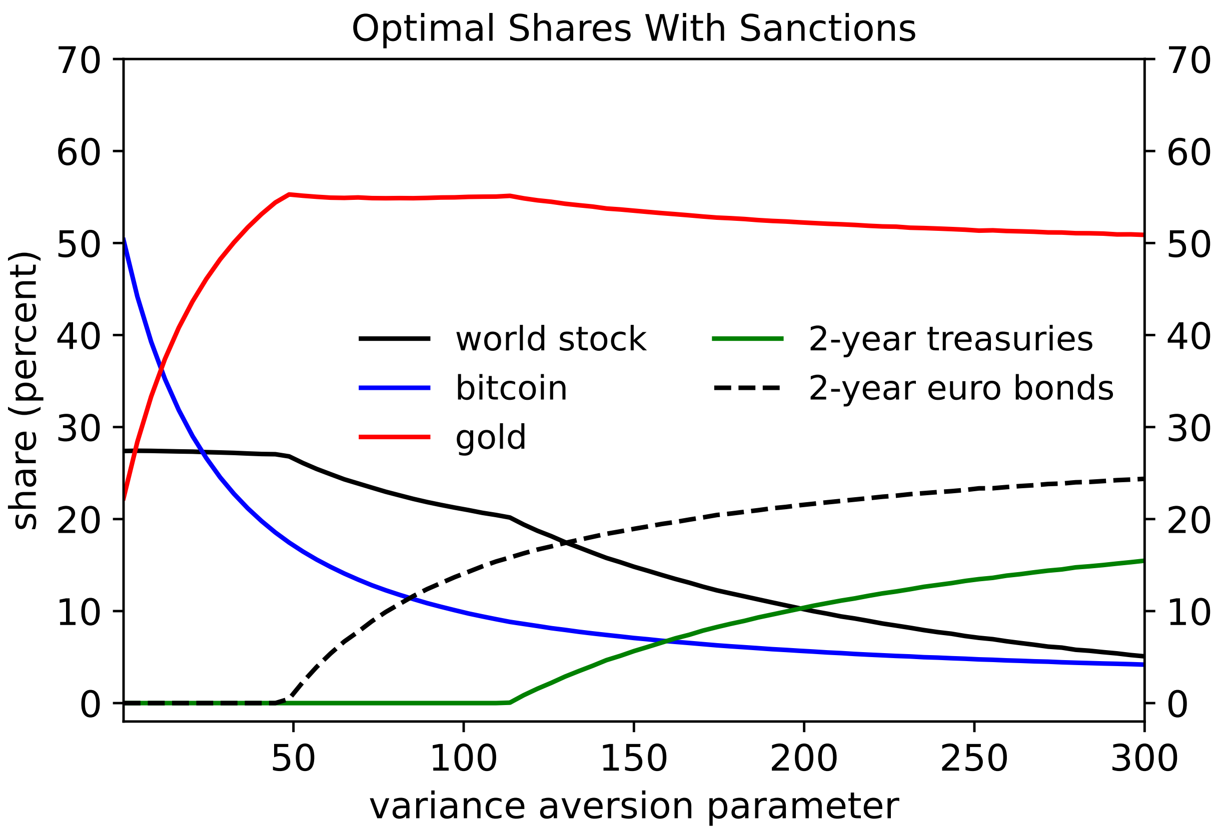

Figure 17: Optimal asset shares with mean-variance preferences and sanctions.

7.3 With Sanctions, Varying Bitcoin Only

The allocations proposed in Figure 17 may not be feasible for some central banks, who cannot acquire or store such a large quantity of physical gold.[26] Furthermore, as previously mentioned, central banks may prefer (to the extent possible) matching their currency composition with the currency composition of their imports and external debts. Therefore, I repeat the portfolio optimization varying only the Bitcoin share, assuming that the rest of the portfolio consists of 12% gold, 51% US Treasuries, 17% Euro bonds, and 20% world stock. These percentages approximately correspond to global averages. Accordingly, this analysis addresses the question of how much Bitcoin a representative central bank might want to add to its existing reserve holdings in order to address its sanctions risk, without

[26] About $150 billion of gold, or 3,000 metric tons, is mined every year.

radically altering its existing holdings.

Figures 18 and 19 display results when varying only the Bitcoin share. In this case, the central bank will hold a significant share of Bitcoin to hedge against sanctions, especially when the US and EU sanctions are highly correlated.

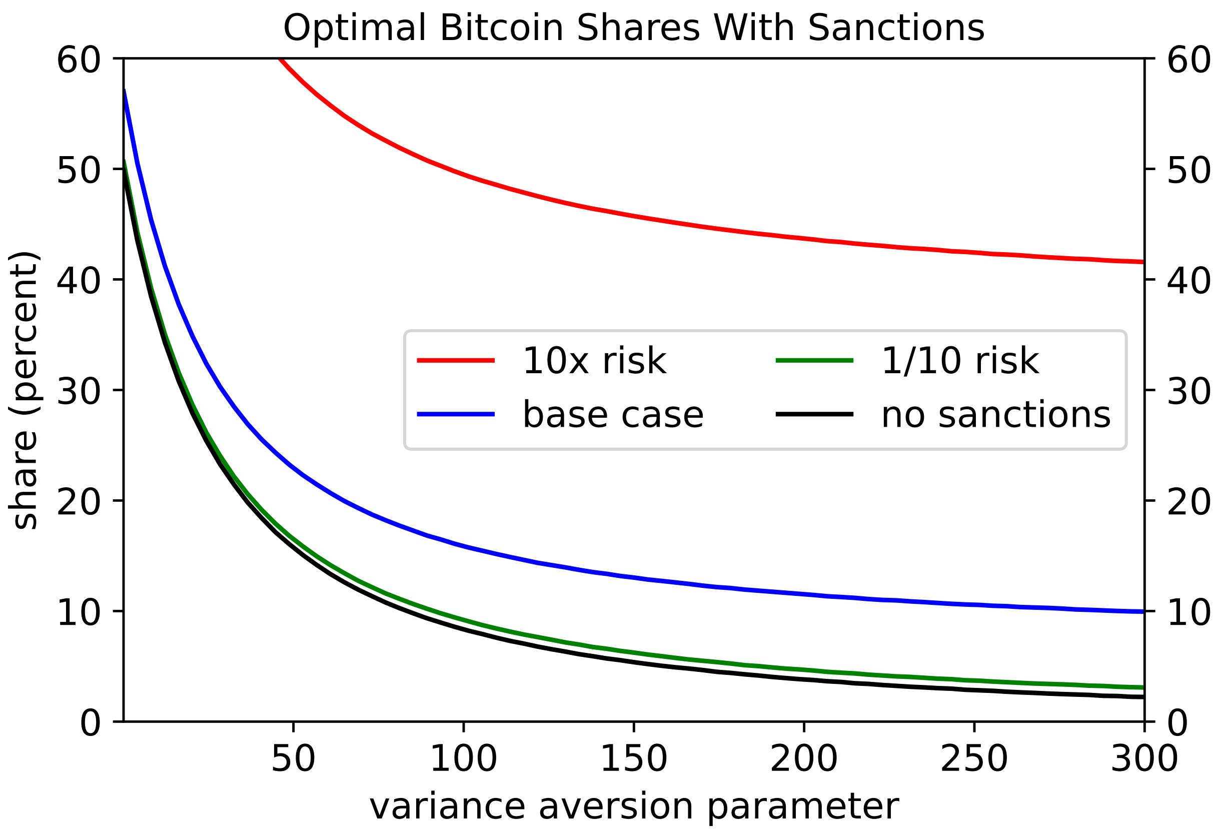

Figure 18: Optimal Bitcoin shares with mean-variance preferences and sanctions, varying only Bitcoin.

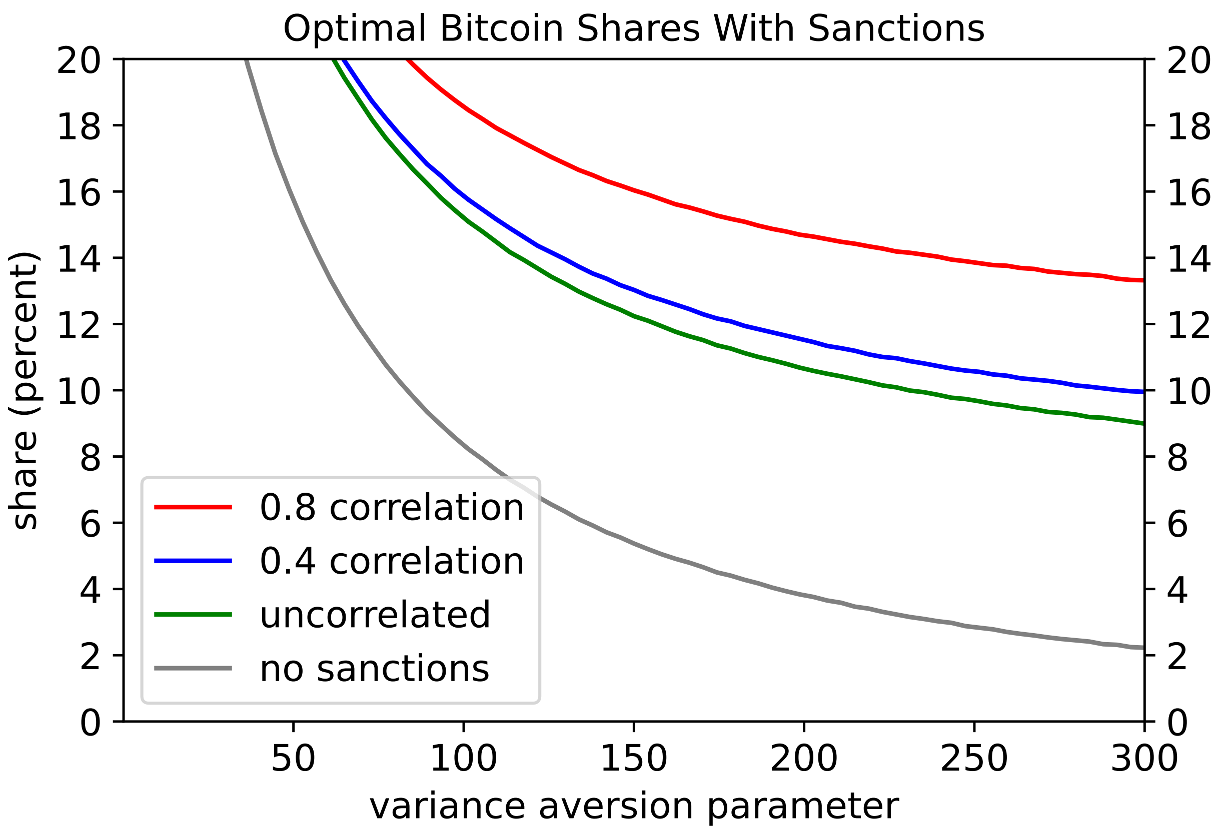

Figure 19: Optimal Bitcoin shares with mean-variance preferences and sanctions, varying only Bitcoin.

7.4 With Sanctions, Varying Gold and Bitcoin

To explore the extent to which gold and Bitcoin are substitutes as hedges against sanctions risk, I vary both the gold and Bitcoin share, assuming that the rest of the portfolio is allocated 58% to Treasuries, 19% to Euro bonds, and 23% to world stock (approximate global averages).

Figures 20 and 21 display the results. Although the central bank holds additional Bitcoin compared to the scenario without sanctions, gold is the preferred asset to hedge against sanctions risk. Specifically, the risk of international sanctions rationalizes holding gold even when Treasuries offer less volatility and a positive real return. Moreover, in the lowest risk case, the optimal gold share rises as a function of risk aversion.

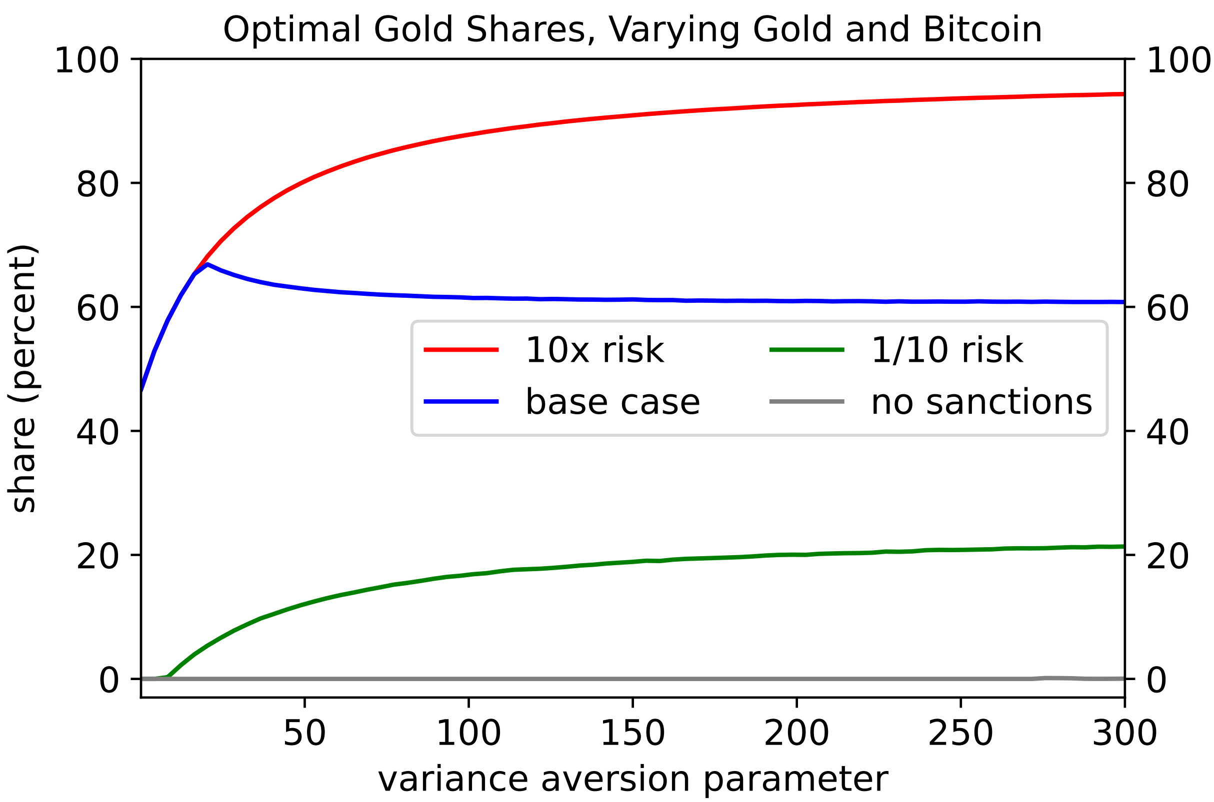

Figure 20: Optimal gold shares with mean-variance preferences and sanctions, varying gold and Bitcoin.

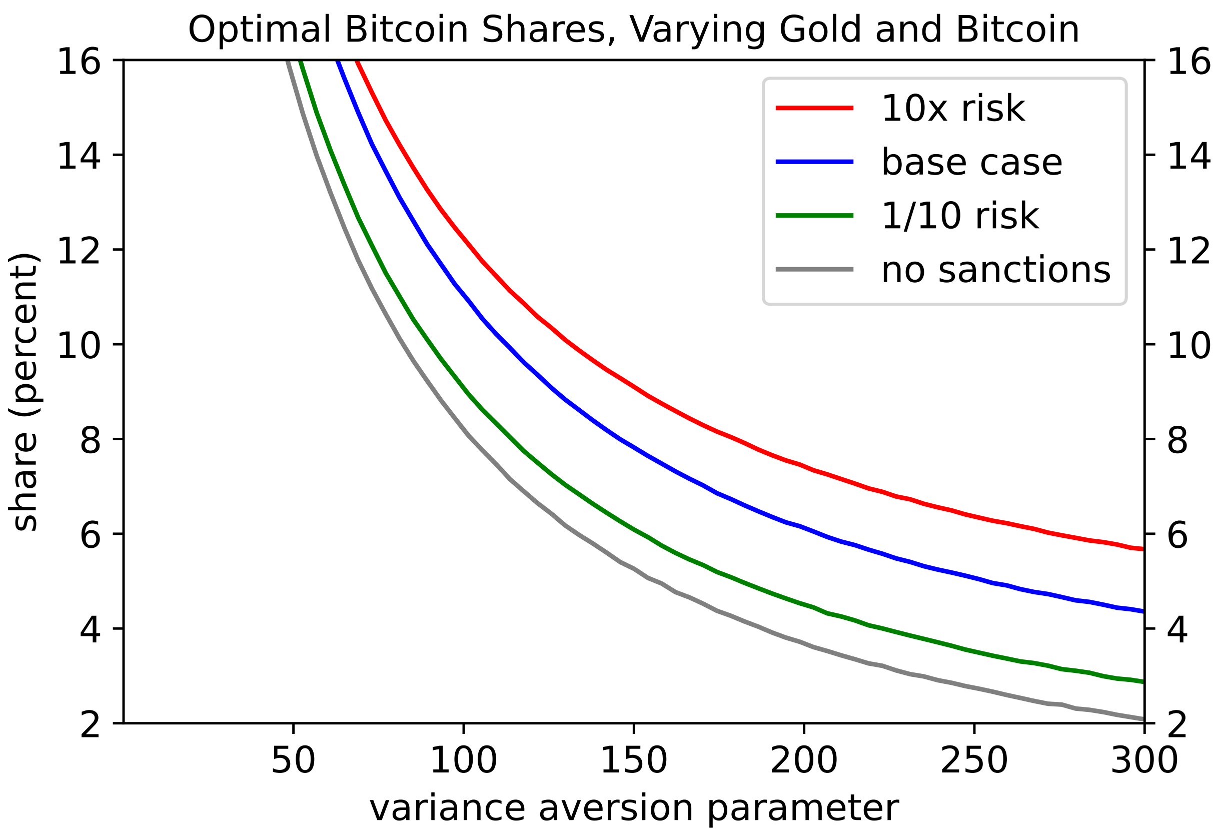

Figure 21: Optimal Bitcoin shares with mean-variance preferences and sanctions, varying gold and Bitcoin.

8 Conclusion

The risk of financial sanctions by major fiat reserve currency issuers has significant implications for central banks, some of whom may not be able to count on US Treasuries or AAA-rated Euro bonds as safe assets. Indeed, in the presence of sanctions, there is no totally safe asset. Cryptocurrencies offer some protection against sanctions, but introduce the risk of high price volatility. The price of gold is also more volatile than that of Treasuries or Euro bonds. Although holding physical gold also provides protection against sanctions, gold is less liquid than fiat assets, and assuming physical custody of gold entails significant logistical and security costs.

There are several avenues for further research. It is possible that different copulas and error distributions may generate even more realistic simulations. The objective function could be modified in many other ways, such as to express loss aversion. A model featuring rebalancing and portfolio adjustment costs might also be more realistic, at the expense of introducing additional parameters. The outputs of this model – optimal portfolio shares as a function of risk aversion and sanctions risk – could be aggregated across a distribution of countries to form demand curves for each asset, creating a general equilibrium model that could be used to estimate the implicit cost of US financial sanctions in terms of their effect on US or EU interest rates.

To the extent that China's economic and political objectives do not align with those of the US and EU, it is conceivable that sanctions risk by the Chinese government is negatively correlated with that of the US and EU. This "feature" of the renminbi may boost its attractiveness as a reserve asset, and reduce the appeal of cryptocurrency for countries that hold renminbi. Like cryptocurrency, the optimal renminbi allocation rises as the correlation between US and EU sanctions rises.

If a central bank does decide to purchase cryptocurrency, the central bank faces a choice of whether to publicly reveal that decision. Choosing to conceal the central bank's Bitcoin allocation might further stymie external attempts to freeze the central bank's assets. Ferranti (2022) discusses the fact that many central banks do not disclose their fiat reserve currency composition. Revealing the central bank's cryptocurrency wallets enables public verification of the central bank's assets, but requires the central bank to accept scrutiny regarding its choice to invest in a highly volatile asset. Aizenman and Inoue (2012) find that central banks tend to underreport their gold holdings to avoid criticism when the price of gold declines.

Lastly, central bank digital currencies may enhance the benefits of cryptocurrency. Theoretically, a central bank could conduct a foreign exchange intervention relying solely upon domestic payments infrastructure by offering to sell the central bank's cryptocurrency to holders of the central bank's digital currency. In that sense, a country's decision to embrace cryptocurrency may boost its resilience to economic shocks.

References

Abramowitz, Milton and Irene Stegun, eds. (1972). *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables*. 10th ed. National Bureau of Standards: Applied Mathematics Series.

Ahn, Daniel P. and Rodney D. Ludema (2020). "The sword and the shield: The economics of targeted sanctions." In: *European Economic Review* 130, p. 103587.

Aizenman, Joshua and Kenta Inoue (2012). *Central Banks and Gold Puzzles*. Working Paper 17894. NBER.

Arslanalp, Serkan, Barry Eichengreen, and Chima Simpson-Bell (Mar. 2022). *The Stealth Erosion of Dollar Dominance: Active Diversifiers and the Rise of Nontraditional Reserve Currencies*. Working Paper WP/22/58. IMF.

Asness, Cliff (Mar. 2021). *The Long Run Is Lying to You*. Tech. rep. AQR Capital Management.

Aysan, Ahmet Faruk et al. (Jan. 2019). "Effects of the geopolitical risks on Bitcoin returns and volatility." In: *Research in International Business and Finance* 47, pp. 511–518.

Biais, Bruno et al. (Feb. 2022). "Equilibrium Bitcoin Pricing." In: *Journal of Finance (forthcoming)*.

Black, Fischer and Robert Litterman (1992). "Global Portfolio Optimization." In: *Financial Analysts Journal* 48(5), pp. 28–43.

Bratspies, Rebecca M. (2018). "Cryptocurrency and the Myth of the Trustless Transaction." In: *Michigan Technology Law Review* 25 (1).

Carpenter, Bob et al. (2017). "Stan: A Probabilistic Programming Language." In: *Journal of Statistical Software* 76 (1), pp. 1–32.

Cheah, Eng-Tuck and John Fry (2015). "Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin." In: *Economics Letters* 130, pp. 32–36.

Chinn, Menzie, Hiro Ito, and Robert McCauley (Aug. 2021). *Do Central Banks Rebalance Their Currency Shares?* Working Paper 29190. NBER.

Demarta, Stefano and Alexander McNeil (2005). "The t-Copula and Related Copulas." In: *International Statistical Review* 73(1), pp. 111–129.

Engle, Robert (2002). "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models." In: *Journal of Business & Economic Statistics* 20 (3), pp. 339–350.

Engle, Robert and Victor Ng (1993). "Measuring and Testing the Impact of News on Volatility." In: *Journal of Finance* 48(5), pp. 1749–1778.

Farago, Adam and Erik Hjalmarsson (June 2019). *Compound Returns*. Working Paper 767. Department of Economics and Centre for Finance, University of Gothenburg.

Felbermayr, Gabriel et al. (2020). "The global sanctions data base." In: *European Economic Review* 129, p. 103561.

Ferranti, Matthew (June 2022). *Estimating the Currency Composition of Foreign Exchange Reserves*. Working Paper 2206.13751. arXiv.

Fioruci, Jose A., Ricardo S. Ehlers, and Francisco Louzada (2014). *BayesDccGarch - An Implementation of Multivariate GARCH DCC Models*. Tech. rep. 1412.2967. arXiv.

Gelman, Andrew (2006). "Prior distributions for variance parameters in hierarchical models." In: *Bayesian Analysis* 1(3), pp. 515–533.

Gelman, Andrew et al. (2013). *Bayesian Data Analysis*. 3rd ed. Chapman and Hall/CRC. ISBN: 1439840954.

Guidelines for Foreign Exchange Reserve Management (2005). Research Report. IMF.

Halaburda, Hanna et al. (2022). "The Microeconomics of Cryptocurrencies." In: *Journal of Economic Literature* 60(3), pp. 971–1013.

Hardle, Wolfgang Karl, Campbell R. Harvey, and Raphael C. G. Reule (2020). "Understanding Cryptocurrencies." In: *Journal of Financial Econometrics* 18(2), pp. 181–208.

Harmson, Stephen (Nov. 1998). *Gold as a Store of Value*. Research Study No. 22. World Gold Council.

Hayes, Adam S. (2019). "Bitcoin price and its marginal cost of production: support for a fundamental value." In: *Applied Economics Letters* 26 (7), pp. 554–560.

Hufbauer, Gary Clyde and Euijin Jung (2020). "What's new in economic sanctions?" In: *European Economic Review* 130, p. 103572.

Hufbauer, Gary Clyde, Jeffrey Schott, et al. (2009). *Economic Sanctions Reconsidered*. 3rd ed. Peterson Institute for International Economics. ISBN: 9780881324129.

Hughson, Eric, Michael Stutzer, and Chris Yung (2006). "The Misuse of Expected Returns." In: *Financial Analysts Journal* 62 (6), pp. 88–96.

Juarez, Miguel and Mark Steel (2010). "Model-Based Clustering of Non-Gaussian Panel Data Based on Skew-t Distributions." In: *Journal of Business & Economic Statistics* 28(1), pp. 52–66.

Klages-Mundt, Ariah and Andreea Minca (May 2021). *While Stability Lasts: A stochastic Model of Stablecoins*. Working Paper 2004.01304v2. arXiv.

Kreuzer, Alexander and Claudia Czado (Nov. 2019). *Bayesian inference for dynamic vine copulas in higher dimensions*. Working Paper 1911.00702. arXiv.

Liu, Yukun and Aleh Tsyvinski (2018). *Risks and Returns of Cryptocurrency*. Working Paper 24877. NBER.

Luckner, Clemens Graf von, Carmen M. Reinhart, and Kenneth S. Rogoff (2021). *Decrypting New Age International Capital Flows*. Working Paper 29337. NBER.

Makarov, Igor and Antoinette Schoar (2020). "Trading and arbitrage in cryptocurrency markets." In: *Journal of Financial Economics* 135, pp. 293–319.

Markowitz, Harry (1952). "Portfolio Selection." In: *Journal of Finance* 7(1), pp. 77–91.

Mulder, Nicholas (2022). *The Economic Weapon: The Rise of Sanctions as a Tool of Modern War*. Yale University Press. ISBN: 0300259360.

Nguyen, Quynh Nga et al. (2020). “Hedging and safe-haven characteristics of Gold against currencies: An investigation based on multivariate dynamic copula theory.” In: *Resources Policy* 68, p. 101766.

Pakel, Cavit et al. (2021). “Fitting vast dimensional time-varying covariance models.” In: *Journal of Business & Economic Statistics* 39 (3), pp. 652–668.

Papaioannou, Elias, Richard Portes, and Gregorios Siourounis (2006). “Optimal Currency Shares In International Reserves: The Impact of the Euro and The Prospects for the Dollar.” In: *Journal of the Japanese and International Economies* 20(4), pp. 508–547.

Patton, Andrew (2012). “A review of copula models for economic time series.” In: *Journal of Multivariate Analysis* 110, pp. 4–18.

Shanaev, Savva et al. (2020). “Cryptocurrency Value and 51% Attacks: Evidence from Event Studies.” In: *Journal of Alternative Investments* 22 (3), pp. 65–77.

Shaw, William (2006). “Sampling Student’s T distribution – use of the inverse cumulative distribution function.” In: *Journal of Computational Finance* 9 (4), pp. 37–73.

Shiferaw, Yegnanew A. (2019). “Time-varying correlation between agricultural commodity and energy price dynamics with Bayesian multivariate DCC-GARCH models.” In: *Physica A: Statistical Mechanics and its Applications* 526, p. 120807.

Sklar, Abe (1959). “Fonctions de Répartition à n Dimensions et Leurs Marges.” In: *Publications de l’Institut Statistique de l’Université de Paris* 8, pp. 229–231.

Sklar, Abe (1973). “Random Variables, Joint Distribution Functions, and Copulas.” In: *Kybernetika* 9(6), pp. 449–460.

Smales, L. A. (2019). “Bitcoin as a safe haven: Is it even worth considering?” In: *Finance Research Letters* 30, pp. 385–393.

So, Mike and Cherry Yeung (2014). “Vine-copula GARCH model with dynamic conditional dependence.” In: *Computational Statistics & Data Analysis* 76, pp. 655–671.

Tang, Chaofeng and Kentaka Aruga (2022). "Relationships among the Fossil Fuel and Financial Markets during the COVID-19 Pandemic: Evidence from Bayesian DCC-MGARCH Models." In: *Sustainability* 14(1).

Woetzel, Jonathan et al. (Nov. 2021). *The rise and rise of the global balance sheet: How productively are we using our wealth?* Tech. rep. McKinsey Global Institute.

Zarate, Juan (2013). *Treasury's War: The Unleashing of a New Era of Financial Warfare.* PublicAffairs. ISBN: 1610391152.

Zhang, Hai (Apr. 2021). *Mean-Variance Portfolio Analysis for Continuous Compound Return.* Working Paper 3830922. SSRN.

A Stablecoins

"Stablecoins" are a different type of cryptocurrency that attempt to maintain a fixed exchange rate against an external quantity, often the US dollar. Stablecoins attempt to do so primarily in two different ways. First, some stablecoin issuers hold US dollar collateral, such as Treasury bills or commercial paper. Second, some stablecoin issuers accept other cryptocurrencies as collateral (including other stablecoins), usually over-collateralizing their US dollar obligations to account for the high price volatility of cryptocurrency. Executing a stablecoin trade over the Ethereum network requires paying a small quantity of Ether as a transaction ("gas") fee. Because stablecoin issuers need to hold and transact in collateral, stablecoins are inherently more centralized than Bitcoin and Ether. In fact, the two largest stablecoin issuers, Circle (which issues US Dollar Coin) and Tether (which issues Tether) retain the ability to block cryptocurrency wallets containing their stablecoins, freezing the stablecoins. As of July 28, 2022, Circle has blocked 48 Ethereum wallets[27] and Tether has blocked 692 Ethereum wallets.[28] No stablecoin currently exists which is resistant to sanctions, backed by sufficient collateral to preclude the risk of losing its peg, and sufficiently liquid to accommodate a central bank's transactions (billions of US dollars).

A stablecoin that is sufficiently decentralized that its issuer could not block transactions probably would need to rely heavily on algorithms for its implementation (introducing various security risks) and exclusively accept other cryptocurrencies as collateral. However, the track record of so-called "algorithmic" stablecoins that are backed by other cryptocurrencies is mixed. In May 2022, the $19 billion stablecoin TerraUSD collapsed when the value of its collateral, another cryptocurrency called Luna, cratered. Many other algorithmic stablecoins, including Basis Cash, Iron Finance, SafeCoin, BitUSD, DigitalDollar, NuBits, and CK USD have also failed.[29] Klages-Mundt and Minca (2021) provide a stochastic model that captures

[27] https://dune.com/phabc/usdc-banned-addresses
[28] https://dune.com/phabc/usdt—banned-addresses
[29] https://indianexpress.com/article/technology/crypto/luna-terra-crash-a-brief-history-of-failed-algorithmic-stablecoins-7934293/

the deflationary deleveraging that the algorithmic stablecoin Dai experienced in March 2020. Speculators who exchanged other cryptocurrencies for Dai were forced to repurchase Dai as their collateral dropped in value, in some cases facing forced liquidation of their collateral at a price of zero due to Ethereum network congestion and illiquidity. As a result, Dai became undercapitalized, while its price rose as high as $1.13. To address this crisis, the Dai governance community (holders of the MKR token) issued additional equity-like MKR tokens to recapitalize Dai, and voted to begin accepting US dollar-backed stablecoins as collateral in order to stabilize the price of Dai. Since March 2020, Dai has successfully maintained its $1.00 soft peg. Dai is currently the largest stablecoin whose governance process is sufficiently decentralized that Dai tokens cannot be blocked or frozen. However, if Dai became widely used as a means of evading sanctions, it would be vulnerable to having its stablecoin collateral frozen, potentially precipitating a collapse in the value of Dai.

Any investor that holds an algorithmic stablecoin backed exclusively by decentralized collateral necessarily assumes intermediation risk. It may be impossible to develop an algorithmic stablecoin that is both resistant to sanctions and capable of maintaining a $1.00 soft peg without directly or indirectly holding fiat collateral. Stablecoins will likely remain an active domain for financial innovation, but they do not currently appear to be suitable as reserve assets.

B Market Characteristics of Bitcoin

B.1 Liquidity

In order to be suitable for institutional investment, such as by central banks, it must be possible to transact in billions of dollars of Bitcoin without incurring extravagant costs. The Bitcoin network structure favors large block trades because transaction fees paid to miners are fixed, not a percentage of the transaction, and those fees are typically just a few dollars (converted from units of Bitcoin), as illustrated in Figure 22. The average daily trading volume of Bitcoin is also on par with other major reserve assets, as displayed in Figure 23. Lastly, Figure 24 shows that the bid / ask spreads of Bitcoin average 0.1% or less across multiple exchanges.

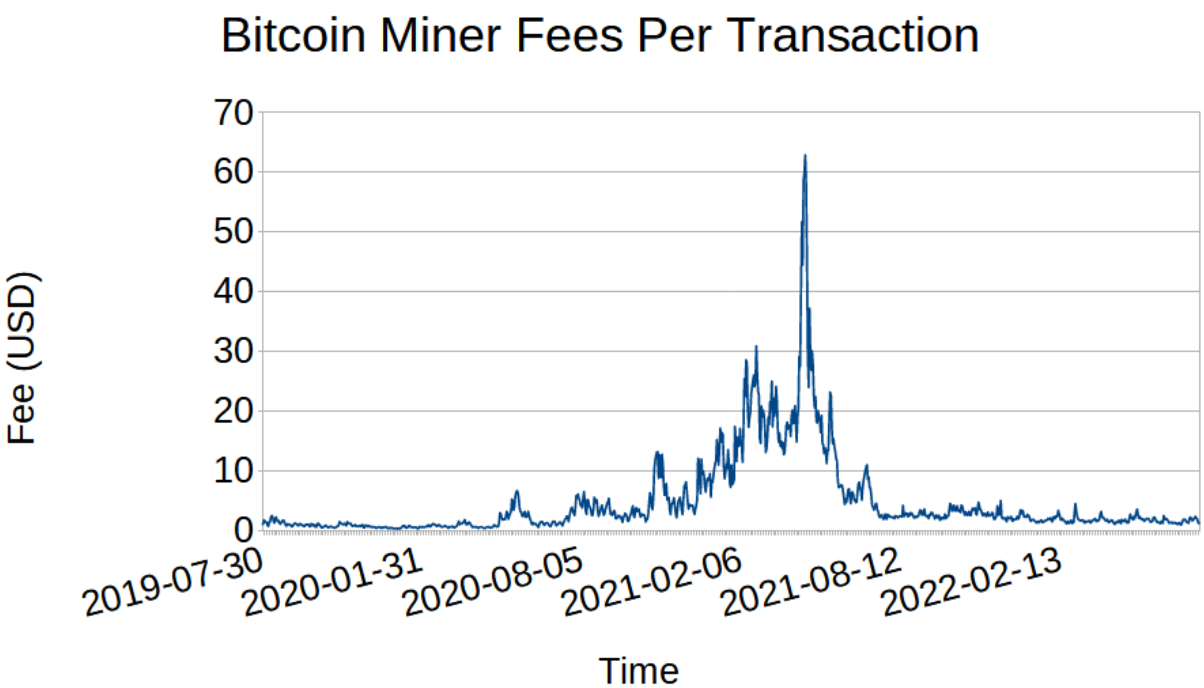

Source: blockchain.com

Figure 22: The average transaction fee for processing an on-chain Bitcoin transaction. The fee is paid in Bitcoin, but converted to US dollars for this chart.

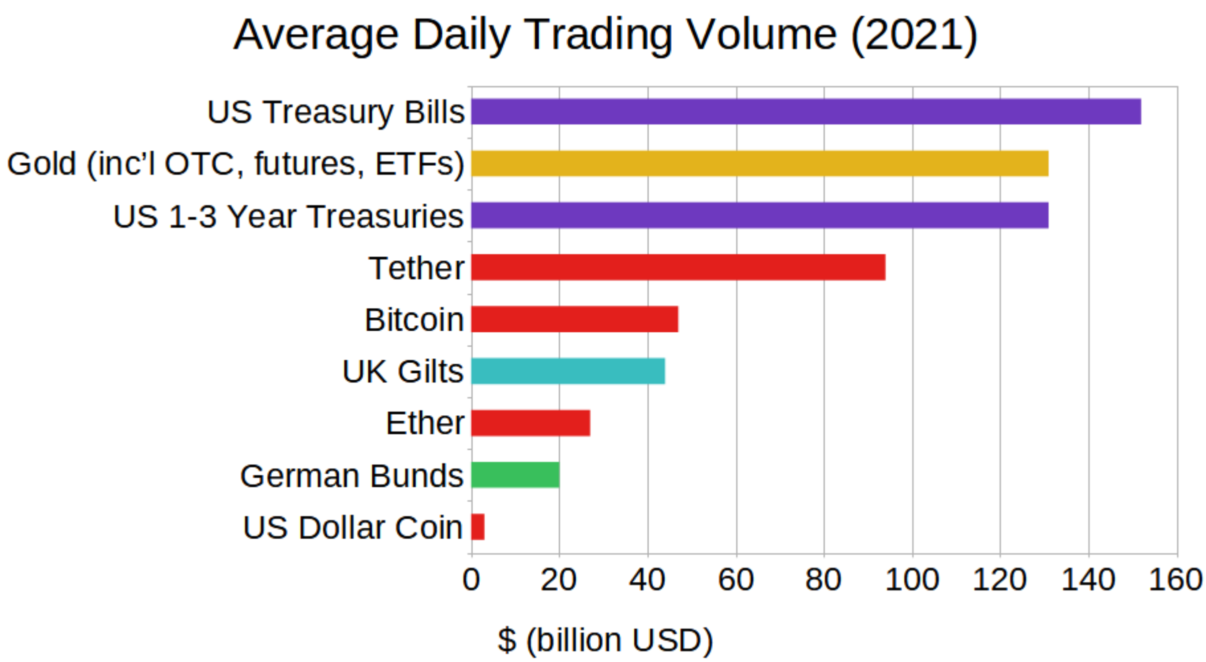

Source: World Gold Council, coinmarketcap.com

Figure 23: Average daily trading volume of major reserve assets and cryptocurrencies, including stablecoins Tether and US Dollar Coin.

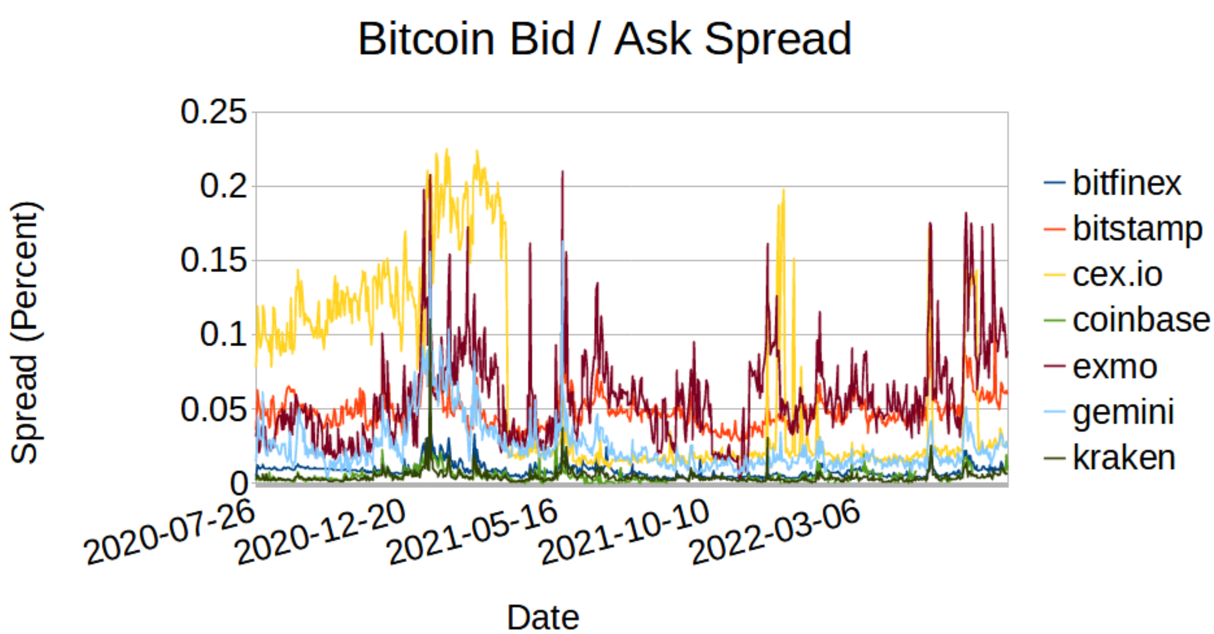

Source: bitcoinity.org

Figure 24: The bid / ask spread of Bitcoin across several cryptocurrency exchanges.

B.2 Flight to Safety

As a risky asset, Bitcoin has historically not exhibited a flight to safety effect; Bitcoin's price tends to fall during periods of economic turmoil. Indeed, Figure 25 shows that Bitcoin declined concurrent with Russia's invasion of Ukraine in February 2022. However, the figure also shows that Bitcoin sharply appreciated immediately following the US Treasury's sanctions against the Central Bank of Russia. Therefore, the decentralized nature of Bitcoin may provide some insurance value against deglobalization shocks, such as the disruption caused by sanctions. This hypothesis is consistent with Aysan et al. (2019), who find that Bitcoin hedges geopolitical risks.

B.3 Store of Value

Estimating the fundamental value of Bitcoin is an active area of research. One strand of literature argues that the fundamental value of Bitcoin is zero. For example, Cheah and Fry (2015) apply statistical tests to Bitcoin's early history and argue that Bitcoin prices constitute a speculative bubble. To the contrary, Biais et al. (2022) argue that Bitcoin should be valued based on its stream of net transactional benefits, including the evasion of government capital controls. Hayes (2019) finds that the price of Bitcoin loosely corresponds to Bitcoin's marginal cost of production.

Indeed, there are several reasons to believe Bitcoin's fundamental value is positive (even if difficult to determine). First, Bitcoin is a resource for which there are no obvious substitutes. Bitcoin is the largest proof-of-work currency by nearly two orders of magnitude. Network effects in the world of cryptocurrency appear to be very strong; efforts to improve Bitcoin by altering aspects of the Bitcoin architecture (such as the 2017 hard fork that spawned "Bitcoin Cash," a version of Bitcoin that can process more on-chain transactions per second) have not seriously rivaled Bitcoin's popularity. Rather than altering the Bitcoin architecture itself, efforts to improve Bitcoin's functionality currently focus on designing methods to transact off-chain, such as the Lightning network, which offers faster transactions and reduced fees.

Source: investing.com

Figure 25: The price returns of reserve assets and Bitcoin, normalized to $1 on February 24, 2022, the day Russia invaded Ukraine. The dashed line indicates 5:00 AM EST on February 28, 2022, when the US Treasury Office of Foreign Assets Control sanctioned the Central Bank of Russia.

Implicitly, these innovative efforts acknowledge the robustness of Bitcoin to usurpation by clones or other rivals.

Second, the total quantity of Bitcoin is capped at approximately 21 million (the total quantity of Ether is not capped, but the proof-of-stake update is expected to result in deflation).[30] After the last Bitcoin is mined in 2140, miners will compete to receive transaction fees for their services, but will not receive any newly minted Bitcoin. Although scarcity does not necessarily imply value, the quantity limit does effectively prevent Bitcoin from

[30]https://blockchain.news/analysis/ethereum-2.0-full-upgrade-will-prompt-a-1-percent-annual-deflation-rate

experiencing devaluation through hyperinflation.

Third, Bitcoin's volatility does not preclude its functioning as a store of value. Gold, a reserve asset widely regarded as a store of value, has experienced significant real losses over prolonged periods of time. Harmson (1998) illustrates that gold tends to maintain its purchasing power over centuries, but the price of gold can experience significant fluctuation over 10- or 20-year periods. Indeed, Figure 26 illustrates that the real price of gold experienced an 86% decline over approximately a 21-year period beginning in 1980. Similarly, since Bitcoin's third "halving" event in July 2016 (which reduced the rate at which Bitcoin is mined), Bitcoin experienced a maximum real drawdown of 83% from December 2017 to December 2018.[31] Superimposing the Bitcoin and gold price graphs reveals several similarities, including the concavity, the pattern of volatility clustering towards the beginning of the series, and the magnitude of the price decline.

[31] Earlier in its history, Bitcoin experienced a 94% nominal drawdown in 2011, declining from $32 to $2. But Bitcoin was a riskier asset in 2011, since it was just two years old and much less widely adopted, making Bitcoin's future prospects in 2011 substantially more uncertain.

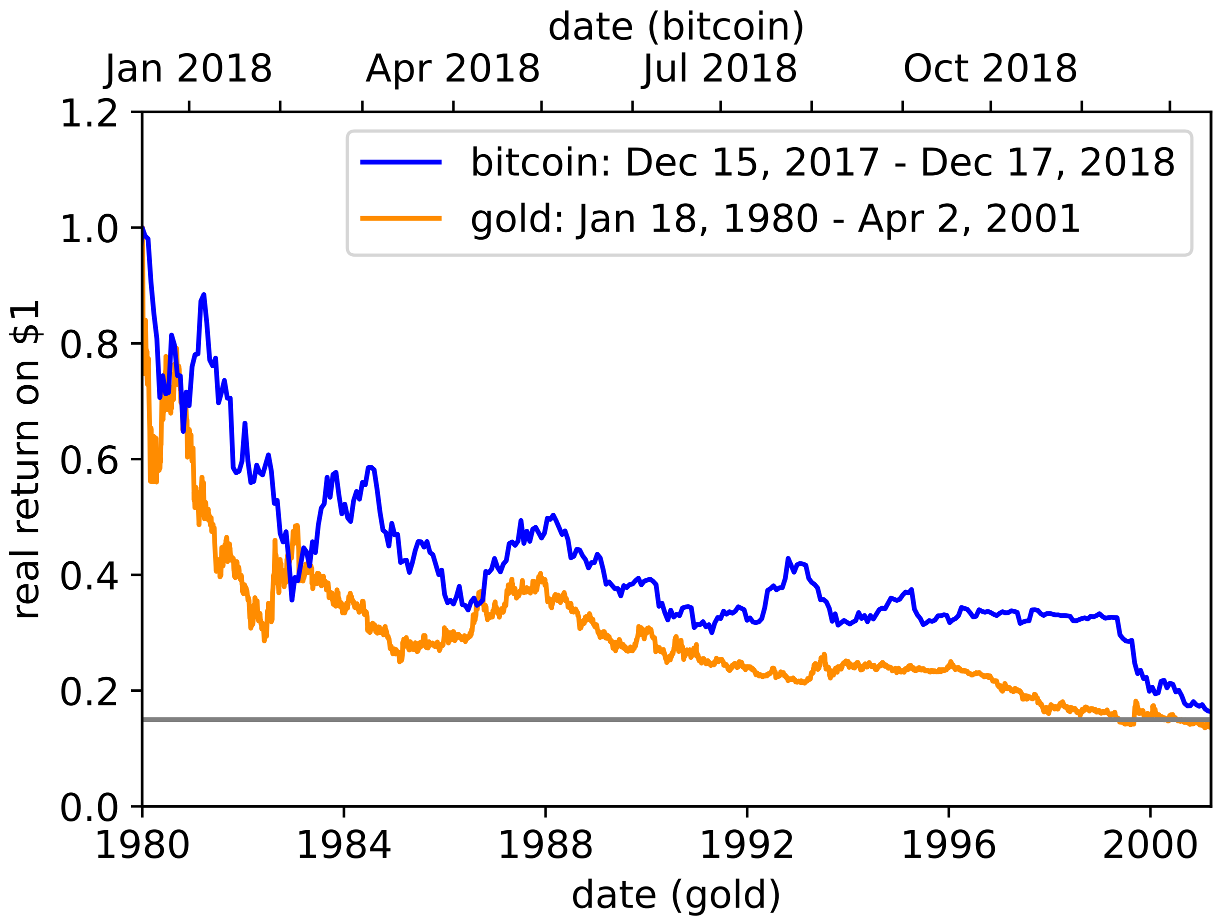

Source: investing.com

Figure 26: The real return of Bitcoin and gold, measured from peak to trough of gold's largest drawdown in the last 50 years, and Bitcoin's largest drawdown since July 2016 (the date of Bitcoin's third "halving" event, which reduces the mining rate of Bitcoin).

C Inverse Student-T Distribution Approximation

The first four terms of the Cornish-Fisher approximation are given in Abramowitz and Stegun (1972). These terms expand the inverse Student-t distribution around the inverse Normal distribution, exploiting the fact that the t-distribution and Normal distribution are "close." The approximation error falls as the degrees of freedom rise, but the error rises in the tails of the Student-t distribution.

$$
\begin{aligned}
T_\lambda^{-1}(x) \approx z & \\
& + \frac{z(z^2+1)}{4\lambda} \\
& + \frac{z(5z^4+16z^2+3)}{96\lambda^2} \\
& + \frac{z(3z^6+19z^4+17z^2-15)}{384\lambda^3} \\
& + \frac{z(79z^8+776z^6+1482z^4-1920z^2-945)}{92160\lambda^4}
\end{aligned}
$$

In order to improve this approximation, I also implement a power series expansion for the tails of the Student-t distribution, where the Cornish-Fisher expansion performs poorly. The tail series performs better when the degrees of freedom are small, making a good pairing with the Cornish-Fisher expansion. Shaw (2006) gives the series expansion, where w is an auxiliary variable and $y_n(w, \lambda)$ is a series of polynomial terms, with n being the order of the expansion. I carry the expansion to six terms.

$$w(x, \lambda) = (1-x)\sqrt{\lambda\pi}\frac{\Gamma(\lambda/2)}{\Gamma((\lambda+1)/2))}$$

$$T_\lambda^{-1}(x) \approx \sqrt{\lambda}(w\sqrt{\lambda})^{-1.0/\lambda} * (1 + y_1(w, \lambda) + y_2(w, \lambda) + \ldots + y_n(w, \lambda))$$

I fit a logistic function, $q = \frac{1}{1+exp(k\lambda)}$ to estimate the crossover quantiles, as a function of the degrees of freedom, where the error in the Cornish-Fisher expansion equals that of the tail series. I find that $k = -1.09080618$.

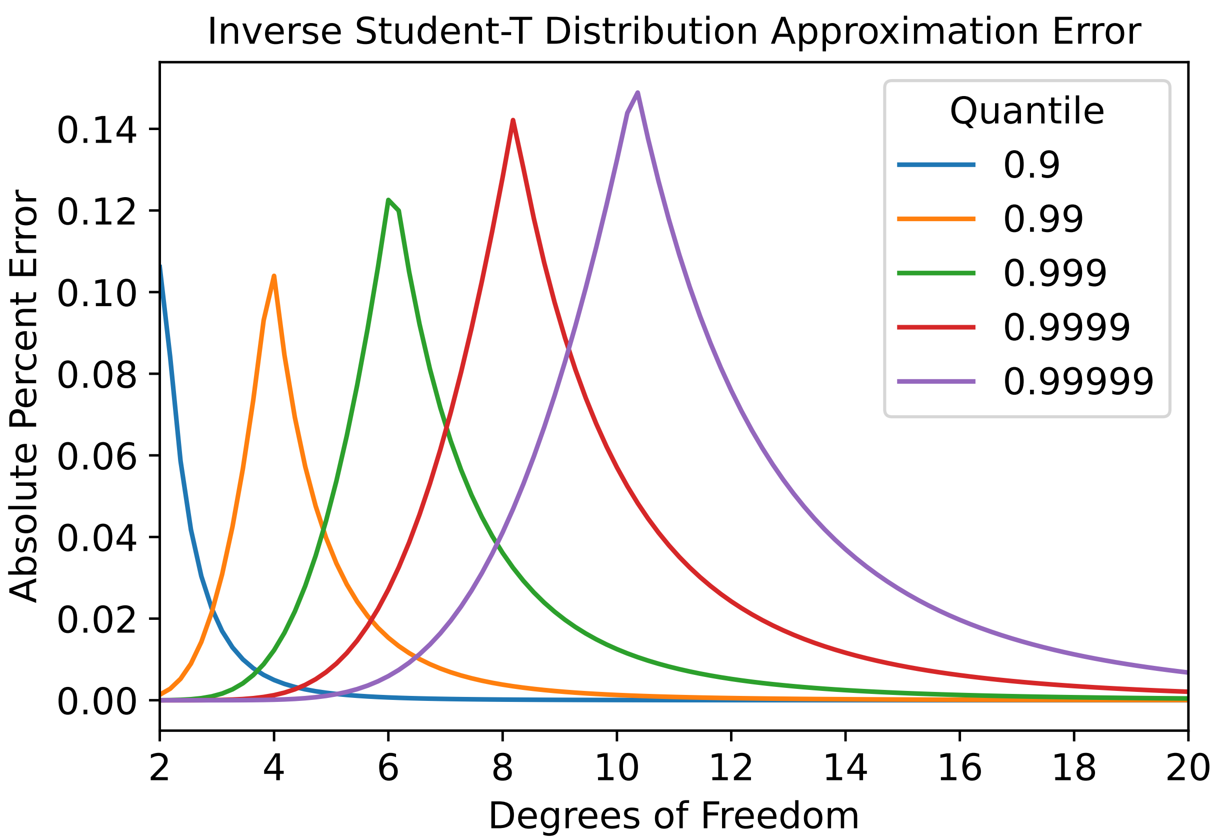

Figure 27: The percent error in the combined inverse Student-t distribution approximation, switching between the Cornish-Fisher and tail series based on the fitted logistic function.