Six Solutions to the Money Supply Problem

Anup Rao
University of Washington
anuprao@cs.washington.edu

September 1, 2025

Abstract

In the money supply problem, the goal is to design a mechanism for controlling the money supply that results in the least distortion of prices. We study 6 solutions to this problem, including 2 novel proposals based on an auction mechanism.

1. The money supply problem

The price of goods and services is determined not just by their own supply and demand, but also by the supply and demand for money. Therefore, the mechanism used to control the money supply can fundamentally alter prices, even when the underlying goods and services remain unchanged. The *money supply problem* is the challenge of developing a mechanism to control the money supply that minimizes unnecessary distortions to prices.

In this article, I explore six approaches to the money supply problem: four standard mechanisms (Section 2) and two new alternatives (Section 3) of my own invention. I discuss how each mechanism interferes with prices, and the trade-offs between them.

2. The four standard solutions

Let us begin by discussing the standard solutions to the money supply problem, starting with the most obvious approaches.

1. Fixed supply The simplest way to control the money supply is to keep it fixed. A fixed money supply, like Bitcoin's ultimate limit of 21 million coins, can cause distortions in prices because the supply of money does not respond to the demand for money. The price distortions are largest when there is a sudden increase in the demand for money.

To illustrate this phenomenon, consider a world with two countries, A and B, separated by a vast ocean. The ocean severely restricts the physical exchange of goods between the two countries, and there is almost no trade. The people of A use Bitcoin, which has reached its supply limit, while the people of B use a government-backed currency. When B's currency becomes worthless due to mismanagement, its government decides to adopt Bitcoin. This decision has significant economic consequences for both A and B.

Since the amount of Bitcoin in B is much less than in A, Bitcoin begins to flow from A to B, and this flow is accompanied by changes to the economic activity in both countries. The imbalance in money supply means that Bitcoin-denominated prices for goods in B are significantly lower than those in A, creating a powerful incentive for new trade. Merchants buy goods in B and transport them across the ocean to sell them for a profit in A. New ships are built. This activity facilitates a flow of Bitcoin from A to B. Over time, this leads to declining prices in A and rising prices in B.

The most immediate and dramatic flow of Bitcoin from A to B is likely to be created by the financial system, since this flow does not rely on physical goods. Anticipating the future trajectory of prices in B, speculators in A amass Bitcoin from investors and use it to acquire real estate and make investments in B. Banks from A that hold Bitcoin are incentivized to open offices in B and offer loans to businesses and individuals. These loans make sense because worthy borrowers in B can expect to acquire more Bitcoin in the future when the supply of Bitcoin in B is larger, and so repay their debt with interest.

This sudden allocation of large quantities of Bitcoin for loans in B means less Bitcoin is available for loans in A, causing interest rates in A to rise, and disrupting prices and economic activity in A. Less Bitcoin becomes available for car loans and mortgages in A, so the revenue of car companies in A suddenly declines, and home construction falters. Car manufacturers shut down some of their factories and lay off workers, and car and home prices drop. This shock spreads through the economy of A, eventually affecting all prices.

In the end, enough Bitcoin will be transferred from A to B for all prices to reach a new equilibrium. One would expect that prices in A will settle at a level lower than they were before Bitcoin was adopted in B, while prices in B will settle at a level higher than they were initially. Car and home sales in A should return, and traffic across the ocean should reset to negligible levels. The unnecessary changes to economic activity before the new equilibrium is achieved are a direct consequence of a fixed money supply; they are caused by the monetary system, not by the physical realities or economic needs of the two countries.

This is not a pathological scenario. With a fixed money supply, any increase in the demand for money in one part of the economy can disrupt transactions in unrelated sectors. For instance, a new technological innovation may lead to new demand for business investment, but a fixed money supply creates a zero-sum constraint. Every Bitcoin allocated to investment immediately reduces the amount available for car loans and mortgages. New transactions in one sector stifle activity in completely independent sectors until prices adjust to accommodate all transactions.

These examples show that under a fixed money supply, changes to the demand for money can lead to strange distortions of prices and unnecessary economic activity. As an aside, a fixed supply also leads to steady deflation as the population grows and the demand for money steadily increases. Historically, deflation has often sparked a vicious deflationary spiral: falling prices lead consumers to delay purchases in anticipation of even lower prices, further depressing prices and reducing transactions. However, it is not entirely clear if this spiral would occur in a world with a known fixed supply, where falling prices have been normalized, and it is not clear that this kind of deflation is a barrier for transactions in such an economy.

2. Gold From the perspective of the money supply problem, gold coins have a feature that Bitcoin does not. Because more gold can be mined, the money supply increases in response to demand. This provides an escape valve for the pressure created by increases in the demand for money.

Let us revisit our example with gold coins. If country A uses gold coins and B adopts them,

the supply in B will be smaller, and prices will be lower. This leads to the same initial effects as with a fixed supply, but they can be dampened by increased gold mining. Because of the increased demand for gold in B, the profitability of mining rises, leading to more gold being mined in both countries.

Just as in the case of a fixed supply, banks may allocate fewer gold coins for car loans so they can fund loans to people in B. However, instead of shutting down some of their car factories, car companies that see a sudden drop in revenue from car sales may pivot some factories to manufacturing equipment for mining gold. The resulting gold can be used to provide some of the missing financing for car loans and mitigate the negative effects of the fixed money supply. Thus, the supply of gold in B increases via a flow of gold from A, as well as via increased mining activity, until the supply of gold in B achieves parity with the supply available in A.

It is difficult for the production of gold to keep pace with sudden changes in the demand for new money, because mining gold takes time and effort. For these reasons, the use of gold coins still leads to significant disruptions in prices for some period of time. However, all things being equal, one would expect faster convergence to new equilibrium prices, because the market has the freedom to increase the money supply if doing so is profitable.

3. Currency backed by gold Instead of gold coins, governments in the last century have issued money in the form of notes that can be redeemed for gold. Notes are certainly easier to use than gold coins, and that is their obvious advantage. But notes also have a subtle advantage: they can be printed *before* the gold backing them is acquired. The money supply now has the additional freedom to expand before new gold is mined. In times when the demand for money increases, governments and banks can issue new notes expecting that the gold required to back the money can later be acquired from the economy via taxation or interest payments.

In our running example, if country A is using notes backed by gold and B suddenly decides to adopt the same notes, banks from A can issue loans to people in B using newly printed notes, even before the new gold is mined. Banks can expect the loans to be repaid with interest because the increased demand for gold means that additional gold will be mined in the future to satisfy the new demand. This instant creation of money minimizes the immediate disruption to prices and activity in A. Prices in A and B can converge to their new equilibrium at the speed of the printing press and the banking system.

Unfortunately, history has shown that this approach is too good to be true. Setting aside the potential for corruption that this system invites, it is very difficult for banks to compute the correct number of new notes, and make good decisions about who should acquire them. If the banks play it safe and issue fewer than the correct number of notes, then we face the same distortions to the economies of A and B as with a fixed money supply. If the banks issue too many notes in the form of new loans, there are dramatic consequences. If loans default and the required gold is not acquired, the banks can no longer redeem all outstanding notes for gold, and we have the beginnings of a liquidity crisis. When people suspect that the banks do not have the gold to back the notes, there is a run on banks and hoarding of gold, which worsens the crisis. Such a crisis is ultimately resolved by devaluing the notes, which leads to dramatic changes in economic activity. So, the banks are forced to perform a delicate balancing act, with negative consequences at either extreme.

The historical record bears witness to the difficulty of maintaining this balance. The United States experienced at least five major banking panics between 1880 and 1933 under the gold standard. Countries have repeatedly been forced to abandon gold backing entirely, as the U.S. did in 1971.

4. Fiat money The fiat money system, adopted by the U.S. dollar since 1971, allows the government to print money when it is needed. Under this system, banks issue deposits backed by reserve currency that can be printed by the government. The government promises to protect banks and their customers by printing reserve currency when required.

As with notes backed by gold, fiat money can be instantly generated to satisfy new demand for money. In the scenario described above, if country A is using a fiat money supply and country B suddenly adopts A's currency, then new money can be generated to satisfy the increased demand. With the approval of the government in A, banks from A can issue loans to people in B at a massive scale without disrupting loans in A. In this way, the money supply in B can quickly reach parity with the money supply in A, and the harmful effects of money supply imbalances can be minimized.

The problem with fiat money supply is that it can lead to an unwarranted and misdirected expansion of the money supply. A misguided bank may look at the initial price differences between A and B and decide that companies that build ships to cross the ocean deserve significant long-term funding. This sudden allocation of newly created money leads to an explosion of activity building new ships that are ultimately destined to be salvaged for scrap metal when the two economies eventually reach equilibrium prices. Although this system avoids runs on the bank, it can lead to a misallocation of resources, and a corresponding disruption of prices.

Under the fiat system, knowing that the government can print the required money, the parties involved in expanding the money supply have little incentive to restrain themselves. During the 2008 financial crisis, banks irresponsibly issued mortgages for overpriced houses that fueled a housing bubble. None of the participants—bankers, real estate agents, buyers, or sellers—had incentives to refuse these loans because they were fueled by newly created money. When the bubble finally burst, the economy was left with a surplus of new houses.

The fiat system concentrates the power to create money in the hands of banks and government officials. The role of the market in creating money is reduced. The privileged few with access to the levers of the money printer are likely to make mistakes and certain to make profits.

3. Solutions based on auctions

An ideal mechanism should respond to new demand for money quickly, to minimize the distortions of a fixed supply. At the same time, the mechanism should be subject to market forces, to avoid the pitfalls of fiat money. I propose an auction-based mechanism with these features. A similar auction was first explored in [1]. In [2], I discussed how banks and central banks can incorporate the mechanism to create a novel decentralized currency for international trade. Here, I will explain how the auctions lead to two solutions for the money supply problem.

Bond auction mechanism The mechanism is based on a financial instrument called $\mathsf{Bond}(r)$. This object always costs \$1, expires 1 day after it is created, and returns $\$(1 + r)$ to the holder upon expiration. $\mathsf{Bond}(r)$ can be thought of as a little money printer. You can buy one for \$1, and after 1 day, it returns your dollar along with $\$r$ of newly printed money. The money supply is controlled by users who bid for these bonds in a daily auction. New bonds are allocated to users daily according to the following rule:

- **Bond rule:** If the total money supply is $\$T$, then $T/2$ bonds are issued using an auction. Users participate by submitting (possibly multiple) bids of the form:

  user U bids for $\mathsf{Bond}(r)$,

indicating that user U is willing to buy $\mathsf{Bond}(r)$ for \$1, with the yield r chosen by the user. The $T/2$ lowest yield bids win the daily auction.

The daily rate at which the money supply grows is determined by the yields of the winning bids in the auction. It is the $T/2$ users that are willing to accept the lowest yields that obtain the newly created money. Because r can be arbitrarily large, all $T/2$ bonds will always sell. We note that a practical implementation of the auction need not involve daily bidding by users. Users can instruct their banks to automatically submit standing orders specifying their desired allocation towards the bids. Banks can aggregate this information for the central bank, which can execute the auctions and distribute newly created money to the winning bidders' accounts.

There are two ways that the bond auction mechanism can be used to regulate the money supply. The first is that the money generated by the mechanism can be used directly. This corresponds to a synthetic version of gold coins, except that new coins are synthetically mined according to the rules of the bond auction given above. Alternatively, banks can issue notes backed by the money generated in the bond auctions. This corresponds to using notes backed by gold. Let us discuss each of these arrangements.

Bond money The bond yields in the auctions represent the market interest rate for a risk-free loan, and these interest rates determine the rate at which new money is created. The issued bonds will certainly sell, because users can place bids with arbitrarily large r. The risk-free interest rate is higher when money is more valuable now than in the future. Thus, the creation of money accelerates exactly when there is an increase in the demand for money in the present.

Let us work through our running example using this new mechanism. Imagine that country A is using dollars generated by the bond auctions when country B abandons its own currency and adopts the currency of A. As before, the immediate effect is that prices in B will be much lower than in A, and money will begin to flow from A to B using a variety of channels. The bids in the bond rule immediately begin to react to the increased demand for money.

When banks from A begin to make loans in B, they face a choice. Each dollar can be allocated to a loan in A, a loan in B, or a bond generated by the bond rule. As demand for loans in B rises, interest rates for loans in A rise, and yields in the bond rule also rise. This rise in bond yields corresponds to the immediate creation of new money that is delivered to holders of bonds. The sudden allocation of money for loans to people in B certainly reduces the money allocated for loans in A, but the effects are mitigated by the newly generated money. Savers obtain these higher yields, which can be routed to borrowers in A and B by the market. Therefore, the increased demand for money leads to the creation of new money to satisfy that demand.

Unlike with gold coins, the rate at which money is created is unlimited and entirely determined by the market. Anticipating the future trajectory of prices, speculators amass money to acquire assetsand make investments in B. This activity can be funded by large quantities of new money generated by the bonds. Speculators sell investment funds to holders of money in A and use the money raised to fund their activities in B. As some savers are drawn to these investments, the savers that continue to bid in the auctions of the bond rule can expect to obtain larger yields, which can be invested in B soon after.

There are no physical limits on the rate at which money is created by the bond rule, and the rate is controlled by the market, so the system finds new equilibrium prices at the speed of market forces. We should expect this to happen much faster than the system based on gold coins.

Currency backed by bond money The money generated by the bond auctions is generated by the action of the market, but there may still be some lag before the market makes the correct allocations. History has shown that if the lag is significant, banks are likely to begin issuing notes that can be redeemed for bond money to close the lag. Just as with notes backed by gold, such notes can be issued *before* the bond money that backs them is generated.

In our running example, if country A uses currency backed by bond money and B abandons its own currency to adopt A's currency, the situation is similar to the use of currency backed by gold. Banks can issue the money required in B instantly, but they run the risk of a liquidity crisis if they issue too much money. In [2], I discussed how central banks can plan for this situation by amassing enough bond money to play a significant role in the bond auction and so force large bond yields in times of crisis. The money so obtained can be routed to banks in the event of a liquidity crisis. The bond auctions given in [2] are a little more complicated than the auctions defined here to accomodate this application.

Unlike with the use of fiat currency, no particular group controls the printing press of bond money. In particular, speculators that anticipate a liquidity crisis can hold onto their bond money and profit from the large yields that emerge when the liquidity crisis is resolved by the central bank. Thus, bad investments made by banks are a source of profit for those that anticipate their failure. Unlike with the use of fiat money, institutions involved in such bad decisions will see a relative loss of their holdings of money, because all holders of money can always keep pace with the growth of the money supply simply by participating in the auctions for bonds.

To summarize, currency backed by bond money can generate new money at the same rate as a fiat system, but unlike with the fiat system, banks and governments that make bad decisions using the money do face consequences in the market.

4. Other money

Sections 2 and 3 contain all ways to control the money supply that are known to me. Currencies sometimes switch between mechanisms. For example, the U.S. dollar has been backed by gold in the past, and is now a fiat currency.

Although Bitcoin currently generates new coins through mining, the money supply of Bitcoin is unaffected by the demand for money, so from the perspective of the money supply problem, it has a fixed supply. Ethereum currently regulates the money supply by decreasing supply in response to increases in transactions, so the supply does respond (mildly) to the demand for money, but in the wrong direction. Therefore, it also generates the same distortions as a fixed supply mechanism. If Bitcoin were to abandon the hard cap and issue coins to miners proportionally to electricity consumed, it would become similar to gold coins from the perspective of the money supply problem.

Stablecoins pegged to the dollar, while ostensibly distinct from the dollar, are the same as fiat currencies from the perspective of the money supply problem. That is because newly created dollars can be converted to stablecoins through the peg. Cryptocurrencies like DAI, which maintain pegs to other instruments, inherit the properties of the mechanisms used to control the supply of those instruments.

5. Conclusions and Open questions

In this article, we defined the money supply problem and investigated six approaches to resolving it. We reasoned about these various approaches using a running example. It would be interesting to find the definitions that can be used to reason about this problem rigorously and in greater generality. Albeit for different purposes, Bitcoin specifies a way for money to be mined using computational effort. It would be interesting to find other solutions to the money supply problem that rely on creative mechanism using such concepts.

References

[1] Rao, Anup. 2023. Elastic Cash. arXiv:2301.04244 [cs.CR]. https://arxiv.org/abs/2301.04244.

[2] Rao, Anup. (2025, June 25). Decentralized money supply: A new paradigm for reserve currencies. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5287712