UNIVERSITÄT
BERN

Faculty of Business, Economics
and Social Sciences
Department of Economics

The Monetary Dynamics of Hyperinflation Reconsidered

Luca Benati

23-05

March, 2023

DISCUSSION PAPERS

Schanzeneckstrasse 1
CH-3012 Bern, Switzerland
http://www.vwi.unibe.ch

The Monetary Dynamics of Hyperinflation Reconsidered*

Luca Benati
University of Bern†

Abstract

Data from 20 hyperinflations provide *no evidence* of a Laffer curve for seignorage: rather, the relationship between money growth and seignorage has been uniformly positive at all inflation rates. Consistent with this, evidence shows that the most plausible money demand specification for hyperinflations is not Cagan's 'semi-log', which imposes a Laffer curve upon the data, but rather either the 'log-log', or a functional form close to it such as Benati, Lucas, Nicolini, and Weber's (2021), which produce a monotonically increasing relationship between money growth and seignorage. Compared to the semi-log, functional forms close to the log-log imply different properties for theoretical models of hyperinflations along two dimensions: (*i*) the equilibria's stability properties are reversed, with the high-inflation equilibrium being unstable under rational expectations; and (*ii*) there is the possibility of explosive inflation even in the presence of well-defined steady-states. I discuss the implications of this for the interpretation of specific historical episodes. Under the Weimar Republic, a plausible interpretation of the macroeconomic consequences of the invasion of the Ruhr is that it pushed the economy beyond the unstable steady-state, into the region of explosive inflation.

*I wish to thank Pierpaolo Benigno, Jesus Fernandez-Villaverde, Ed Herbst, Dirk Krueger, Francisco Ruge-Murcia, and Harald Uhlig for useful discussions, suggestions, and/or comments. Special thanks to Eugene White, Tara McIndoe-Calder, Liuyan Zhao, Carlos Gustavo Machicado, and Zorica Mladenovic and Bent Nielsen, for kindly providing data for the French Revolution's, Zimbabwe's, China's, Bolivia's and Yugoslavia's hyperinflations, respectively. Usual disclaimers apply.

†Department of Economics, University of Bern, Schanzeneckstrasse 1, CH-3001, Bern, Switzerland. Email: luca.benati@vwi.unibe.ch

1 Introduction

Since the classic study of Phillip Cagan (1956), the literature on hyperinflation has been consistently dominated by a single narrative, featuring two main elements:

(I) the relationship between money growth and seignorage is *hump-shaped*—i.e., it exhibits a *Laffer-curve* property—and

(II) historically, governments have near-uniformly *inflated in excess of the revenue-maximizing rate*, i.e., they have been on the 'wrong side' of the Laffer curve.

This view, first introduced by Cagan (1956), can be found in all of the most prominent subsequent papers in the literature, from Sargent and Wallace (1973, 1987), to Sargent (1977), and Salemi and Sargent (1979). The dominance of this view is further testified by the fact that it can be found in all of standard graduate textbooks' treatments of hyperinflation, from (e.g.) Blanchard and Fischer (1990, Chapter 4, pp. 195-201), to Obsteld and Rogoff (1996, Chapter 8, pp. 515-530), to Walsh (2017, Chapter 4, pp. 153-162).

In spite of its ubiquity and its dominance, this narrative suffers from a crucial problem: as I show, it is sharply at odds with the data.

In this paper I revisit the monetary dynamics of hyperinflations based on data from 20 episodes, from the French Revolution to Venezuela under Nicolas Maduro. Evidence of a Laffer curve for seignorage is virtually *non-existent*: rather, in nearly all cases the relationship between money growth and seignorage had, and has been positive at all inflation rates.

Although, in principle, several alternative rationalizations for the absence of a Laffer curve for seignorage could be possible, I argue that in fact the explanation is straightforward. Following Cagan (1956), the literature on hyperinflations has been consistently dominated by the *semi-log* money demand specification, relating the *logarithm* of real money balances to a measure of the *level* of an opportunity cost of money (within the present context, expected inflation). As it is well known,[1] a Laffer curve for seignorage is a mathematical property of the semi-log specification that is independent of the specific value taken by the semi-elasticity of money demand. As a result, when a researcher estimates a semi-log (s)he is not 'discovering' a Laffer curve: (s)he is *imposing* it upon the data. It should therefore come as no surprise that the notion of a Laffer curve for seignorage has been part of macroeconomists' conventional wisdom for the last six decades. As my evidence shows, however, such a notion is incorrect.

In fact, evidence clearly suggests that the most plausible money demand specification for hyperinflations is not Cagan's (1956) semi-log, but rather either Allan Meltzer's (1963) 'log-log', or a functional form close to it such as Benati, Lucas, Nicolini, and Weber's (2021) benchmark specification, which for empirically plausible values of the structural parameters produce a *monotonically increasing* relationship between money growth and seignorage.

[1] See e.g. Lucas (2000).

In several cases the data's preference for the log-log (or a specification close to it) is so strong that it is apparent even to the naked eye. Stark illustrations of this are provided by Yugoslavia, Zimbabwe, Germany, post-WWII Hungary and Greece, and, to a slightly lesser extent, China: for all these episodes, the *logarithm* of inflation tracks log real money balances remarkably closely,[2] whereas its *level* exhibits a scant connection to it.

Overall, evidence therefore suggests that both the finding of a Laffer curve for seignorage, and Cagan's paradox of policymakers inflating in excess of the revenue-maximizing rate during hyperinflations, are the product of the literature's focus on the semi-log functional form, which automatically imposes a Laffer curve upon the data.

I show that, compared to the semi-log, functional forms close to the log-log imply materially different properties for theoretical models of hyperinflations along two dimensions:

(i) the equilibria's stability properties are reversed, with the high-inflation equilibrium being unstable under rational expectations; and

(ii) as a consequence of (i) there is the possibility of explosive inflation even in the presence of well-defined steady-states—which is instead not possible with the semi-log—if a shock pushes the economy to the right of the high-inflation equilibrium, into the region of 'runaway inflation'.

I discuss the implications of (i) and (ii) for the interpretation of specific historical episodes. In particular, under the Weimar Republic inflation had been very high, but not explosive, until the end of 1922. The occupation of the Ruhr on the part of France in January 1923, however, caused a dramatic increase in the German government's budget deficit. A plausible interpretation of the explosive inflationary dynamics that followed the occupation of the Ruhr is therefore that the dramatic and sudden deterioration of Germany's fiscal position pushed the economy to the right of the unstable high-inflation steady-state, into the region of explosive inflation.

The paper is organized as follows. The next section discusses the relationship between money growth, inflation, and seignorage within three alternative models of the demand for real money balances, proposed respectively by Cagan (1956), Meltzer (1963), and Benati et al. (2021). Section 3 briefly discusses the previous literature. Section 4 discusses the main features of the 20 hyperinflation episodes in terms of duration, and of median and maximum inflation, whereas the data and their sources are discussed in detail in Appendix A. Section 5 explores the empirical relationship between money growth and seignorage based on either simple visual inspection of the raw data, or panel regressions with country-specific fixed effects. Evidence clearly suggests that across all of the 20 episodes the two series have been near-uniformly positively correlated. Section 6 shows that both for the comparatively short-lived hyperinflation episodes, and for longer samples characterized by high, or very high inflation rates, the most plausible functional form for the demand for real money

[2] Quite obviously, once appropriately rescaled.

balances is either the log-log, or a specification close to it. On the other hand, Cagan's semi-log specification clearly appears to be at odds with the data. Section 7 presents evidence from panel VARs for the logarithms of inflation and real money balances, whereas Section 8, working within the same framework, estimates the elasticity of money demand. Section 9 discusses the implications of these findings for the interpretation of specific historical episodes, focusing in particular on Weimar's hyperinflation. Section 10 concludes, and outlines three directions for future research.

2 Money Growth and Seignorage Within Alternative Models of Money Demand

By defining the money stock, real GDP, and the price level as M_t, Y_t, and P_t, the instantaneous revenue from money creation—i.e. seignorage—expressed as a fraction of GDP is defined as

$$\xi_t \equiv \frac{dM_t}{dt}\frac{1}{P_tY_t} = \underbrace{\left(\frac{dM_t}{dt}\frac{1}{M_t}\right)}_{\theta_t}\frac{M_t}{P_tY_t} = \theta_t\frac{M_t}{P_tY_t} \tag{1}$$

where θ_t is instantaneous money growth, and $M_t/(P_tY_t)$ is the demand for real money balances expressed as a fraction of GDP.

I start by comparing the relationship between money growth and seignorage within the two models of the demand for real money balances that have dominated the post-WWII literature, i.e. Cagan's (1956) and Meltzer's (1963). Since a limitation of these models is that neither of them had been derived from first principles, I then consider the micro-founded model proposed by Benati et al. (2021), which produces a functional form for the demand for real money balances very close to the log-log.

As the opportunity cost of money I take a nominal instantaneous interest rate, $R_t = r^* + \pi_t^e$, where r^* and π_t^e are respectively the real interest rate, which without loss of generality I assume to be constant, and expected inflation. Since for hyperinflations r^* is negligible compared to π_t^e, in what follows I ignore it, and I set $R_t = \pi_t^e$. I also assume real GDP to be constant. Again, the rationale is that its fluctuations are negligible compared to fluctuations in money growth and inflation.

2.1 Cagan (1956) versus Meltzer (1963)

The demand for real money balances as a fraction of GDP is described by either Cagan's (1956) semi-log functional form,

$$\ln\left(\frac{M_t}{P_tY_t}\right) = \beta - \alpha\pi_t^e, \tag{2}$$

or Meltzer's (1963) log-log,

$$\ln\left(\frac{M_t}{P_t Y_t}\right) = \phi - \gamma \ln(\pi_t^e) \tag{3}$$

where α, β, γ and ϕ are positive constants. Expected inflation evolves according to

$$\frac{d\pi_t^e}{dt} = \delta(\pi_t^e - \pi_t) \tag{4}$$

with δ>0. In what follows I focus on the case of perfect foresight, which is obtained for $\delta \to \infty$.

From (1) and (2) Bruno and Fischer (1990) show that for the semi-log specification the first equilibrium condition of the model (in Bruno and Fischer's notation, the 'GG curve') is given by

$$\ln \xi = \ln \theta_t - \alpha \pi_t^e \tag{5}$$

By the same token, from (1) and (3) it can be shown that for the log-log the GG curve is given by

$$\ln \xi = \ln \theta_t - \gamma \ln \pi_t^e \tag{6}$$

Since (1) represents the government's budget constraint, which ought to be satisfied at each point in time, the economy always ought to be on the GG curve. Within this context ξ, the amount of seignorage the government needs to raise, is assumed to be exogenously given, and acts as a shifter for either GG curve.

Finally, by taking time-derivatives of either (2) or (3) and setting them to zero, we obtain that for either functional form in equilibrium (i.e., when $d\pi_t^e/dt = d\pi_t/dt = d\theta_t/dt = 0$)

$$\pi_t = \pi_t^e = \theta_t \tag{7}$$

The equilibria of the economy lie at the intersection of the '45 degree line' (7), and of either of the two GG curves, i.e. either (5) or (6).

2.1.1 Money growth and seignorage in the steady-state

With semi-log money demand steady-state seignorage is given by $\xi = \theta \exp[\beta - \alpha\theta]$, which traces out a Laffer curve as a function of θ. Specifically, (i) for $\theta = 0$, $\xi = 0$; (ii) for $\theta = 0$, $d\xi/d\theta = \exp(\beta)$>0, so that starting from a steady-state with $\theta = \xi = 0$, an increase in θ generates a positive amount of seignorage; (iii) ξ reaches a maximum corresponding to $\theta = \alpha^{-1}$; and (iv) for $\theta \to +\infty$, $\xi \to 0$.

With log-log money demand, on the other hand, steady-state seignorage is $\xi = \exp(\phi)\theta^{1-\gamma}$, which implies that, as long as $0 < \gamma < 1$, seignorage increases monotonically with θ. Specifically, (i) for $\theta = 0$, $\xi = 0$; (ii) for $\theta \to 0$, $d\xi/d\theta \to +\infty$; and (iii) for $\theta \to +\infty$, $\xi \to +\infty$.

In fact, as I will discuss in Section 8, panel estimation of γ for all of the 20 episodes considered jointly produce estimates that are either virtually identical or very close

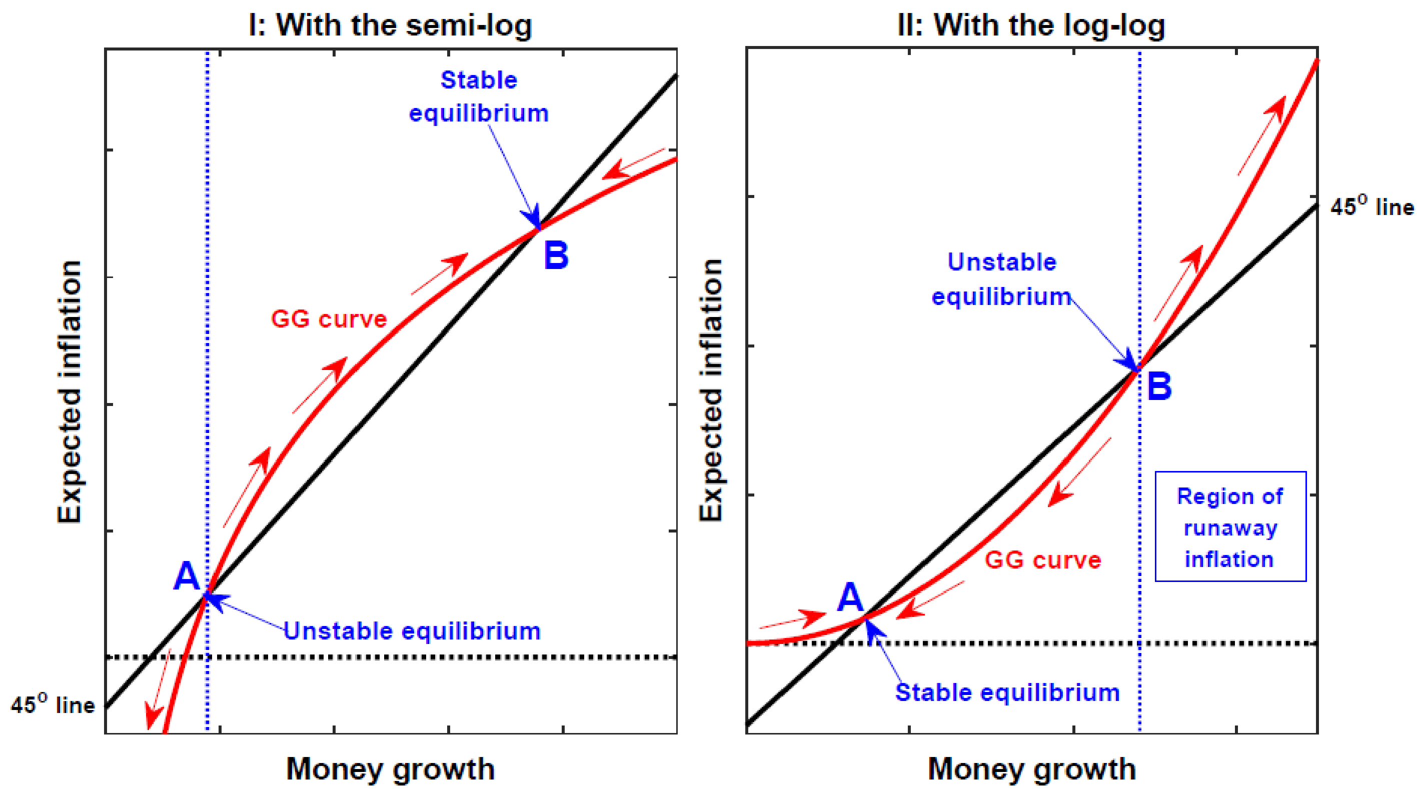

Figure 1 Money growth and seignorage for Cagan's (1956) semi-log and Meltzer's (1963) log-log specifications for the demand for real money balances

to Baumol and Tobin's theoretical value of 1/2. Further, the fraction of bootstrap replications for which γ is estimated to be greater than 1 is consistently equal to 0. Splitting the overall sample into the 10 episodes with the lowest and respectively the highest median inflation, on the other hand, produces clear evidence that for the latter group γ has consistently been higher than for the former. Crucially, however, even for the 10 episodes with the highest median inflation there is little to no evidence that γ may have been greater than 1.

Overall, a plausible interpretation of the evidence is that γ is indeed an increasing function of θ (and π). Therefore, if money growth crosses a certain threshold θ^*, we can safely expect that γ will become greater than 1, with the implication that for all $\theta > \theta^*$ seigniorage will be a *decreasing* function of money growth (i.e., we will obtain a Laffer curve even with the log-log). The crucial point is that, historically, the threshold θ^* appears never to have been crossed. I now turn to a discussion of the dynamical properties of the system under the two money demand specifications.

2.1.2 The dynamical properties of the system

The left-hand side panel of Figure 1 shows the equilibria for the semi-log specification, whereas the right-hand side panel shows the corresponding equilibria for the log-log. Although either case features two equilibria, the dynamical properties of the system under the two money demand specificatons are very different. It can be trivially shown that for the semi-log specification A is unstable and B is stable, which is the classic paradox first highlighted by Cagan (1956): the government could raise the same amount of revenue at a lower inflation rate in equilibrium A. However, for the log-log specification the opposite is true: A is stable, and B is unstable. A key implication is that if a shock pushes the economy beyond B we enter a region of 'runaway inflation', and inflation explodes without limits. With the semi-log on the other hand this is not possible, because the high-inflation equilibrium B is stable.

Discussion In several cases inflation quite clearly exhibits an explosive dynamics towards the end of the hyperinflation. This is the case, e.g., for the seven episodes reported in Figure 2. With the semi-log functional form, the *only* way for the model to generate explosive inflation is for the government to attempt to collect an amount of seignorage greater than the maximum feasible (i.e., that associated with the peak of the Laffer curve). Under these circumstances, the GG curve and the 45 degrees line in the left hand-side panel of Figure 1 do not touch, so that there is no steady-state to speak of, and as a result inflation simply explodes to infinity.

With the log-log functional form, on the other hand, explosive inflation can result even if in fact *any* amount of seignorage is feasible, in the sense that it can be collected in a steady-state. Consider the economy described by the right hand-side panel of Figure 1: if A and B are sufficiently close, a large shock may cause the economy to jump to a point on the GG curve to the right of B. In Section 7 I will discuss a

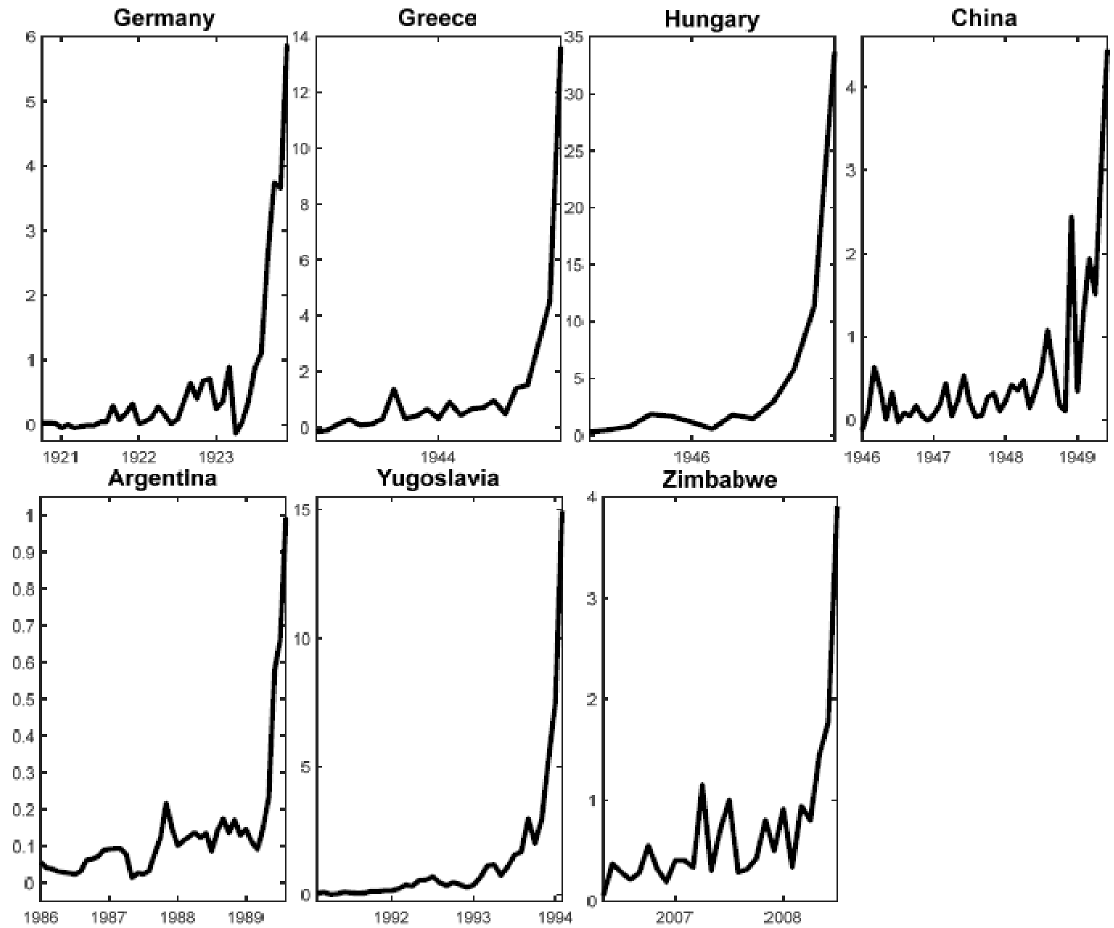

Figure 2 Selected evidence on explosive inflation in the latest stages of hyperinflations

plausible historical example of such a shock, i.e. the invasion of the Ruhr on the part of France in January 1923, and its devastating impact on Germany's finances.

As mentioned, a limitation of either Cagan's (1956) or Meltzer's (1963) model is that they had not been derived from first principles. I therefore now turn to a micro-founded model of the transaction demand for money.

2.2 Benati, Lucas, Nicolini, and Weber (2021)

Appendix A describes in detail the model of the transaction demand for money proposed by Benati et al. (2021). The model generalizes the framework proposed by Baumol (1952) and Tobin (1956) by allowing for several alternative functional forms for the cost of making transactions as a function of the number of 'trips to the bank', n_t. Notice that within this framework n_t is the velocity of money, i.e. the ratio between nominal GDP and nominal money balances, and its inverse is therefore the demand for money balances as a fraction of GDP:

$$\frac{1}{n_t} = \frac{M_t}{P_t Y_t}. \tag{8}$$

Whereas Baumol and Tobin assumed that the cost of making transactions increases linearly with n_t, Benati et al.'s (2021) benchmark functional form is given by

$$\theta(n_t) = \psi n_t^{\sigma} \tag{9}$$

where ψ and σ are positive constants (with $\sigma = 1$ we obtain the Baumol-Tobin case). Notice that ψn_t^{σ} is the welfare cost of inflation expressed as a fraction of maximum potential output.

When the cost of making transactions is described by (9), the solution for n_t is

$$\sigma\psi \frac{n_t^{\sigma+1}}{1 - \psi n_t^{\sigma}} = R_t = \pi_t^e \tag{10}$$

This expression implicitly defines a solution for n_t as a function of π_t^e. As discussed by Benati et al. (2021, p. 46), at low inflation, and therefore low interest rates the welfare costs ψn_t^{σ} are negligible, so that $1 - \psi n_t^{\sigma} \simeq 1$, and the solution (10) becomes $\sigma\psi n_t^{\sigma+1} \simeq \pi_t^e$. Then, taking logarithms we obtain

$$\ln \frac{1}{n_t} = \ln \frac{M_t}{P_t Y_t} \simeq \frac{1}{\sigma + 1} \left[\ln \sigma\psi - \ln \pi_t^e\right] \tag{11}$$

which is Meltzer's log-log specification, with elasticity $1/(\sigma + 1)$.

For the present purposes, a crucial point to stress is that although expression (11) provides a good approximation to the exact solution (10) only at low inflation levels, in fact expression (10) exhibits features very close to those of the log-log at *all* inflation levels. Figure 3 provides simple evidence on this for Bolivia, which experienced a

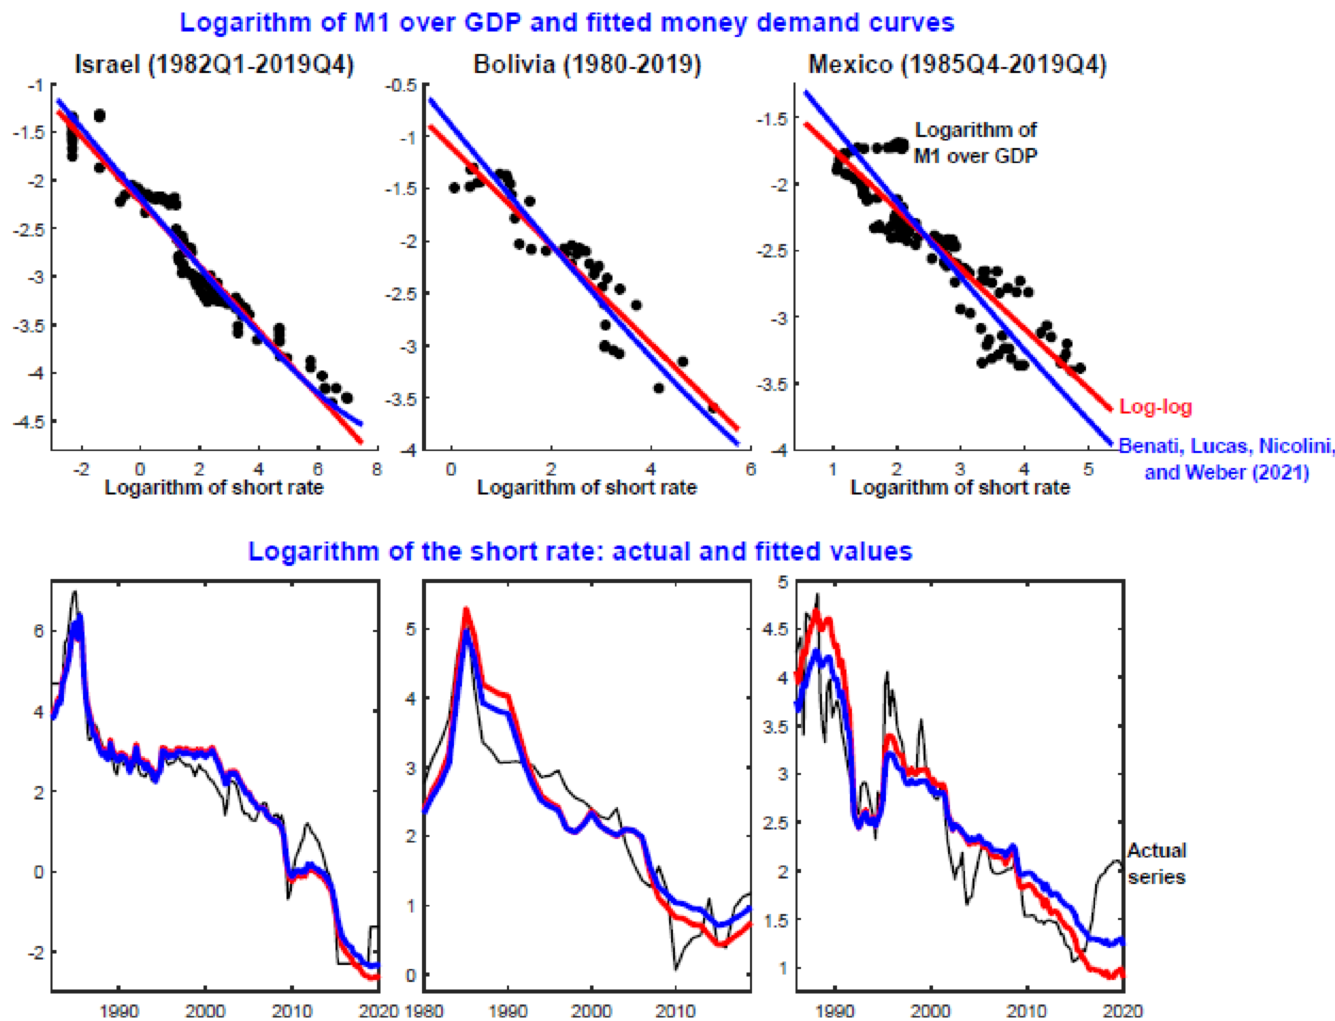

Figure 3 Evidence on the similarity between Meltzer's (1963) log-log and Benati, Lucas, Nicolini and Weber's (2021) money demand specification

hyperinflation in the mid-1980s, and for Israel and Mexico, which although never experienced hyperinflations went through episodes of very high inflation in the early 1980s. The top row shows scatterplots of the logarithms of a short-term nominal interest rate and the ratio between nominal M1 and nominal GDP, together with estimates of log-log and Benati et al.'s (2021) demand curves for real money balances.[3] The bottom row shows the actual logarithm of the short rate together with the values predicted based on either of the two estimated money demand curves. The overall message from Figure 3 is that for any of the three countries the estimates of the two functional forms are very close.

In order to replicate the fall in real money balances associated with increases in the inflation rate, and therefore in expected inflation, in expression (10) it ought to be the case that $1 - \psi n_t^{\sigma} > 0$, which implies that the welfare cost of inflation expressed as a fraction of maximum potential output ought to be smaller than one. Since all of the 20 hyperinflations analyzed herein have been characterized by dramatic collapses in real money balances, in what follows I assume that this condition is satisfied.

By combining expressions (1) and (8), log seignorage is given by

$$\ln \xi_t = \ln \theta_t - \ln n_t \tag{12}$$

Taking logarithms of (10), and then taking derivatives with respect to time, we obtain

$$\frac{d \ln \pi_t^e}{dt} = \frac{(1+\sigma) - \psi n_t^{\sigma}}{1 - \psi n_t^{\sigma}} \left[\frac{d \ln n_t}{dt} \right] \tag{13}$$

By the same token, taking logarithms of (8), and then taking derivatives with respect to time, we obtain

$$\frac{d \ln n_t}{dt} = \pi_t - \theta_t \tag{14}$$

Combining (13) and (14) we obtain

$$\frac{d \ln \pi_t^e}{dt} = \underbrace{\frac{(1+\sigma) - \psi n_t^{\sigma}}{1 - \psi n_t^{\sigma}}}_{\Psi(n_t)} [\pi_t - \theta_t] = \Psi(n_t) [\pi_t - \theta_t] \tag{15}$$

As previously discussed, the empirically relevant case is $1 - \psi n_t^{\sigma} = \Psi(n_t) > 0$. The solution for money growth, money velocity (and therefore its inverse, the demand for real money balances as a fraction of GDP), inflation, and seignorage is fully characterized by equations (10), (12), and (15).

[3] The log-log demand curves have been estimated via a simple OLS regression of log M1 velocity on the logarithm of the short rate. Since Elliot et al.'s (1996) unit root tests strongly suggests that both series are I(1) for all countries, this regression is in fact a cointegrating regression. Benati et al.'s demand curve, on the other hand, has been obtained by estimating (10) via non-linear least squares.

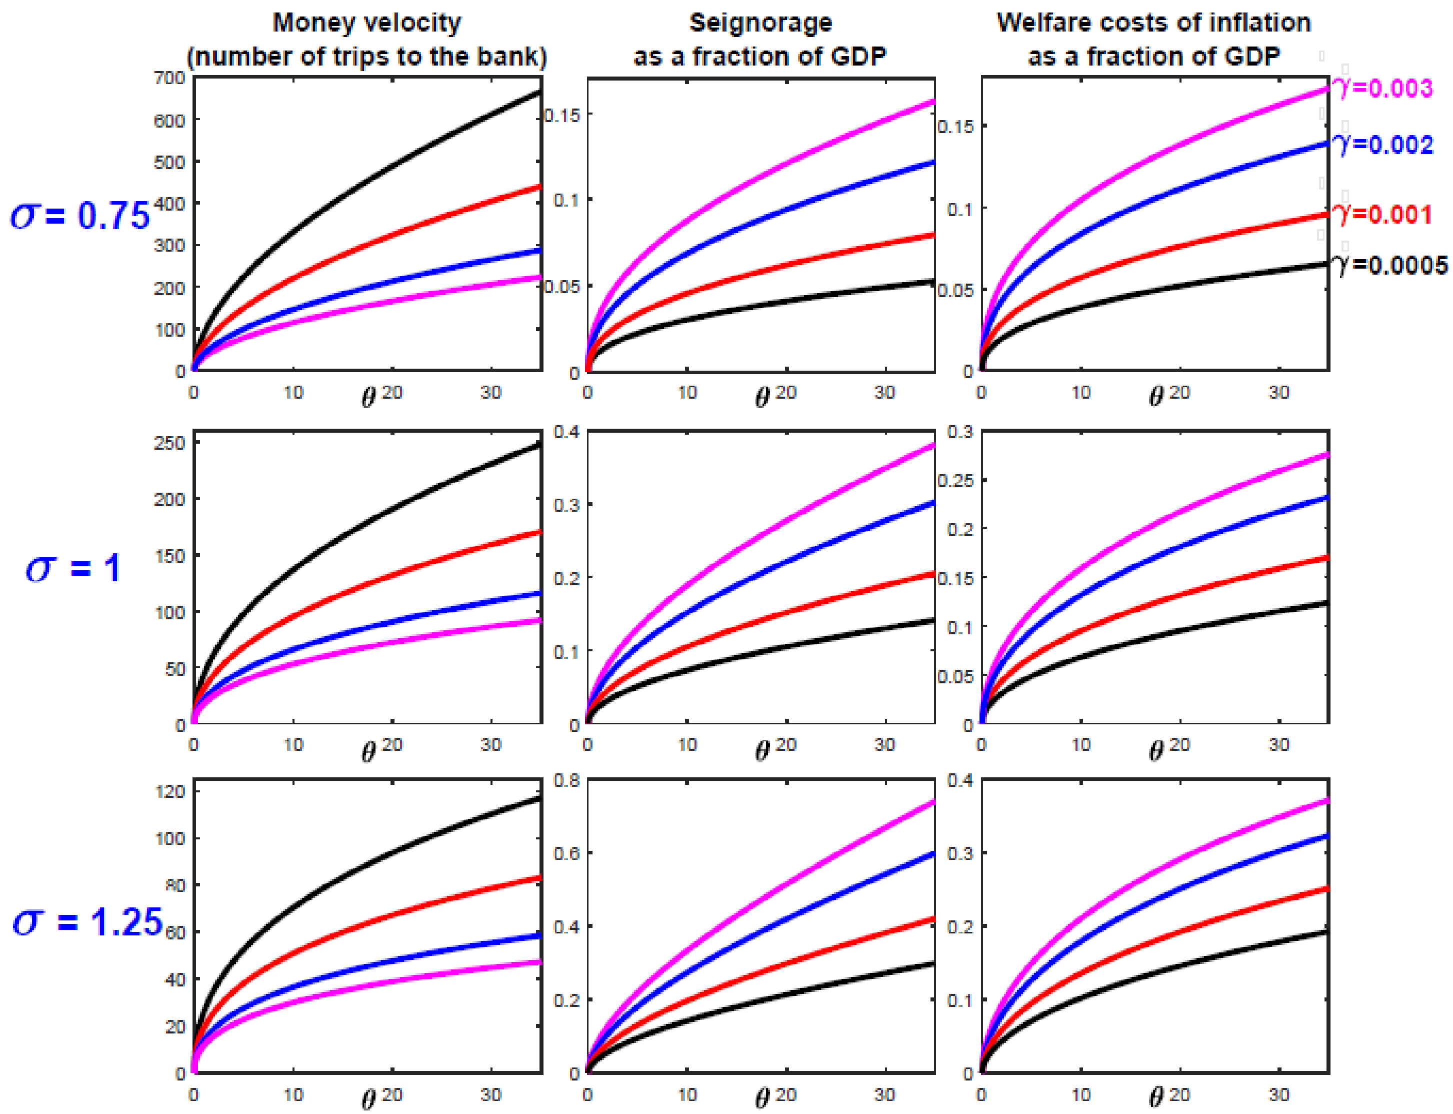

Figure 4 Money velocity, seignorage, and the welfare costs of inflation in the model of Benati, Lucas, Nicolini, and Weber (2021)

2.2.1 Steady-state and dynamics

In the steady-state $d\ln\pi_t^e/dt = 0$ and $\pi_t = \pi_t^e$, so that once again expression (7), $\pi_t = \pi_t^e = \theta_t$. holds. Within the present context the GG curve becomes

$$\pi_t^e = \sigma\psi\frac{\theta_t^{\sigma+1}\xi^{-(\sigma+1)}}{1-\psi\theta_t^{\sigma}\xi^{-\sigma}} \tag{16}$$

where once again ξ, which is assumed to be exogenously given, acts as a shifter for the GG curve. The steady-state equilibrium lies at the intersection between this curve and the 45 degree line (7).

It can be trivially shown that, from a qualitative point of view, both the shape of the GG curve, and the dynamical properties of the system, are exactly the same as those for the log-log that are shown in the right hand-side panel of Figure 1. In particular, since in equation (15) $\Psi(n_t) > 0$, this expression implies once again that when the economy's position on the GG curve is below (above) the 45 degrees line, so that $\pi_t - \theta_t < 0$ (> 0), $d\ln\pi_t^e/dt < 0$ (> 0). Once again the implication is that the steady-state A is stable, whereas the steady-state B is unstable, and beyond it lies a region of explosive inflation.

2.2.2 Money growth, velocity and seignorage in the steady-state

Figure 4 shows money velocity (i.e. n_t), together with seignorage and the welfare costs of inflation, both expressed as a fraction of GDP, as functions of θ for alternative values of σ and ψ. For θ I consider values from 0 to 35, which is just slightly higher than the maximum inflation rate ever recorded (for Hungary's post-WWII episode) of 33.67 (see the next section). Irrespective of the values of σ and ψ, all of the three objects are monotonically increasing in θ. Notice that the welfare costs of inflation as a fraction of maximum output, ψn^{σ}, are uniformly smaller than one, and as a result of this velocity is monotonically increasing in θ. For our purposes the feature of main interest is the fact that seignorage is uniformly increasing in money growth, so that there is no evidence of a Laffer curve for seignorage at any value of θ.

I now turn to a brief overview of the literature.

3 Related Literature

Cagan's (1956) paper spawned a vast literature. In this section I provide a brief overview, by narrowly focusing on two groups of studies: (1) classic papers, such as those of Sargent and Wallace (1973) and Sargent (1977), and (2) the very few studies containing results in line with this paper's position. Before delving into this, however, I start by providing a brief summary of Cagan's discussion of the most appropriate functional form for the demand for real money balances.

3.1 Cagan (1956) on the most appropriate functional form for the demand for real money balances

Cagan (1956) did not derive the semi-log specification (2) within a micro-founded framework, but rather simply postulated it.[4] In reaction to the empirical shortcoming of the postulated specification for the latests stages of hyperinflations, for which the models' fit had typically been worse than for the initial stages,[5] he speculated that an alternative functional form may be needed in order to provide a better characterization of the data. In particular, he conjectured[6]

> '[...] that the function that determines the demand for real cash balances does not conform to [the semi-log functional form]. To be consistent with the data, this hypothesis requires that all observations that lie to the right of the linear regression shall fall in order along some curved regression function [...].'

In practice, this means that the alternative specification he was speculating about should have been either a log-log, or a functional form close to it. In the end, however, Cagan's solution[7] was neither to use a log-log, nor a specification close to it, but rather to simply exclude, in some cases, the latest observations from the empirical analysis:

> 'The periods covered by the statistical analysis exclude some of the observations near the end of the hyperinflations. The excluded observations are from the German, Greek, and second Hungarian hyperinflations [...]. All the excluded observations lie considerably to the right of the regression lines, and their inclusion in the statistical analysis would improperly alter the estimates of α and β derived from the earlier observations of the hyperinflation.'

It is to be noticed that the three episodes whose latest observations Cagan excluded from the analysis are the *most extreme* in his dataset, i.e. those which, for the purpose of discriminating between the semi-log and the log-log, are the *most informative*.

I now turn to briefly discuss the two previously mentioned groups of studies. An important point to stress is that all of the four classic studies discussed in the next

[4] See Cagan (1956, p. 35).

[5] This problem has been repeatedly documented by several authors. E.g., for Yugoslavia's episode, see Petrovic and Mladenovic (2000).

[6] To be precise, Cagan (1956, p. 55) also discussed an alternative possible explanation, based on the notion that, towards the latest stages of hyperinflations, agents may come to expect a currency reform. Although, in principle, perfectly plausible, this alternative explanation suffers from the crucial limitation that it appears as implausible that expectations of a currency reform should have been such as to make the historical paths for inflation and real money balances, in several instances (Yugoslavia, Zimbabwe, post-WWI Greece and Germany, post-WWII Hungary, and, to a slightly lesser extent, China), to so clearly conform to those implied by a log-log specification (see Section 6.1).

[7] See Cagan (1956, Section IV.3, 'Observations that do not fit the regression', p. 55).

sub-section, as well as the overwhelming majority of existing studies of the demand for money during hyperinflations, have been based on Cagan's semi-log specification.

3.2 Classic studies

Sargent and Wallace (1973) pointed out that Cagan's (1956) estimator of the semi-elasticity of money demand was inconsistent under rational expectations, and documented how, in Cagan's dataset, inflation Granger-caused money growth, whereas money growth failed to Granger-cause inflation.

Sargent (1977) showed how, under rational expectations, the semi-elasticity of money demand could in fact be identified by assuming that shocks to money demand and money supply be contemporaneously uncorrelated. A key result he obtained was that estimates of the semi-elasticity of money demand based on Cagan's dataset were characterized by a very large uncertainty. In particular,

'[t]he estimates are so loose that confidence bands of two standard errors on each side of them include values that would imply that the creators of money were inflating at rates that maximized their command over real resources, thus maybe resolving [Cagan's] paradox [...].'

Salemi and Sargent (1979) postulated a vector autoregressive (VAR) representation for the joint dynamics of inflation and money growth, and estimated it *via* maximum likelihood conditional on the rational expectations restrictions implied by Cagan's semi-log functional form. Consistent with Sargent (1977), a main finding was that the extent of econometric uncertainty surrounding the point estimates of the semi-elasticity of money demand was much more substantial than for Cagan's (1956) estimates, which, again, could be taken to provide a possible explanation for Cagan's paradox.

Taylor (1991) introduced cointegration methods to the study of the demand for real money balances during hyperinflations. As he first pointed out, if both inflation and real money balances are I(1), and under the minimal assumption that the forecast errors are I(0), cointegration allows to test for the presence of a demand for real money balances—rather than postulating it, as it was done in the previous literature—and to estimate it *via* maximum likelihood. Following Taylor (1991), several papers have applied cointegration techniques to the study of the demand for money during hyperinflations. A key limitations of these studies is that, to the very best of my knowledge, they have all been based on asymptotic critical values, which, in small samples, have been shown to be essentially unreliable.[8] In the next sub-section I discuss a specific example, Zhao's (2018) study of the Chinese hyperinflation.

[8] E.g. Johansen (2002), with reference to his trace and maximum eigenvalue tests, showed that asymptotic critical values are essentially unreliable in small samples.

3.3 Studies conceptually in line with the present work

Barro (1970), still working within a pre-rational expectations framework, developed a sophisticated model for the demand for real money balances allowing for currency substitution, which produced a markedly better fit than Cagan's (1956) semi-log specification.[9] For the present purposes, the crucial point is that Barro's empirical specification—see his equations (74)-(75)—boiled down to a linear relationship between log real money balances and the *logarithm* of the sum of expected inflation, the real interest rate, and additional terms. The superior fit of Barro's specification, compared to Cagan's (1956), is therefore in principle compatible with the present work's results.

Zhao (2017) is, to the very best of my knowledge, the paper which is closest, in terms of its main objective, to the present work. Based on data for China's hyperinflation, it uses cointegration techniques in order to address the issue of which, among the semi-log and the log-log functional forms, provides a a better characterization of the data. Taken at face value, his results are in line with mine: whereas he detects cointegration between log real money balances and the logarithm of inflation, he does not detect it between log real balances and inflation's level. A limitation of Zhao's results is that his cointegration tests are based on asymptotic critical values, which should be regarded as unreliable because of the short sample length. In fact, performing the same Johansen's tests reported in Zhao's (2017) Table 6, but bootstrapping them as in Cavaliere, Rahbek, and Taylor (2012), I obtain p-values for the trace and maximum eigenvalue test statistics equal to 0.245 and 0.156, respectively, based on the logarithm of inflation, and equal to 0.514 and 0.617 based on its level. This suggests that although, from the perspective of the present work, Zhao (2017) did obtain the correct result, in fact he produced it based on an unreliable procedure.

I now turn to discussing the data and exploring the empirical relationship between inflation and seignorage.

4 The Data

Table 1 reports, for each individual episode, the number of observations, together with the maximum and median inflation rate.[10] Whenever possible, I follow Cagan (1956), and I set the end of the hyperinflationary episode at 12 months after inflation had last exceeded the threshold he proposed, of 0.5 on a log scale.[11] The dataset features

[9]As he pointed out (see p. 1257), '[i]n general, the average errors in Cagan's form are about twice as large as those [based on Barro's specification], and the serial correlation of residuals is substantially more pronounced'.

[10]Based on monthly data, inflation has been computed as the month-on-month log-difference (in natural logarithms) of the relevant price index, whereas based on weekly data it has been computed as the week-on-week log-difference multiplied by 4.

[11]In fact, in several cases this was not possible due to either lack of, or discontinuities in one or more series. E.g., for China and Yugoslavia the dataset ends in May 1949 and January 1994,

Table 1 Maximum and median inflation during hyperinflationary episodes[a]

	Sample length	Maximum	Median
	Monthly data		
France (January 1794-June 1796)	27	0.761	0.154
Based on Cagan's (1956) data			
Austria (January 1921-August 1922)	20	0.852	0.212
Germany (September 1920-November 1923)	39	5.885	0.157
Hungary (July 1922-February 1924)	20	0.683	0.215
Hungary (July 1945-July 1946)	13	33.670	1.677
Poland (April 1922-January 1924)	22	1.322	0.302
Russia (January 1922-February 1924)	26	1.142	0.420
Greece (January 1943-November 1944)	23	13.659	0.470
Based on Barro's (1970) data			
Austria (January 1921-December 1922)	24	0.852	0.149
Poland (January 1922-January 1924)	25	1.229	0.289
Germany (January 1921-August 1923)	32	2.931	0.136
Hungary (October 1921-February 1924)	29	0.683	0.152
Other datasets			
Germany[b] (December 1919-December 1923)	60	5.736	0.090
China (August 1945-May 1949)	45	4.446	0.240
Chile (January 1972-November 1974)	35	0.575	0.106
Bolivia (February 1983-August 1986)	43	1.039	0.155
Argentina (January 1987-April 1991)	52	0.992	0.135
Brazil (August 1988-March 1991)	33	0.592	0.199
Peru (January 1987-September 1991)	57	1.512	0.147
Yugoslavia (January 1991-January 1994)	37	11.290	0.496
Congo (May 1991-July 1995)	51	2.144	0.241
Angola (December 1995-January 1998)	26	0.610	0.061
Bulgaria (January 1996-February 1998)	26	1.230	0.039
Zimbabwe (January 2004-June 2008)	53	3.912	0.300
Venezuela (January 2016-March 2019)	39	1.313	0.200
	Weekly data		
Germany (September 7, 1922-November 15, 1923)	57	13.285	0.804
Hungary (December 31, 1945-July 23, 1946)	26	35.280	3.280

[a] Based on monthly data, inflation is computed as the month-on-month log-difference of the price level. Based on weekly data, it is computed as the week-on-week log-difference multiplied by 4. [b] Based on Graham's (1930) data.

a dramatic extent of variation of inflationary experiences, ranging from post-WWII Hungary's peak of 33.67;[12] to the second and third most extreme episodes, Greece and Yugoslavia, with peak inflation rates of 13.66 and 11.29, respectively; down to the mildest, Chile, with a maximum inflation rate of 0.58. Exactly two thirds of the episodes had inflation rates in excess of 1, whereas 18.5 and 11.1 per cent had rates in excess of 5 and 10, respectively.

Appendix B provides a detailed description of the data and of their sources for all of the episodes in chronological order, from the French Revolution to Venezuela. Since evidence based on weekly data is uniformly in line with that based on monthly data, in what follows I will mention it only briefly, and I will instead near exclusively focus on the results based on monthly data.

I now turn to exploring the empirical relationship between inflation and seignorage, both at the level of individual countries, and based on panel regressions with country-specific fixed effects.

5 Exploring the Empirical Relationship Between Inflation and Seignorage

5.1 Evidence for Cagan's (1956) dataset based on Sargent and Wallace's (1973) measure of seignorage

The top row of Figure 5 shows, for the seven episodes in Cagan's (1956) dataset, scatterplots of the logarithm of Sargent and Wallace's (1973, Table 6, p. 345) measure of seignorage plotted against the logarithm of money growth from Cagan's dataset,[13] whereas the bottom row shows the evolution of the two series over time (with different scales for the left and right hand-side axes). With the single exception of Russia,[14] for all other six countries the relationship between the two series clearly appears to have been positive *at all levels* of money growth.

respectively, when the inflation rate (month-on-month log-difference) had been equal to 4.44 and 8.35.

[12] Hungary's peak of 33.67 was reached in July 1946. The data for Hungary shown in Figure 1 stop in June 1946 (for which the inflation rate was 11.34) because Sargent and Wallace's (1973) Table 6 does not report the seignorage figure for the month of July.

[13] To be precise, Cagan's dataset features, for each month, the log-difference of the price level and the logarithm of the ratio between the price level and the money stock. Based on these series, however, the log-difference of the money stock can be trivially computed.

[14] The evidence for Russia should be treated with some caution. E.g., Barro (1970) eschewed Russian data because (see his footnote 36) 'the assumption of constant real income appeared unreasonable and adequate income data was unavailable'. In what follows I will instead use Cagan's data for Russia, because since we are here dealing with *hyper*inflations, even an unaccounted-for deep recession should reasonably be thought of as only introducing a minor distortion in the estimates.

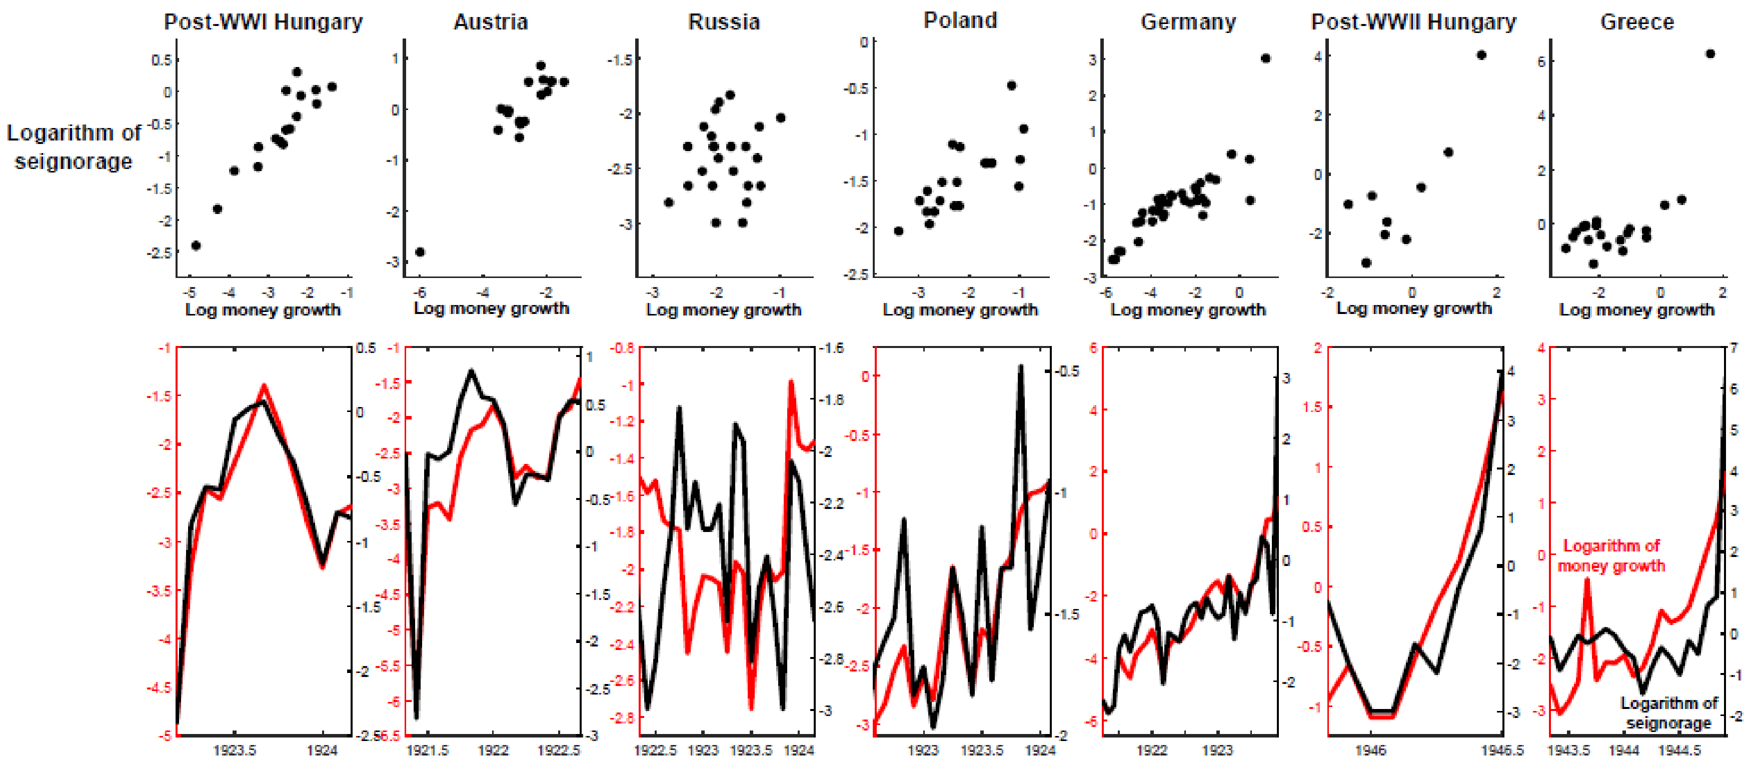

Figure 5 The logarithm of Sargent and Wallace's (1973) measure of seignorage plotted against the logarithm of money growth from Cagan (1956)

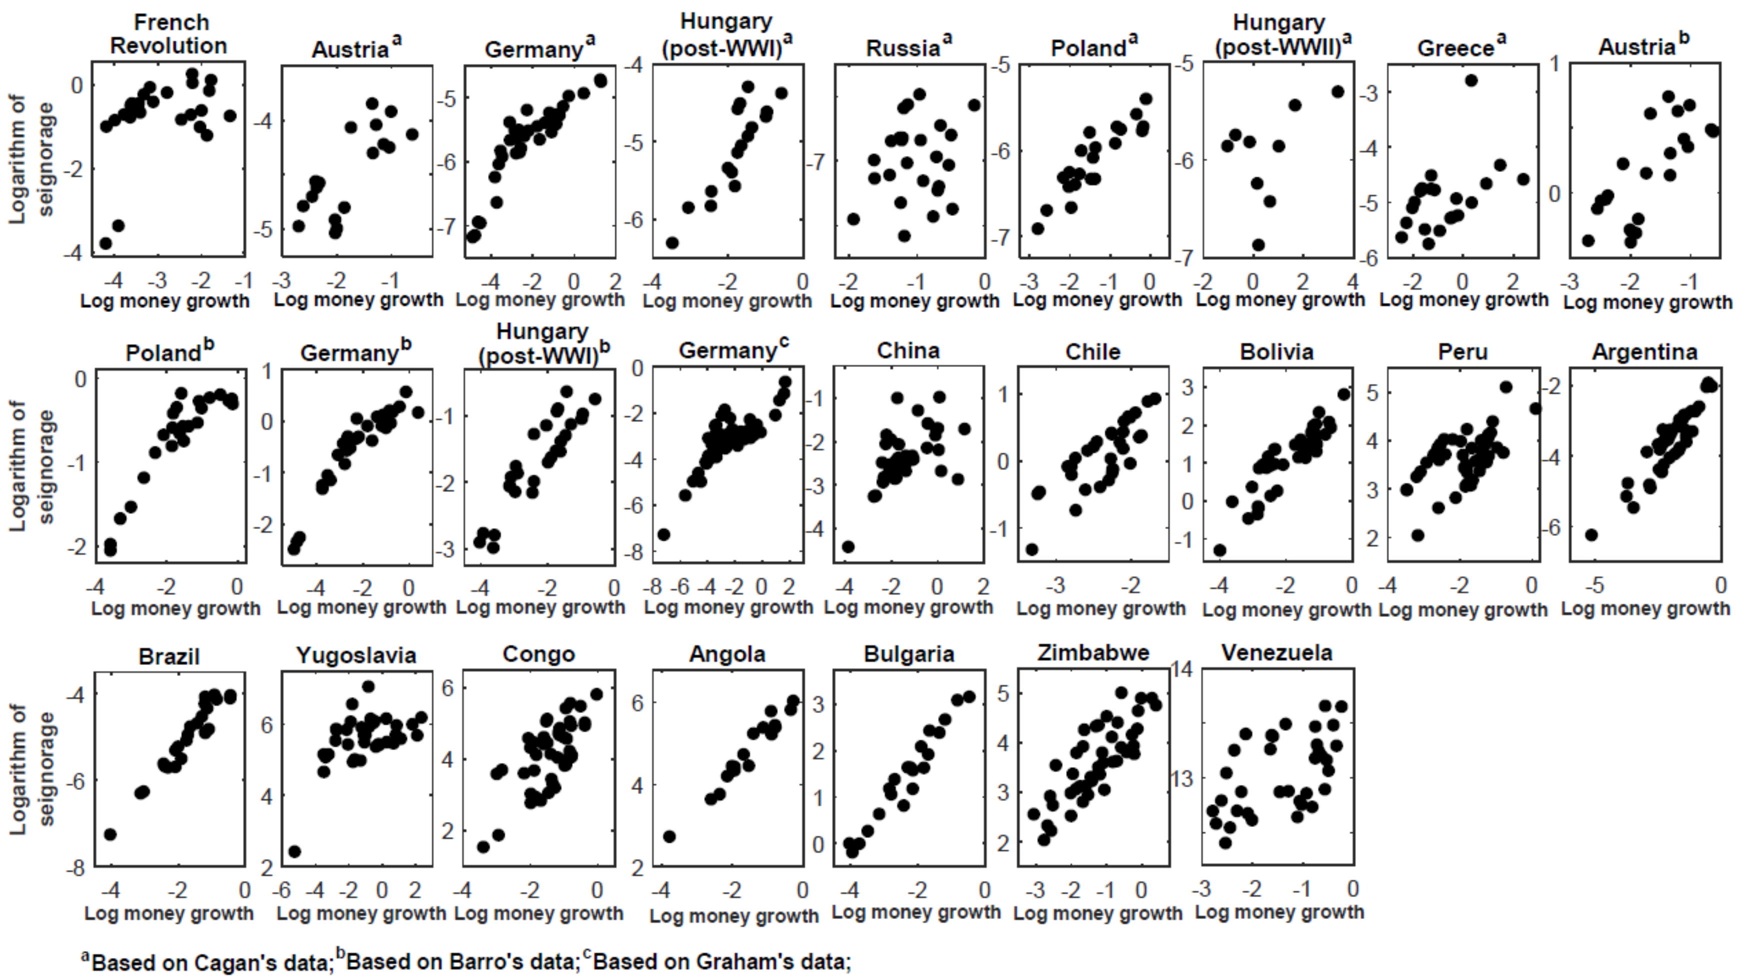

Figure 6 The logarithms of money growth and seignorage for the 20 hyperinflations (monthly data)

5.2 Evidence based on a measure of seignorage in the spirit of Bresciani-Turroni (1937)

One possible caveat to the evidence in Figure 5 is that Sargent and Wallace (1973, see p. 341, and Table 6) computed their measure of seignorage based on the *arithmetic* average

$$\tilde{\xi}_t \equiv \frac{M_t - M_{t-1}}{\frac{1}{2}\left(P_t + P_{t-1}\right)}. \tag{17}$$

where M_t and P_t are the money stock and the price level in month t. As discussed by Bresciani-Turroni's (1937, pp. 146-154), however, within the context of hyperinflations, which expecially in their very last stages are characterized by explosive growth in money and the price level, *geometric* averages should logically thought of as superior. In Appendix C I perform an extensive investigation of this issue, and I show that Bresciani-Turroni's conjecture is indeed correct. In particular, I show that under a wide range of empirically plausible circumstances the geometric average

$$\xi_t^* \equiv \theta_t^* \left[\left(\frac{M_{t-1}}{P_{t-1}} \right)^{\omega} \left(\frac{M_t}{P_t} \right)^{1-\omega} \right] \tag{18}$$

where $\theta_t^* \equiv \ln(M_t)$-$\ln(M_{t-1})$=$m_t$-$m_{t-1}$, and ω=0.5, exhibits a vastly superior performance compared to the arithmetic average (17). This is the case, in particular, when money and prices evolve according to *exponential*, as opposed to linear trends, as it seems quite clearly to be the case under hyperinflations. In the rest of the paper I will therefore exclusively work with measures of seignorage computed based on the geometric average (18).

Figure 6 shows for all 20 episodes, and when possible based on alternative datasets,[15] the same evidence shown in the top row of Figure 5, with seignorage computed based on (18). Once again, the impression of a uniformly positive relationship between money growth an seignorage at all levels of money growth is very clear. This is the case in particular for Germany, Hungary, Poland, Chile, Bolivia, Argentina, Brazil, Congo, Angola, Bulgaria, Zimbabwe, and Venezuela. On the other hand, evidence is less than clear-cut for the French Revolution, Russia, and China. Even these three episodes, however, provide essentially no evidence of a Laffer curve for seignorage.[16]

[15] E.g., for Germany based on either Graham's (1930), Cagan's (1956), or Barro's (1970) data.

[16] Figure A.4 in the Appendix shows evidence for the Confederacy during the U.S. Civil War based on data from Lerner (1956). Although neither the Union nor the Confederacy experienced a hyperinflation (for the Confederacy the maximum quarterly log-difference of the price level was equal to 0.465), the evidence in Figure A.4 is starkly in line with that in Figures 5 and 6.

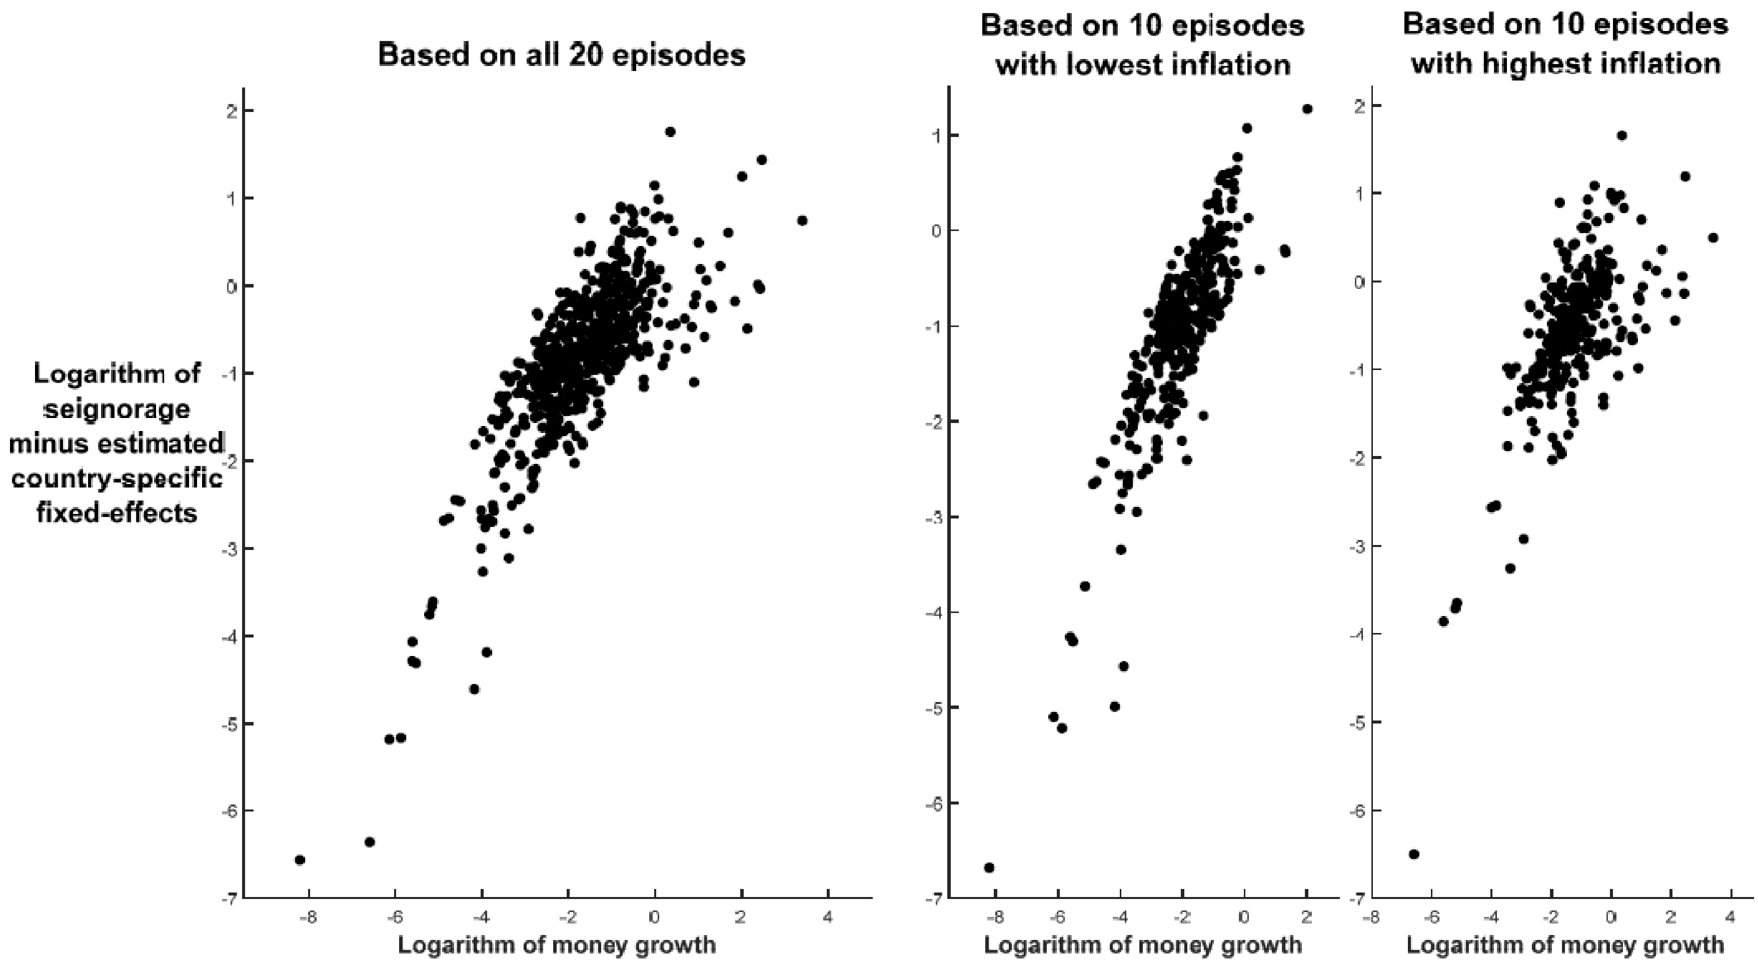

Figure 7 Evidence from regressing log seignorage on log money growth: logarithm of seignorage minus estimated country-specific fixed effects

5.3 Evidence from panel regressions with country-specific fixed effects

Figure 7 reports evidence from Least Absolute Deviations[17] (LAD) panel regressions of the logarithm of seignorage on the logarithm of money growth and country-specific fixed-effects,[18]

$$\ln \xi_{i,t} = \psi_1 \ln \Delta m_{i,t} + \psi_2 (\ln \Delta m_{i,t})^2 + \lambda_i, \tag{19}$$

where i indexes the country and t indexes the month, based on both all of the 20 episodes considered jointly, and the 10 episodes with either the highest or the lowest median inflation rates.[19] The number of observations is 636 based on all of the 20 episodes, and 291 and 345 based on the two datasets with the high- and low-inflation episodes, respectively. The figure shows scatterplots of $\ln \Delta m_{i,t}$ and $[\ln \xi^*_{i,t}\text{-}\hat{\lambda}_i^{LAD}]$—i.e. log seignorage minus the estimated country-specific fixed effects—together with the LAD regression line. Results are robust to adding a cubic term, $\psi_3(\ln \Delta m_{i,t})^3$, in the regression (these results are available upon request).

The evidence in the figure is quite clear: first, the LAD regression lines are monotonically increasing, showing no evidence of the hump shape associated with a Laffer curve; second, the three scatterplots clearly suggest that, historically, higher money growth has been associated on average with higher seignorage *at all inflation levels*. These results, together with the evidence in Figure 5 and 6, provide a clear refutation of the notion of a Laffer curve for seignorage.

5.4 Rationalizing the absence of a Laffer curve for seignorage

How can we explain the absence of a Laffer curve in the data from 20 hyperinflations? Although in principle several alternative rationalizations might be possible, Occam's razor suggests to consider the simplest one. As discussed in Section 2, a Laffer curve for seignorage is a mathematical property of the semi-log specification that is independent of the specific value taken by the semi-elasticity of money demand (see e.g. Lucas, 2000.). As a result, when a researcher estimates a semi-log money demand curve (s)he is not 'discovering' a Laffer curve: (s)he is *imposing* it upon the data. Following Cagan (1956), the overwhelming majority of the papers in the literature—

[17] Evidence from the corresponding OLS regressions is qualitatively the same as that in Figure 2, and it is available upon request.

[18] There are two reasons for including country-specific fixed effects in the regressions. First, in order to control for country-specific idiosyncratic factors. Second, whereas inflation and money growth, being defined as the log-differences of the relevant objects, are pure numbers, (M_t/P_t) in the expression for ξ^*_t is the ratio between a quantity expressed in national currency units and an index number, and it is therefore an index number itself. As a result, ξ^*_t is also an index number, and its value is therefore defined up to a factor of proportionality.

[19] For Austria, Germany, post-WWI Hungary, and Poland, for which I have multiple data sources, I used Cagan's data.

in particular, the most notable and influential among them,[20] i.e. Sargent and Wallace (1973), Sargent (1977), and Salemi and Sargent (1979)—have been based on the semi-log functional form. It should therefore come as no surprise that the notion of a Laffer curve for seignorage has become part of macroeconomists' conventional wisdom. As I have shown, however, evidence quite clearly suggests that such a notion is incorrect.

Based on Meltzer's log-log, on the other hand, the relationship between inflation and seignorage is positive at all inflation levels for any value of the elasticity of money demand smaller than one in absolute value, whereas based on Benati et al.'s (2021) functional form it is uniformly positive at all inflation levels in the empirically relevant in which velocity is increasing in inflation. A straightforward explanation for the previously documented absence of a Laffer curve in the data is therefore that either the true functional form for the demand for real money balances is Benati et al.'s (2021), or it is the log-log with the elasticity being smaller than one in absolute value.

In fact, as I show in the next section, the data provide strong support to this position.

6 Which Functional Form Best Describes the Data?

6.1 The dynamics of real money balances and inflation

Figure 8 shows, for the six most extreme hyperinflations, the logarithm of real money balances together with either the *level* of inflation (in the top row), or its *logarithm*[21] (in the bottom row), whereas Figures A.5 and A.6 in the Online Appendix show the same evidence for the remaining episodes. The evidence in the top and bottom rows therefore corresponds to a semi-log and, respectively, a log-log specification for the demand for real money balances, relating log real balances to either the level or the logarithm of inflation. The episodes have been ranked in decreasing order based on the median inflation values reported in Table 1, from post-WWII Hungary (with

[20] An important exception is represented by Barro (1970), who developed a model of the demand for real money balances allowing for currency substitution, which produced a substantially better fit than the semi-log specification (see Barro (1970, p. 1257): '[i]n general, the average errors in Cagan's form are about twice as large [...], and the serial correlation of residuals is substantially more pronounced'). For the present purposes, the crucial point is that Barro's specification—see his equations (74)-(75)—boiled down to a linear relationship between log real money balances and the *logarithm* of the sum of expected inflation and additional terms. The superior fit of Barro's specification compared to Cagan's is therefore compatible with the notion that the true money demand specification is not the semi-log, but rather the log-log (or a functional form close to it).

[21] Because of the very high-frequency of the data, even if we are here dealing with hyperinflationary episodes, in a few instances inflation turned out to be negative. In all of these cases, the corresponding observations for log inflation are plotted as missing. It is important to stress that based on either (3) or (10), the demand for real money balances depends on *expected* inflation, which within the context of a hyperinflation should be thought to have been uniformly positive.

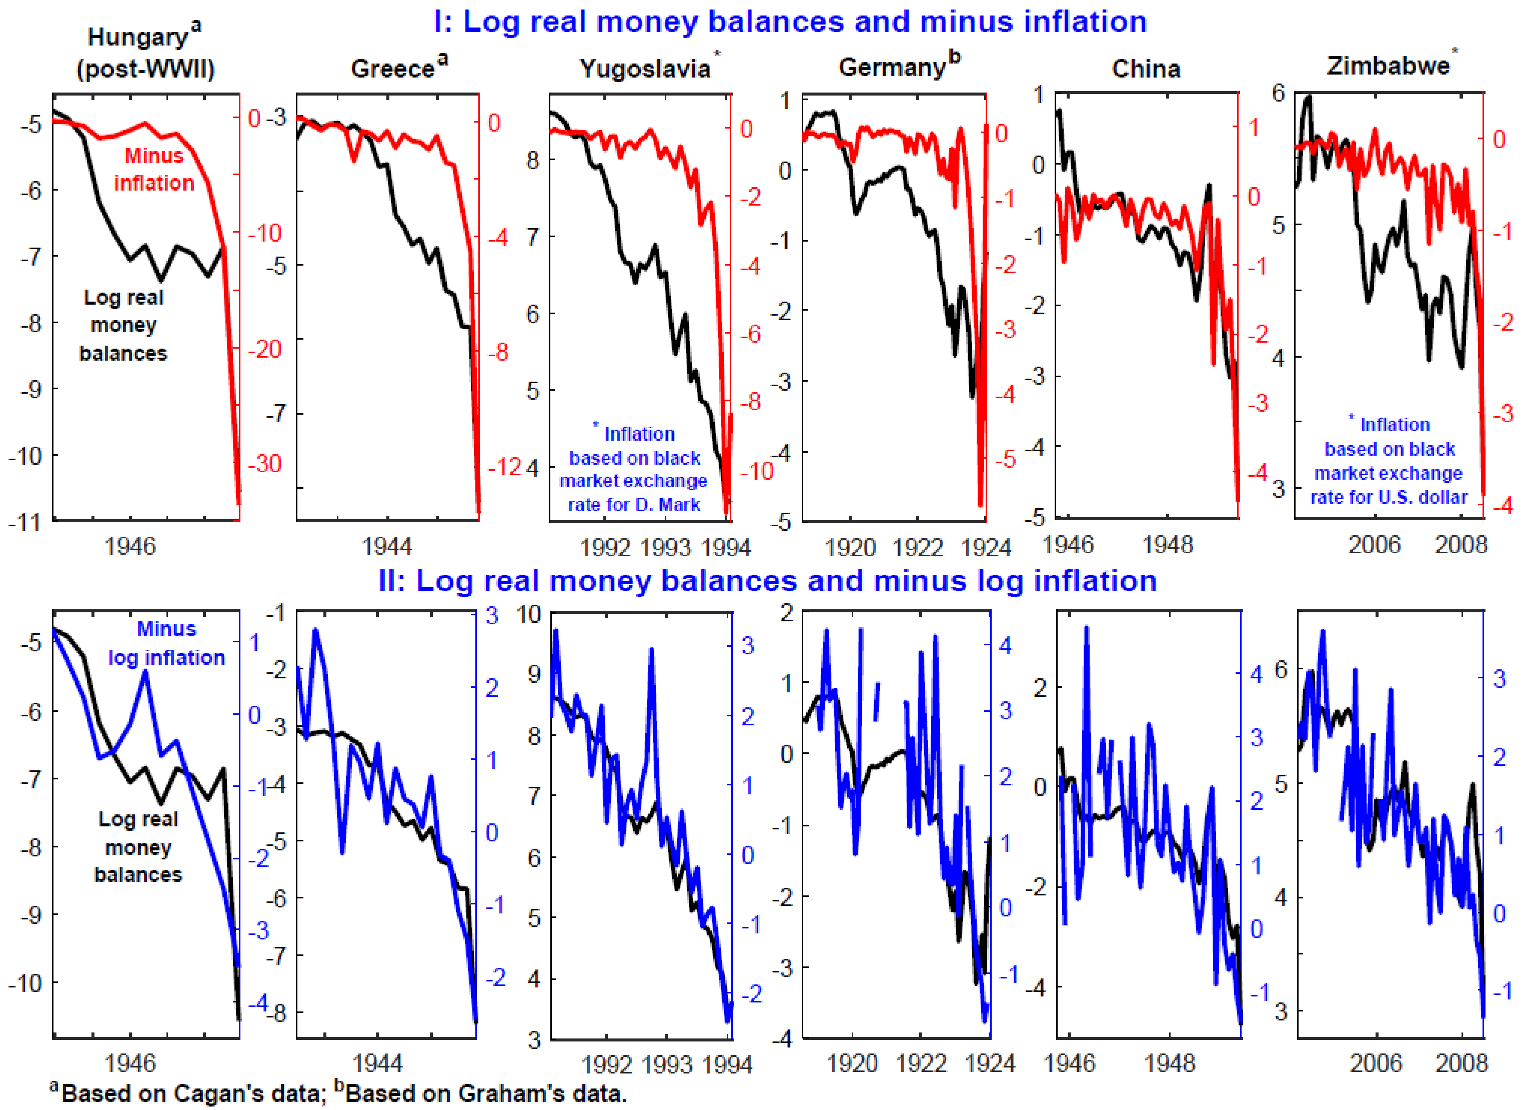

Figure 8 Monthly raw data for log real money balances and (log) inflation for extreme hyperinflations

median inflation equal to 1.677) to Chile (0.106). The following facts emerge quite clearly from the figures:

(1) for the most extreme episodes the logarithm of inflation tracks log real money balances remarkably closely,[22] whereas its level exhibits a uniformly weaker, or even scant connection to it. This is especially clear for post-WWII Hungary, Greece, Yugoslavia, Germany, and Zimbabwe, whereas evidence for China in slightly weaker. This suggests that for extreme hyperinflations—i.e., the episodes which should be regarded as the most informative—the log-log specification (or a functional form close to it, such as Benati et al.'s) provides a more plausible description of the joint dynamics of the two series than the semi-log.

(2) For all other episodes, visual evidence suggests that the two specifications are on an essentially equal footing in terms of their ability to characterize the joint dynamics of the data.

Overall, the evidence in Figures 6, A.5 and A.6 therefore suggests that if we had to choose which, between the semi-log and the log-log, best describes the joint dynamics of inflation and real money balances, the natural choice would be the log-log.

6.2 Long-run evidence on money velocity and nominal interest rates

Figure 9 provides additional evidence based on long-run data for M1 velocity[23] and a short-term nominal interest rate for six high-inflation countries from Benati et al.'s (2021) dataset. Specifically, the figure shows the logarithm of M1 velocity together with either the *level* of a short-term rate (in the top row), or its *logarithm* (in the bottom row). The evidence in the top and bottom rows therefore corresponds to a semi-log and, respectively, a log-log specification for the demand for real money balances with unitary income elasticity, relating log velocity to either the level, or the logarithm, of a short-term rate. Once again, the episodes have been ordered from the most to the least extreme (based on the maximum value taken by the nominal interest rate).

The evidence in the figure is once again vely clear: whereas fluctuations in log M1 velocity bear a uniformly weak, or even scant connection to movements in the *level* of the short-term rate, they are typically strongly correlated with its *logarithm*. This is starkly apparent for Israel and Bolivia, for which the logarithms of M1 velocity and

[22] Up to a scale factor, which in the three figures is accounted for by allowing the right hand-side and the left hand-side scales to differ.

[23] I focus on velocity (defined as the ratio between nominal GDP and nominal M1), rather than real money balances, because of the sizeable increases in GDP that have taken place over such long sample periods. On the other hand, following Cagan (1956), in the literature on hyperinflation real GDP is typically assumed to be costant. Since, historically, hyperinflations have consistently been short-lived episodes, this assumption is, for the purpose of these studies, innocuous. When considering longer sample periods, however, changes in real GDP cannot be ignored, and the most appropriate variable becomes velocity.

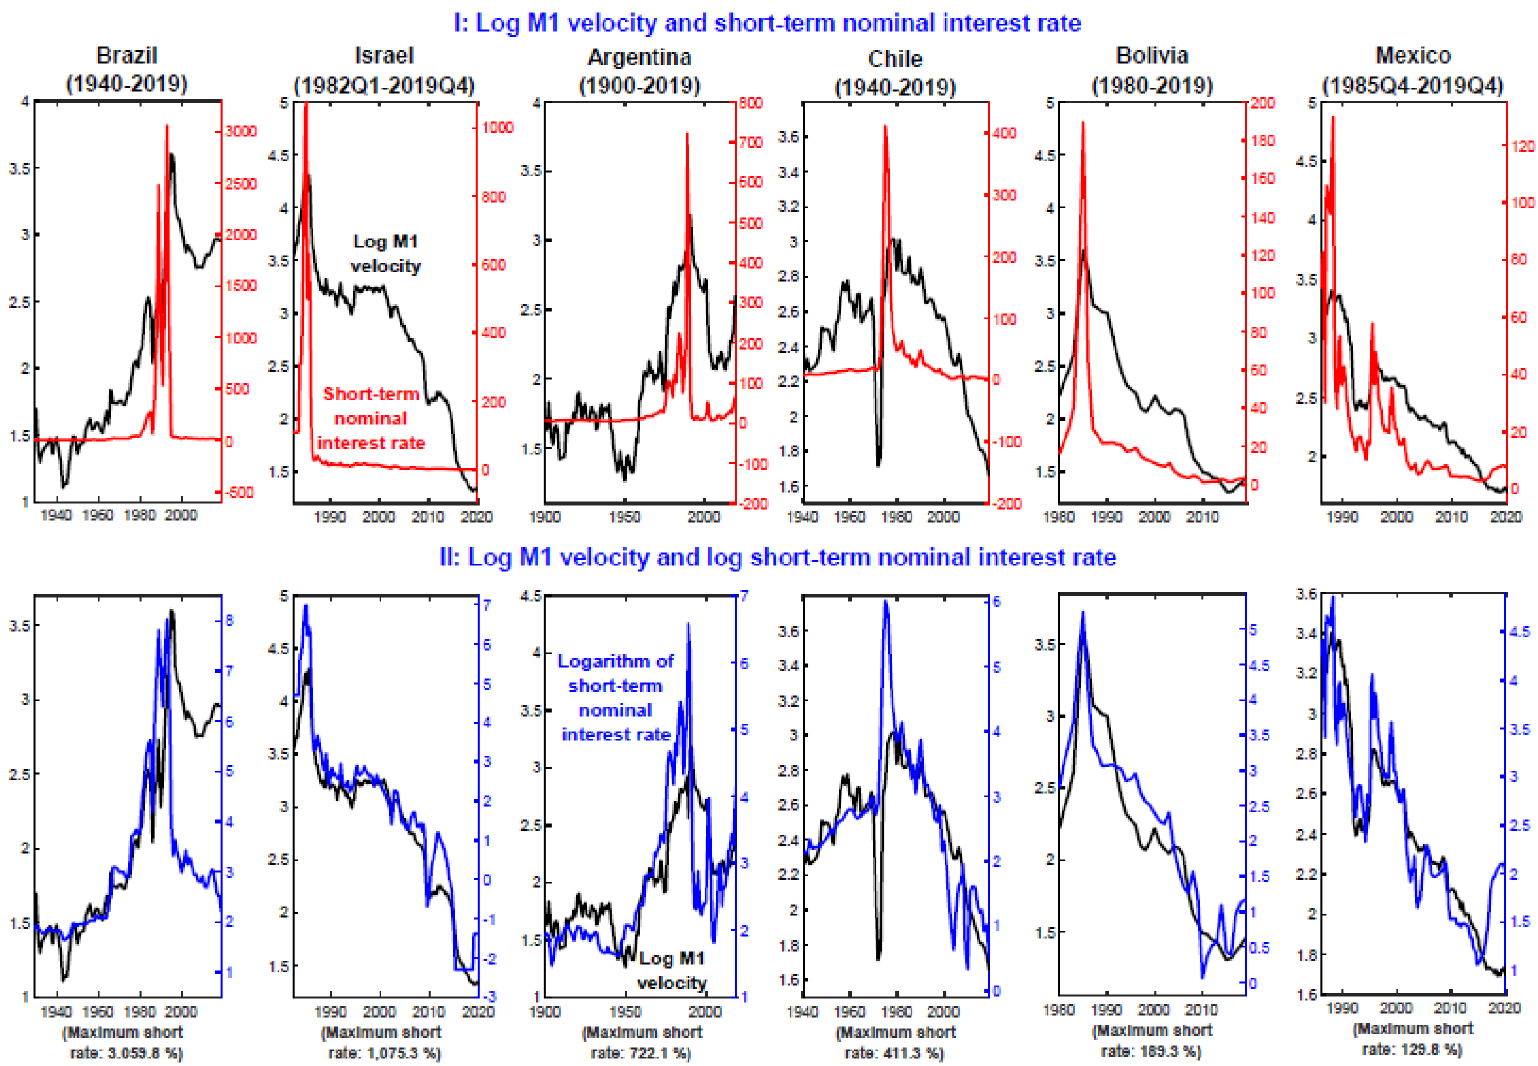

Figure 9 Long-run data for M1 velocity and a short-term nominal interest rate for high-inflation countries

of the short-term rate have very closely co-moved over the entire sample periods; it is just slightly less so for Argentina—which had exhibited some temporary deviations between the two series around 1950 and following the disinflation of the end of the 1980s—and for Chile with the single exception of the early 1970s, a period of exceptional economic and social turmoil which culminated in Augusto Pinochet's military *coup* of 1973; and it is also apparent for Mexico, and for Brazil during the entire course of the XX century, whereas the two series have been moving out of synch since the start of the new millenniun.

The reason why this evidence is so stark is straightforward. Historically, hyperinflations have uniformly been short-lived episodes, lasting at most a few years. Over such short periods of time it may therefore be difficult to discriminate between alternative functional forms for the demand for real money balances. The longer the sample period, however, the more extreme the range of inflationary experiences—from hyperinflation to (near) price stability—typically becomes, with the result that the inferiority of the semi-log specification becomes manifestly apparent even to the naked eye.

A counter-argument to this is that, as the economy approaches, and then enters a full-blown hyperinflation, the demand for real balances may change in fundamental ways, as phenomena such as currency substitution, which are either negligible or second-order at lower inflation rates, become more and more relevant, thus causing a progressive increase in the elasticity of money demand. Under these circumstances the log-log specification, with its constant elasticity of the demand for real balances, might provide a less accurate characterization of the data than the semi-log, for which the elasticity is increasing (in absolute value) with the opportunity cost of money. Under this interpretation, the evidence in Figure 9 should be regarded as uninformative, and only data pertaining to hyperinflationary episodes should be regarded as relevant for the issue of determining which functional form best describes hyperinflation data. As shown in the previous sub-section, however, even when narrowly focusing on hyperinflation episodes, evidence quite clearly suggests that the most plausible description of the data is provided by the log-log, rather than by the semi-log.

The bottom line is therefore that whether we narrowly focus on the comparatively short windows of time associated with hyperinflations, or we consider much longer sample periods of several decades, there is simply no evidence in favor of the semi-log, whereas there is strong evidence in favor of the log-log, or a functional form close to it.

6.3 Evidence from a VAR-based model comparison exercise

The evidence in the previous two sub-sections is especially persuasive because it is based on the raw data. In this section I complement it with the following model comparison exercise. Based on both all of the 20 episodes considered jointly, and the 10 episodes with either the highest or the lowest median inflation rates, I estimate *via* maximum likelihood the following two specifications for the joint dynamics of the logarithm of real money balances and inflation:

$$\begin{bmatrix} \tilde{m}_{i,t} \\ \pi_{i,t} \end{bmatrix} = \begin{bmatrix} c_i^{\tilde{m}} + A(L)\tilde{m}_{i,t\text{-}1} + B(L)\pi_{i,t\text{-}1} + \epsilon_{i,t}^{\tilde{m}} \\ c_i^{\pi} + C(L)\tilde{m}_{i,t\text{-}1} + D(L)\pi_{i,t\text{-}1} + \epsilon_{i,t}^{\pi} \end{bmatrix} \tag{20}$$

and

$$\begin{bmatrix} \tilde{m}_{i,t} \\ \pi_{i,t} \end{bmatrix} = \begin{bmatrix} c_i^{\tilde{m}} + A(L)\tilde{m}_{i,t\text{-}1} + B(L)\tilde{\pi}_{i,t\text{-}1} + \epsilon_{i,t}^{\tilde{m}} \\ c_i^{\pi} + \exp\left\{C(L)\tilde{m}_{i,t\text{-}1} + D(L)\tilde{\pi}_{i,t\text{-}1}\right\} + \epsilon_{i,t}^{\pi} \end{bmatrix} \tag{21}$$

where i indexes the country; t indexes the month; $\tilde{m}_{i,t} \equiv \ln(M_{i,t}/P_{i,t})$; $\tilde{\pi}_{i,t} \equiv \ln(\pi_{i,t})$; $A(L)$, $B(L)$, $C(L)$, and $D(L)$ are polynomials in the lag operator, L; $c_i^{\tilde{m}}$ and c_i^{π} are country-specific intercepts; and $\epsilon_{i,t}^{\tilde{m}}$ and $\epsilon_{i,t}^{\pi}$ are country-specific residuals, which I postulate to follow a bivariate normal distribution with a non-diagonal covariance matrix. Equation (20) describes a panel VAR model with country-specific fixed-effects for log real money balances and inflation, and it therefore corresponds to the semi-log specification. Expression (21), on the other hand, postulates that, up to country-specific dummies and random disturbances, the joint dynamics of real money balances and inflation is described by

$$\begin{bmatrix} \tilde{m}_{i,t} \\ \tilde{\pi}_{i,t} \end{bmatrix} = \begin{bmatrix} A(L) & B(L) \\ C(L) & D(L) \end{bmatrix} \begin{bmatrix} \tilde{m}_{i,t\text{-}1} \\ \tilde{\pi}_{i,t\text{-}1} \end{bmatrix} \tag{22}$$

corresponding to the log-log specification. By casting (22) into the form (21)—i.e, taking as the dependent variable, in the second equation, the *level* of inflation, rather than its logarithm—it is possible to meaningfully compare, in terms of log-likelihood, which of the two functional forms provides the most plausible description of the data. As we will see, evidence overwhelmingly favors the log-log.[24]

I estimate either (20) or (21) *via* maximum likelihood. Specifically, I estimate the mode of the log-likelihood *via* simulated annealing as described in Corana, Marchesi, Martini, and Ridella (1987), and I then proceed to stochastically map the log-likelihood's surface *via* Random-Walk Metropolis (RWM). The only difference between the 'standard' RWM algorithm which is routinely used for Bayesian estimation and what I am doing here is that the jump to the new position in the Markov chain is accepted or rejected based on a rule which does not involve any Bayesian priors, as it uniquely involves the likelihood of the data.[25] All other estimation details are

[24]For the log-log I have also considered an alternative to (21) in which, in the equation for $\pi_{i,t}$, the country-specific dummies are inside the exponential function, i.e. $\pi_{i,t} = \exp[c_i^{\pi} + C(L)\tilde{m}_{i,t\text{-}1} + D(L)\tilde{\pi}_{i,t\text{-}1}] + \epsilon_{i,t}^{\pi}$. The results produced by this alternative specification are qualitatively the same as those produced by (21), and they are available upon request.

[25]So, to be clear, the proposal draw for β, $\tilde{\beta}$, is accepted with probability $\min[1, r(\beta_{s-1}, \tilde{\beta} \mid Y, X)]$, and rejected otherwise, where β_{s-1} is the current position in the Markov chain, and

$$r(\beta_{s-1}, \tilde{\beta} \mid Y, X) = \frac{L(\tilde{\beta} \mid Y, X)}{L(\beta_{s-1} \mid Y, X)}$$

identical to Benati (2008), to which the reader is referred to. I use one million draws for the burn-in period, and twenty million for the ergodic distribution, which I 'thin' by sampling every 10,000 draws in order to reduce the draws' autocorrelation, thus obtaining 2,000 draws for the ergodic distribution. For any of the three panels of countries, either of the two lag orders I consider (2 or 4), and based on either (20) or (21), the fractions of accepted draws are uniformly very close to the 23 per cent ideal acceptance rate in high dimensions.[26] Figure A.7 in the Appendix reports, for each individual model parameter, the inefficiency factor[27] for the draws from the ergodic distribution. The inefficiency factors are uniformly around 1-2, i.e. significantly below the value of 20-25 which is typically taken as signalling problems in convergence.

Table 2 Evidence from the model comparison exercise: minima, maxima, and medians of the distributions of the log-likelihoods produced by Random-Walk Metropolis

	Model (20)			Model (21)		
	Min	Median	Max	Min	Median	Max
	Based on all 20 episodes					
p=2	303.1	324.8	337.1	557.9	583.1	596.2
p=4	219.5	245.4	259.7	470.0	494.6	508.4
	Based on 10 episodes with highest median inflation					
p=2	-6.4	9.6	20.3	102.0	119.0	129.0
p=4	-34.9	-13.3	-1.6	81.0	109.0	122.2
	Based on 10 episodes with lowest median inflation					
p=2	599.8	625.4	635.5	660.0	678.8	689.3
p=4	549.5	570.7	582.9	646.6	667.9	678.8

Table 2 reports the minima, medians, and maxima of the ergodic distributions of the log-likelihood produced by RWM.[28] The key result in the table is that the minima of the distributions based on model (21) are uniformly greater than the corresponding

which uniquely involves the likelihood. With Bayesian priors it would be

$$r(\beta_{s-1}, \tilde{\beta} \mid Y, X) = \frac{L(\tilde{\beta} \mid Y, X)P(\tilde{\beta})}{L(\beta_{s-1} \mid Y, X)P(\beta_{s-1})}$$

where $P(\cdot)$ would encodes the priors about β.

[26] See Gelman, Carlin, Stern, and Rubin (1995).

[27] The inefficiency factor is defined as the inverse of the relative numerical efficiency measure of Geweke (1992), $RNE = (2\pi)^{-1} \frac{1}{S(0)} \int_{-\pi}^{\pi} S(\omega) d\omega$, where $S(\omega)$ is the spectral density of the sequence of draws from the ergodic distribution for the parameter of interest at the frequency ω. I estimate the spectral densities based on the Fast Fourier transform.

[28] Notice that with p=2 the number of overall observations is equal to 555 for the panel of 20 countries, and to 247 and 308 for the panels of 10 countries with the highest and lowest inflation rates, whereas the corresponding figures with p=4 are 481, 206, and 275. This explains why, *ceteris paribus*, log-likelihoods are uniformly greater for p=2 than they are for p=4. In terms of likelihood

maxima of the distributions based on model (20), thus suggesting that the log-log provides a more plausible characterization of the data than the semi-log.

I now turn to estimating panel VARs for the logarithms of real money balances and inflation.

7 Evidence from Panel VARs for the Logarithms of Inflation and Real Money Balances

I estimate *via* OLS the following panel VAR specification for the logarithms of the two series,

$$\left[\begin{array}{c} \tilde{m}_{i,t} \\ \tilde{\pi}_{i,t} \end{array}\right] = \left[\begin{array}{c} c_i^{\tilde{m}} \\ c_i^{\pi} \end{array}\right] + \left[\begin{array}{cc} A(L) & B(L) \\ C(L) & D(L) \end{array}\right] \left[\begin{array}{c} \tilde{m}_{i,t\text{-}1} \\ \tilde{\pi}_{i,t\text{-}1} \end{array}\right] + \left[\begin{array}{c} \epsilon_{i,t}^{\tilde{m}} \\ \epsilon_{i,t}^{\pi} \end{array}\right] \quad (23)$$

where the notation is the same as before, characterizing uncertainty about the estimates *via* standard bootstrapping methods,[29] based on 10,000 bootstrap replications. I set the lag order to either 2 or 4.

Table 3*a* reports, for any of three panels and for either lag order, the point estimates of the four largest eigenvalues of the VAR, together with the fractions of bootstrap replications for which they are estimated to be smaller than one. The main findings in the table can be summarized as follows:

(1) in all cases, all eigenvalues except the largest one are significantly smaller than one. In particular, (*i*) their point estimates range between 0.279 and 0.859 and (*ii*) the fractions of bootstrap replications for which they are estimated to be smaller than one are consistently greater than 99 per cent, and in fact in most cases they are equal to 100 per cent.

(2) As for the largest eigenvalue of any of the VARs, (*i*) in nearly all cases it is not possible to reject the null hypothesis that it is equal to one at the 10 per cent level, and (*ii*) the point estimates are exactly equal to one based on the 10 most extreme episodes; they are equal to 0.98 based on all of the 20 episodes; and they are somehow smaller (0.926 and 0.945) only for the 10 least extreme episodes.

These results are compatible with the notion that the joint dynamics of the logarithms of real money balances and inflation *is driven by a single unit root process*, which implies that the two series are cointegrated. As I will discuss in Section 9.2, a natural conjecture is that the unit root in the system originates from permanent variation in the amount of seignorage the government is attempting to raise, which,

per data point (i.e., dividing the log-likelihood by the number of observations), the difference between p=2 and p=4 is, *ceteris paribus*, small to negligible.

[29] To be clear, the methodology is exactly the same as (e.g.) Barsky and Sims (2011), with the only difference that they performed estimation for a single country, whereas I am doing it for panels of either 20 or 10 (and I therefore have country-specific intercepts).

in turn, is due to the often explosive behavior of the government budget deficit during hyperinflations.

Table 3*a* Evidence from panel VARs for the logarithms of real money balances and inflation: point estimates for the four largest eigenvalues, and fractions of bootstrap replications for which the eigenvalues are greater than 1

	Point estimates for the four largest eigenvalues				Fractions of bootstrap replications[a] for which λ_i<1			
	λ_1	λ_2	λ_3	λ_4	λ_1	λ_2	λ_3	λ_4
	Based on all 20 episodes							
p=2	0.314	0.314	0.672	0.980	1.000	1.000	1.000	0.770
p=4	0.582	0.582	0.834	0.981	1.000	1.000	0.992	0.741
	Based on 10 episodes with highest median inflation							
p=2	0.313	0.313	0.459	1.009	1.000	1.000	1.000	0.283
p=4	0.505	0.505	0.509	1.009	1.000	1.000	1.000	0.269
	Based on 10 episodes with lowest median inflation							
p=2	0.279	0.279	0.765	0.926	1.000	1.000	1.000	0.935
p=4	0.731	0.731	0.859	0.945	1.000	1.000	0.999	0.838

[a] The bootstrapped distributions have been rescaled so that the median of the distribution is equal to the point estimate.

7.1 Comparison with the semi-log

It is instructive to compare the results in Table 3*a* with those produced by the corresponding VARs for the logarithm of real money balances and the *level* of inflation. These results are reported in Table 3*b*. The main findings are that (*i*) estimates of the largest eigenvalue range between 1.249 and 1.729, and (*ii*) in most cases, they are estimated to be significantly greater than one. This suggests that the semi-log functional form distorts the inference along an additional, and crucial dimension, by spuriously suggesting that the joint dynamics of real money balances and inflation is *explosive*, rather than simply possessing a unit root.

Table 3*b* Evidence from panel VARs for inflation and the logarithm of real money balances: point estimates for the four largest eigenvalues, and fractions of bootstrap replications for which the eigenvalues are greater than 1								
	Point estimates for the four largest eigenvalues				Fractions of bootstrap replications[a] for which λ_i<1			
	λ_1	λ_2	λ_3	λ_4	λ_1	λ_2	λ_3	λ_4
	Based on all 20 episodes							
p=2	0.330	0.330	0.937	1.728	0.876	0.816	0.534	0.124
p=4	0.631	0.631	0.924	1.680	0.795	0.746	0.550	0.063
	Based on the 10 episodes with highest median inflation							
p=2	0.061	0.094	0.936	1.729	0.953	0.891	0.539	0.035
p=4	0.674	0.969	0.969	1.584	0.799	0.536	0.522	0.105
	Based on the 10 episodes with lowest median inflation							
p=2	0.367	0.367	0.970	1.249	1.000	1.000	0.761	0.000
p=4	0.550	0.610	1.008	1.273	0.961	0.913	0.431	0.000

[a] The bootstrapped distributions have been rescaled so that the median of the distribution is equal to the point estimate.

I now turn to estimating the elasticity of the demand for real money balances, starting from (23), which I re-estimate by imposing that the largest eigenvalue is exactly equal to one, which amounts to imposing one cointegration vector.

8 Estimating the Elasticity of Money Demand

8.1 Evidence from panel VARs

I start by estimating the panel VARs (23) in its vector error-correction (VECM) form, imposing in estimation one cointegration vector. As mentioned, this is equivalent to imposing the restriction that the largest eigenvalue in (23) is exactly equal to one, which is near-uniformly accepted by the data. I estimate the models *via* the two-stage procedure discussed by Luetkepohl (1991, Section 11.2.2*b*, pp. 370-372), which is appropriate when, as in the present case, there is just one cointegration vector. Specifically, in the first stage I estimate the cointegration vector between the logarithms of real money balances and inflation based on Stock and Watson's (1993) 'dynamic OLS' (DOLS) estimator, setting k, the number of leads and lags in the DOLS procedure, to either 1 or 2. In the second stage I then estimate the VECM form of (23) *via* OLS, imposing in estimation the cointegration residual obtained in the first stage. I characterize uncertainty about the estimates by bootstrapping the estimated VECM as in Cavaliere, Rahbek, and Taylor (2012). I estimate the elasticity of money demand as the second element of the normalized cointegration vector (with the normalization being performed on log real money balances). Table 4 reports the point estimates of the elasticity, the 90 per cent coverage bootstrapped confidence intervals, and the fraction of the bootstrapped distribution below -1, which, as mentioned, is the threshold below which an increase in inflation is associated with a decrease in seignorage with log-log money demand. For reasons of space I only report results based on p=2. Results for the alternative lag order (p=4) are numerically very close, and are available upon request.

Table 4 Evidence from panel cointegrated VARs for the logarithms of real money balances and inflation: point estimates of the elasticity, 90% bootstrapped confidence interval, and fractions of bootstrap replitions for which the elasticity is smaller than -1

	Point estimates, and 90% confidence interval	Fractions of bootstrap replications for which elasticity is below -1
	Based on all 20 episodes	
k=1	-0.502 [-0.593; -0.412]	0.000
k=2	-0.554 [-0.656; -0.454]	0.000
	Based on 10 episodes with highest median inflation	
k=1	-0.792 [-0.896; -0.681]	0.001
k=2	-0.926 [-1.041; -0.796]	0.143
	Based on 10 episodes with lowest median inflation	
k=1	-0.295 [-0.392; -0.191]	0.000
k=2	-0.312 [-0.412; -0.202]	0.000

Evidence is reported in Table 4. Based on the joint panel of all 20 episodes the estimated elasticity is either equal to Baumol and Tobin's benchmark value of -1/2, or very close to it, and there is no evidence that it might have been smaller than -1. The implication is that, historically, the relationship between inflation and seignorage has consistently been monotonically increasing at all inflation levels.

Splitting the overall sample into the most and the least extreme 10 episodes, however, uncovers two important differences between the two sets.

First, the most extreme episodes have consistently been characterized by a larger elasticity (in absolute value), with the difference between the point estimates for the two groups of episodes being equal to -0.497 based on k=1, and to -0.614 based on k=2. Further, as Figure A.8 in the Appendix shows, there is essentially no overlapping between the bootstrapped distributions for the estimated elasticities for the two sets of episodes. This suggests that the most extreme episodes have been characterized by a significantly more elastic demand for real money balances than comparatively milder episodes. Two possible, non mutually exclusive explanations for this pattern are that the most extreme the hyperinflation, (i) the more widespread the phenomenon of currency substitution becomes, so that, *ceteris paribus*, the elasticity of the demand for real money balances in national currency progressively increases; and (*ii*) agents devote more and more resources to forecasting inflation, with the result that they get surprised less and less (i.e., expectations become 'more and more rational').

Second, there is some, although admittedly weak evidence that for the most extreme episodes the elasticity might have been greater than one in absolute value. As we will now see, these results are confirmed by two alternative approaches to the estimation of the elasticity of money demand.

8.2 Evidence from two alternative approaches

Salemi and Sargent (1979) postulated a VAR representation for the joint dynamics of inflation and money growth, and estimated it *via* maximum likelihood conditional on the rational expectations restrictions implied by Cagan's semi-log functional form. In Appendix D I adopt an approach combining Salemi and Sargent's (1979) insight of postulating a time-series representation for the series of interest, and imposing upon it the restrictions implied by a theoretical specification for the demand for real money balances, with elements borrowed from Hamilton (1985) and Burmeister and Wall (1982, 1987). Specifically, I postulate a VAR representation for the logarithms of real money balances and the expectation of inflation at $t+1$ conditional on information at t, and I estimate it *via* MLE by imposing the restrictions implied by Meltzer's log-log specification. As in Section 6.3 I characterize uncertainty about the estimates by stochastically mapping the log-likelihood's surface *via* RWM. Consistent with the use of country-specific dummies in Sections 5.2, 6.3, 7, and 8.1 I allow the intercepts in the VAR representation for $\ln(M_t/P_t)$ and $\ln(\pi_{t+1|t})$ to be country-specific, whereas I impose that all of the remaining parameters be the same for all of the countries in the panel. The first panel of Figure A.10 in the Appendix shows the distributions of the draws from the ergodic distributions for the elasticity of real money balances, both for all of 20 the episodes considered jointly, and for the 10 episodes with either the highest or the lowest median inflation rates, respectively. Evidence is qualitatively in line with, and quantitatively very close to, that based on panel VARs discussed in the previous sub-section. In particular,

(*i*) evidence strongly suggests that higher inflation has consistently been associated with a larger (in absolute value) elasticity. In particular, whereas based on all of the 20 episodes the median and the 5th and 95th percentiles of the distribution of the elasticity are equal to -0.573 [-0.669; -0.486], the corresponding figures for the 10 most and least extreme episodes are -0.932 [-1.094; -0.759] and -0.306 [-0.135; -0.468], respectively.

(*ii*) Only for the panel featuring the 10 most extreme episodes there is some evidence that the elasticity may have been greater than one in absolute value, with 22 per cent of the draws being associated with values of the elasticity below -1. On the other hand, for either the 10 least extreme episodes, or all of the 20 episodes considered jointly, all of the draws are associated with values of the elasticity greater than -1.

Finally, Appendix D.2 reports evidence from a third approach in which, at each point in time, expected inflation next period is computed based on either an AR(p) model for inflation, or a VAR(p) model for inflation and money growth, which are estimated *via* constant-gain recursive least squares. The gain parameter is estimated *via* MLE together with the other structural parameters of the model, and uncertainty about the estimates is characterized, once again, *via* RWM. Evidence based on the univariate and bivariate schemes for forecasting inflation are qualitatively the same, and numerically very close, which is in line with Sargent and Wallace's (1973) classic result that, within a bivariate VAR representation for money growth and inflation, the former does not Granger-cause the latter. The second panel of Figure A.10 in the Appendix shows the distributions of the draws from the ergodic distributions for the elasticity of real money balances, based on the univariate forecasting model for inflation (results based on the bivariate model are available upon request). Evidence is broadly in line with that discussed so far, with the only difference that, for any of the three sets of episodes, the estimated elasticities are uniformly greater in absolute value. In particular, for the panel of 10 most extreme episodes, there is strong evidence that the elasticity may have been slightly below -1, with the median and the 5th and 95th percentiles of the distribution being equal to -1.034 [-1.055; -1.016], and the fractions of draws below -1 being equal to 0.999. On the other hand, based on either all of the 20 hyperinflations jointly considered, or the 10 least extreme episodes, not a single draw is associated with values of the elasticity greater than one in absolute value, with the median and the 5th and 95th percentiles of the distribution being equal to -0.905 [-0.918; -0.892], and -0.741 [-0.759; -0.722, respectively.

I now turn discussing the implications of my findings, starting from the interpretation of historical episodes of hyperinflation.

9 Implications

9.1 The standard view of hyperinflation is incorrect

As mentioned in the Introduction, since Cagan's (1956) landmark study the literature on hyperinflation has been dominated by a narrative featuring two main elements: (1) the relationship between money growth and seignorage exhibits a Laffer-curve property, and (2) historically, governments have near-uniformly inflated in excess of the revenue-maximizing rate, i.e., they have been on the 'wrong side' of the Laffer curve. This view of the world can be found in all graduate textbooks—from from (e.g.) Blanchard and Fischer (1990, Chapter 4, pp. 195-201), to Obsteld and Rogoff (1996, Chapter 8, pp. 515-530), to Walsh (2017, Chapter 4, pp. 153-162)—and it has been taught in PhD courses for decades. As my evidence has shown, this view is incorrect: the relationship between money growth and seignorage had, and has been uniformly positive at all inflation rates. This suggests that previous studies of

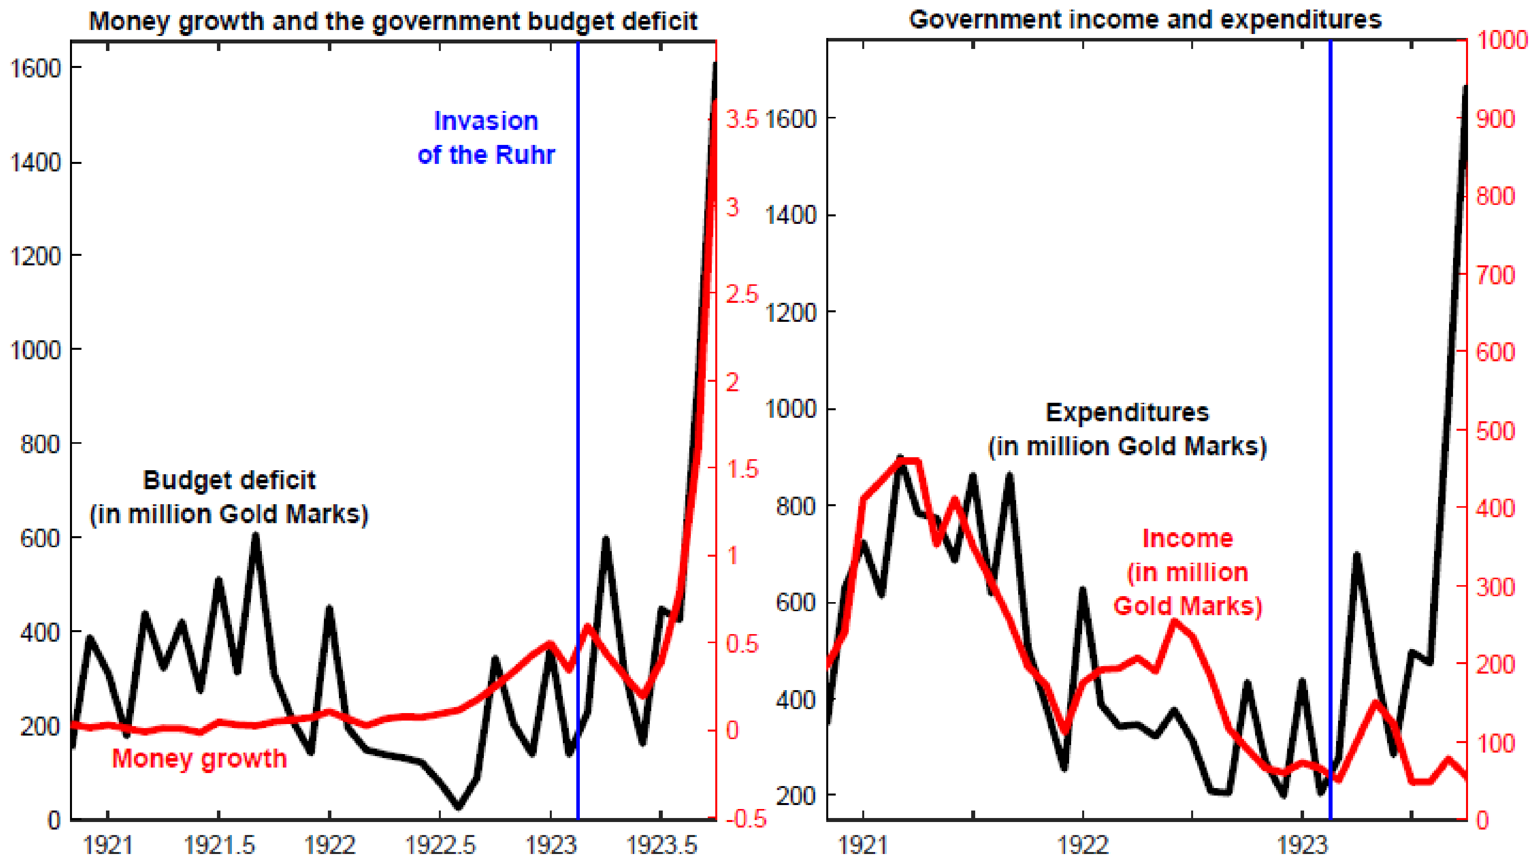

Figure 10 The budget deficit, money growth and seignorage in the Weimar Republic

hyperinflation are incorrect along several dimensions.

9.2 Interpreting historical episodes of hyperinflation

If the true functional form for the demand for real money balances were Cagan's semi-log, what we would need are theories explaining why a rational government may end up on the 'wrong' side of the Laffer curve for seignorage. The standard explanation, due to Sargent and Wallace, is that under rational expectations only the high-inflation steady-state is stable, whereas the low-inflation one is unstable.

If, on the other hand, the true functional form is Meltzer's log-log (or a specification close to it such as Benati et al.'s), and in the light of the fact that historically the relationship between seignorage and inflation had been uniformly positive, what we rather need are theories explaining why a rational government may have decided to move the economy along such a monotonically increasing relationship.

Consider e.g. once again the last four panels of Figure 5. Why would the Polish, German, Hungarian, and Greek governments have decided to create money at ever-increasing rates, thus raising ever-increasing amounts of seignorage? In line with Bresciani-Turroni (1937) and Sargent (1982), a plausible explanation is that they were attempting to plug ever-increasing holes in the government budget, i.e. to finance progressively increasing deficits *via* money creation.

Some of the available evidence on the evolution of budget deficits during hyperinflations is compatible with this notion. For Germany's episode, the data from Bresciani-Turroni's (1937) Table I, page 437, plotted in the first panel of Figure 10, show for the latest stages of the hyperinflation (i.e., since early 1923) a strong positive correlation between money growth and the budget deficit, computed as the difference between the expenditures and the income (i.e., essentially taxes) in million of Gold Marks (i.e., in real terms).[30]

One possible explanation for such a strong correlation could be the phenomenon described by Bresciani-Turroni (1937) and by Olivera and Tanzi: since taxes are specified in nominal terms at time t, and they are due at a later date, $t+\tau$, an inflation outburst between t and $t+\tau$ would have reduced their real value, thus increasing, *ceteris paribus*, the budget deficit. The second panel of Figure 10 shows evidence on this, based on the same data from Bresciani-Turroni (1937). The inflation explosion of 1923 had indeed been associated with a decrease in the government's income. Such a decrease, however, was minor compared to both previous decreases in 1921 and 1922, when inflation had been comparatively more stable, and especially the contemporaneous evolution of expenditures in 1923, which literally skyrocketed. Even if we are willing to attribute the entire decrease in income in 1923 to the Bresciani-

[30] For countries other than Germany evidence is not clear cut, because (to the very best of my knowledge) the data for government income and expenditures are routinely reported in nominal terms for individual fiscal *years* (see e.g. the data in Sargent (1982)). This implies that it is not possible to reconstruct the month-on-month evolution of the real government budget deficit.

Turroni-Olivera-Tanzi effect, still the explosion in expenditures appears to have played the dominant role in the corresponding explosion in the budget deficit.

In fact, as discussed by Bresciani-Turroni (1937, Section VIII, pp. 62-64) the key event under this respect was the occupation of the Ruhr:

> 'In the first months of 1923 the occupation of the Ruhr gave the coup de grâce to the national finances and the German mark. Because of it some important sources of income were lost to the State [...]. In addition, the German Governments [...] did not think to cover the heavy expenses caused by Passive Resistance with new taxes. It conceded very large credits to the Ruhr industry to put it in a position to continue production.'.

The evidence in Figure 10 therefore suggests that the Bresciani-Turroni-Olivera-Tanzi effect had only played a minor role in the explosion of the budget deficit that took place in 1923, and that the dominant cause had rather been the invasion of the Ruhr on the part of France.

9.3 Interpreting explosive inflation in the latest stages of hyperinflations

As shown in Figure 2, in several cases inflation exhibits an explosive dynamics towards the end of hyperinflations. With Cagan's semi-log the *only* way for the model to generate explosive inflation is for the government to attempt to collect an amount of seignorage greater than the maximum feasible (i.e., that associated with the peak of the Laffer curve). Under these circumstances, the GG curve and the 45 degrees line in the left hand-side panel of Figure 1 do not touch, so that there is no steady-state to speak of, and as a result inflation simply explodes to infinity. With the log-log or (e.g.) Benati et al.'s (2021) functional form, on the other hand, explosive inflation can result even if in fact any amount of seignorage is feasible, in the sense that it can be collected in a steady-state. Consider the economy described by the right hand-side panel of Figure 1: if A and B are sufficiently close, a large shock may cause the economy to jump to a point on the GG curve to the right of B.

10 Conclusions, and Three Directions for Future Research

Since it was first documented by Cagan (1956), policymakers' (alleged) tendency to inflate in excess of the revenue-maximizing rate during hyperinflations has been confirmed by several subsequent studies, to the point that it has nearly achieved the status of a stylized fact in empirical macroeconomics. Based on data from 20 hyperinflations I have shown that in the data there is nearly no evidence of a Laffer curve for seignorage. I have argued that a likely explanation for this is that the money demand specification which best describes the joint dynamics of inflation and real money balances during hyperinflations is not Cagan's semi-log, which features the Laffer curve as a mathematical property independent of the value of the semi-elasticity of money demand, but rather Meltzer's log-log (or a specification close to it), which for values of the elasticity smaller than 1 in absolute value produces the monotonically increasing relationship between inflation and seignorage seen in the data. My evidence suggests that Cagan's paradox is the product of the literature's predominant focus on the semi-log functional form.

An obvious direction for future research is to explore the issue issue of whether, during hyperinflations, the economy may have been operating under indeterminacy, so that hyperinflatonary episodes may have been influenced by sunspots. Sargent and Wallace (1987) developed a model of monetary financing of the government budget deficit *via* the inflation tax allowing for the possibility of indeterminate equilibria, but they did not take it to the data, neither (to the very best of my knowledge) they did in subsequent work. Based on the evidence reported in the present work, the starting point should therefore be to perform a theoretical analysis along the lines of Sargent and Wallace's (1987) based on the log-log functional form, and to then estimate the model based on the dataset I have used in the present work, possibly expanded with additional series pertaining to government finances.

A second natural extension is to apply to the context of hyperinflations the methodology proposed by Ascari, Bonomolo, and Lopes (2019) for computing temporarily unstable paths within rational expectations DSGE models. Within the context of the U.S. Great Inflation of the 1970s, Ascari et al. (2019) show that when allowing for temporarily unstable paths the data quite clearly select them as the most plausible explanation of the Great Inflation. Given the extreme nature of hyperinflations, Ascari et al.'s (2019) results logically suggests that we should expect to obtain the same results.

A third possible extension is to apply to the context of hyperinflations the methodology proposed (e.g.) by Angeletos, Collard, and Dellas (2018) in order to introduce 'hall of mirrors' effects into DSGE models, thus allowing to capture the impact of higher-order expectations on macroeconomic dynamics. The very nature of hyperinflations logically suggests that such effects may have played an important role, since individuals' propensity to 'run away from money' (i.e., shrink their real money balance in order to minimize the inflation tax) crucially depends on what they expect other individuals to do, which in turns depends on what other individuals expect other individuals to do, and so on.

References

Alvarez, F., and F. Lippi (2009): "Financial Innovation and the Transactions Demand for Cash," *Econometrica*, 77(2), 363–402.

Anderson, R. B., W. A. Bomberger, and G. E. Makinen (1988): "The Demand for Money, the Reform Effect, and the Money Supply Process in Hyperinflation: The Evidence from Greece and Hungary II Reexamined," *Journal of Money, Credit and Banking*, 20(4), 653–672.

Angeletos, G., F. Collard, and H. Dellas (2018): "Quantifying Confidence," *Econometrica*, 86(5), 1689–1726.

Ascari, G., P. Bonomolo, and H. F. Lopes (2019): "Walk on the Wild Side: Temporarily Unstable Paths and Multiplicative Sunspots," *American Economic Review*, 109(5), 1805–1842.

Barro, R. J. (1970): "Inflation, the Payments Period, and the Demand for Money," *Journal of Political Economy*, 78(6), 1228–1263.

Barsky, R., and E. Sims (2011): "News Shocks and Business Cycles," *Journal of Monetary Economics*, 58(3), 273–289.

Baumol, W. J. (1952): "The Transactions Demand for Cash: An Inventory Theoretic Model," *Quarterly Journal of Economics*, 66, 545–372.

Benati, L. (2008): "Investigating Inflation Persistence Across Monetary Regimes," *Quarterly Journal of Economics*, 123(3), 1005–1060.

Benati, L., R. E. LucasJr., J.-P. Nicolini, and W. Weber (2019): "International Evidence on Long-Run Money Demand," *Minneapolis FED Working Paper 587, June 2019*.

——— (2021): "International Evidence on Long-Run Money Demand," *Journal of Monetary Economics*, 117(1), 43–63.

Bresciani-Turroni, C. (1937): *The Economics of Inflation.* John Dickens and Co. Northampton.

Bruno, M., and S. Fischer (1990): "Seigniorage, Operating Rules, and the High Inflation Trap," *Quarterly Journal of Economics*, 105(2), 353–374.

Burmeister, E., and K. D. Wall (1982): "Kalman Filtering Estimation of Unobserved Rational Expectations with an Application to the German Hyperinflation," *Journal of Econometrics*, 20, 255–284.

——— (1987): "Unobserved Rational Expectations and the German Hyperinflation with Endogenous Money Supply," *International Economic Review*, 28(1), 15–32.

CAGAN, P. (1956): "The Monetary Dynamics of Hyperinflation," *in M. Friedman, ed., Studies in the Quantity Theory of Money, University of Chicago Press*, pp. 25–120.

CAVALIERE, G., A. RAHBEK, AND A. M. R. TAYLOR (2012): "Bootstrap Determination of the Cointegration Rank in Vector Autoregressive Models," *Econometrica*, 80(4), 1721–1740.

CHRISTIANO, L. J. (1987): "Cagan's Model of Hyperinflation Under Rational Expectations," *International Economic Review*, 28(1), 33–48.

CORANA, A., M. MARCHESI, C. MARTINI, AND S. RIDELLA (1987): "Minimizing Multimodal Functions of Continuous Variables with the Simulated Annealing Algorithm," *ACM Transactions on Mathematical Software*, 13, 262–80.

DRAZEN, A. (1985): "A General Measure of Inflation Tax Revenues," *Economics Letters*, 17, 327–330.

FLOOD, R. P., AND P. M. GARBER (1980): "Market Fundamentals versus Price-Level Bubbles: The First Tests," *Journal of Political Economy*, 88(4), 745–770.

GELMAN, A., J. CARLIN, H. STERN, AND D. RUBIN (1995): *Bayesian Data Analysis*. Chapman and Hall, New York.

GEWEKE, J. (1992): "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments," *in J. M. Bernardo, J. Berger, A. P. Dawid and A. F. M. Smith (eds.), Bayesian Statistics, Oxford University Press, Oxford*, pp. 169–193.

GRAHAM, F. D. (1930): *Exchange, Prices, and Production in Hyperinflation*. Princeton University Press.

HAMILTON, J. (1994): *Time Series Analysis*. Princeton, NJ, Princeton University Press.

HAMILTON, J. D. (1985): "Uncovering Financial Market Expectations of Inflation," *Journal of Political Economy*, 93(6), 1224–1241.

HARVEY, A. (1989): *Forecasting, Structural Time Series Models, and the Kalman Filter*. Cambridge University Press.

JOHANSEN, S. (2002): "A Small Sample Correction for the Test of Cointegrating Rank in the Vector Autoregressive Model," *Econometrica*, 70(5), 1929–1961.

LERNER, E. (1956): "Inflation in the Confederacy, 1861-65," *in M. Friedman, ed., Studies in the Quantity Theory of Money, University of Chicago Press.*

LUCASJR., R. E. (2000): "Inflation and Welfare," *Econometrica*, 68(2), 247–274.

MCINDOE-CALDER, T. (2018): "Hyperinflation in Zimbabwe: Money Demand, Seigniorage and Aid Shocks," *Applied Economics*, 50(15), 1659–1675.

MELTZER, A. H. (1963): "The Demand for Money: The Evidence from the Time Series," *Journal of Political Economy*, 71(3), 219–246.

PETROVIC, P., AND Z. MLADENOVIC (2000): "Money Demand and Exchange Rate Determination under Hyperinflation: Conceptual Issues and Evidence from Yugoslavia," *Journal of Money, Credit and Banking*, 32(4), 785–806.

SALEMI, M. K., AND T. J. SARGENT (1979): "The Demand for Money During Hyperinflation under Rational Expectations: II," *International Economic Review*, 20(3), 741–758.

SARGENT, T. J. (1977): "The Demand for Money during Hyperinflations under Rational Expectations: I," *International Economic Review*, 18(1), 59–82.

——— (1982): "The End of Four Big Inflations," *in R. E. Hall, ed., 'Inflation: Causes and Effects', University of Chicago Press*, pp. 41–97.

SARGENT, T. J., AND F. R. VELDE (1995): "Macroeconomic Features of the French Revolution," *Journal of Political Economy*, 103(3), 474–518.

SARGENT, T. J., AND N. WALLACE (1973): "Rational Expectations and the Dynamics of Hyperinflation," *International Economic Review*, 14(2), 328–350.

——— (1987): "Inflation and the Government Budget Constraint," *in Assaf Razin and Efraim Sadka (editors), Economic Policy in Theory and Practice, MacMillan Press.*

STOCK, J. H., AND M. W. WATSON (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61(4), 783–820.

TAYLOR, M. P. (1991): "The Hyperinflation Model of Money Demand Revisited," *Journal of Money, Credit and Banking*, 23(3), 327–351.

TOBIN, J. (1956): "The Interest Elasticity of Transactions Demand for Money," *Review of Economics and Statistics*, 38, 241–247.

WHITE, E. N. (1990): "The Stock of Paper Money During the French Revolution," Working Paper N. 2 of the Financial and Monetary History of the French Revolution Project, January 1988.

——— (1991): "Measuring the French Revolution's Inflation: the Tableaux de dépréciation," *Histoire et Mesure*, 6(3), 245–274.

Wu, G. (1958): *The Historical Materials of Inflation in Old China*. Shanghai Renmin Publishing House.

Zhao, L. (2017): "The Behavior of Money Demand in the Chinese Hyperinflation," *China Economic Review*, 42, 145–154.

Zhao, L., and L. Li (2015): "Interest Rate, Money Demand and Seigniorage: The Chinese Hyperinflation 1946-1949," *China Economic Review*, 34, 169–179.

Appendix

A Benati, Lucas, Nicolini, and Weber's (2021) Model of the Transaction Demand for Money

The model of the transaction demand for money I briefly summarize in Section 2 of the main text is the one analyzed in Benati, Lucas, Nicolini, and Weber (2017, 2021). The model features a labor-only, representative agent economy in which making transactions is costly. Preferences are described by

$$E_0 \sum_{t=0}^{\infty} \beta^t U(x_t),$$

where $\beta < 1$, x_t is consumption at date t, and the function U is differentiable, increasing in its argument, and concave. In each period the agent is endowed with a unit of time, with l_t allocated to goods production and 1-l_t used to carry out transactions. The goods production technology is given by $y_t = x_t = z_t l_t$, where z_t is an exogenous stochastic process. Households choose the number n_t of 'trips to the bank' as in the classic Baumol-Tobin model. At the beginning of a period, a household begins with some nominal wealth that can be allocated to money M_t or to risk-free government bonds B_t. During the first of the n_t subperiods, one member of the household uses money to buy consumption goods. During this same initial subperiod, another member of the household produces and sells goods in exchange for money. At the end of the subperiod, producers transfer to the bank the proceeds from their transactions. Thus, the situation at the beginning of the second subperiod exactly replicates the situation at the beginning of the first. This process is repeated n_t times during the period. The choice of this variable n_t is the only economically relevant decision made by households. Purchases over a period are subject to a cash-in-advance constraint, $P_t x_t \leq M_t n_t$.

Notice that n_t is the velocity of money, and its inverse in equilibrium is the money-to-output ratio, or the demand for real money balances. Baumol and Tobin assumed that the cost of carrying out these transactions increases linearly in the number of trips. Benati et al. (2017, 2021) consider more general specifications in which the total cost of making transactions, measured in units of time, is given by either

$$\theta(n_t) = \psi n_t^{\sigma} \tag{A.1}$$

where ψ and σ are positive constants, or

$$\theta(n_t) = \frac{k + b \ln \varepsilon}{\varepsilon} - \frac{k + b \ln(n_t + \varepsilon)}{n_t + \varepsilon}. \tag{A.2}$$

Expression (A.1) becomes the Baumol-Tobin linear case when $\sigma = 1$. In expression (A.2) it is assumed that $k > b\,(1 - \ln \varepsilon)$, so that the function is always increasing in n_t. This function is also concave as the one before. The main difference between (A.2) and (A.1) is that it asymptotes a constant as the number of trips grows arbitrarily large.

When the cost of making transactions is given by (A.1), Benati et al. (2021) show that the solution for n_t is

$$\sigma\psi\frac{n_t^{\sigma+1}}{1-\psi n_t^{\sigma}}=R_t \tag{A.3}$$

where R_t is the nominal interest rate. Note that ψn_t^{σ} is the cost of inflation in units of time, and therefore it represents the welfare cost of inflation as a ratio of maximum potential output. Taking logs we obtain

$$\ln\sigma\psi+\ln n_t^{\sigma+1}-\ln(1-\psi n_t^{\sigma})=\ln R_t. \tag{A.4}$$

As discussed in the main text of the paper, at low inflation rates the welfare costs of inflation are negligible, so that $1-\psi n_t^{\sigma}\simeq 1$, $\ln(1-\psi n_t^{\sigma})\simeq 0$, and (A.4) therefore becomes

$$\ln\frac{1}{n_t}=\ln\frac{M_t}{P_tY_t}\simeq\frac{1}{\sigma+1}\left[\ln\sigma\psi-\ln R_t\right]. \tag{A.4}$$

which is Meltzer's log-log with elasticity $1/(\sigma+1)$. The Baumol-Tobin case is the one obtained by assuming that $\sigma=1$, which implies an interest rate elasticity of $1/2$.

By the same token, when the cost of making transactions is given by (A.2) the solution is

$$\frac{n_t^2}{(n_t+\varepsilon)^2}\frac{\left[b\left(\ln(\varepsilon+n_t)-1\right)+k\right]}{1-\theta(n_t)}=R_t. \tag{A.5}$$

Ignoring as before the term $1-\theta(n_t)$, and considering relatively small values for ε produces a linear relationship between the log of velocity and the interest rate,

$$\ln\frac{1}{n_t}=\ln\frac{M_t}{P_tY_t}\simeq\frac{k-b}{b}-\frac{1}{b}R_t. \tag{A.6}$$

which corresponds to the semi-log specification.

B The Data

Here follows a detailed description of the data and of their sources.

B.1 The data for hyperinflations

B.1.1 Monthly data

Data for the French Revolution have been generously provided by Eugene White. A monetary aggregate labelled as 'Total *assignats* in circulation less demonetized issues'[31] is from Table 9 of White (1990). The corresponding price index, labelled as 'French Treasury exchange rate: market rate. French paper *assignats* per gold French *livre*', is from Table 2 of White (1991), and it represents the 'conversion rate' of paper *assignats* which had been issued at a specific date into gold French *livre*, meaning that, in fact, this was an *assignats*-specific price index.[32]

Cagan's (1956) data for Austria, Germany, Hungary (both post-WWI and post-WWII), Poland, Russia and Greece are from his Appendix B.[33]

Barro's (1970) data for Austria, Germany, post-WWI Hungary, and Poland are from Tables A1-A4 in the Appendix, and feature the logarithm of real money balances and inflation, computed as the log-difference of the relevant price index.

Graham's data for Germany are from Table XII of Graham (1930), and feature indices of wholesale prices and of total monetary circulation.

The data for China are from Zhao and Li (2015), and they have been kindly provided by L. Zhao. As detailed in Zhao and Li (2015), the data for currency are from Wu (1958, pp. 92 and 122), whereas the price index is from Wu (1958, pp. 160-163).

For Chile and Argentina, data for M1 and the CPI are from the *Banco Central do Chile* and the *Banco Central de la República Argentina*, respectively.

For Bolivia, data for M1 are from Bolivia's central bank, and they have been kindly provided by Carlos Gustavo Machicado, whereas data for the CPI are from the *Instituto Nacional de Estadistica*.

For Brazil the CPI is from the *Instituto Brasileiro de Geografia e Estatística* (IBGE), whereas M1 is from the *Banco Central do Brazil*.

For Peru the CPI is from the *Banco Central de Reserva del Peru*, whereas a monetary aggregate defined as 'Money plus quasi money' is from the International Monetary Fund's *International Financial Statistics* (henceforth, *IMF* and *IFS*).

The data for Yugoslavia—a retail price index (RPI), M1, and the black market exchange rate of the Yugoslav dinar *vis-à-vis* the Deutsche Mark—are from Petrovic and Mladenovic (2000; henceforth PM), and they were kindly provided by Z. Mlade-

[31] To be clear, what this label means is that the stock was computed as the sum of the total amount of *assignats* which had been issued by the Revolutionary government, minus the amount which had been retired from circulation and destroyed (as it was periodically done).

[32] I ignore the other currency issued by the Revolutionary government, the *mandat*, because between February 1796, when the *mandats* were first issued, and June 1796, when my sample ends, the stock of *mandats* consistently represented a tiny fraction of the stock of *assignats* (ranging between 0.06 per cent in February, and 3.8 per cent in June).

[33] Cagan's dataset features the logarithms in base 10 of (P_t/M_t) and of (P_t/P_{t-1}), where M_t and P_t are the nominal money stock and the price level, respectively, for month t. I converted the original data to natural logarithms. Then, I computed the index of log prices as the cumulative sum of the log-difference of the price level, and based on this, and the logarithm of real money balances, I recovered an index for log nominal money balances.

novic.[34] Since, as discussed by PM,[35] several observations of the RPI are unreliable, in what follows I will work with the exchange rate, which, as discussed by PM in footnote 10, was collected directly by them from daily newspapers.

Data for Congo, Angola, and Bulgaria are from the *IMF*'s *IFS*. For Congo they feature a series for 'Money' and the exchange rate *vis-à-vis* the U.S. dollar; for Angola the central bank's 'Reserve Money' and the CPI; for Bulgaria, M1 and the CPI.

For Zimbabwe, a series for 'Reserve Money' is from the *IMF*'s *IFS*, whereas the black-market exchange rate of Zimbabwe's dollar *vis-à-vis* the U.S. dollar, which was used in McIndoe-Calder (2018), has been generously provided by Tara McIndoe-Calder. As discussed by McIndoe-Calder (2018, Section III, pp. 1661-1663) the official CPI series (available from the *IMF*'s *IFS*) is unreliable, and in what follows I will therefore exclusively focus on the black-market exchange rate.

For Venezuela, data for M1 are from the *Banco Central de Venezuela*. As for the price index, the government stopped publishing official CPI figures in December 2015, and it resumed publishing them in June 2019. Although, strictly speaking, an official monthly CPI series is available on a continuous basis since December 2007, I have preferred to resort, once again, to the black-market exchange rate (in the present case, for the Bolivàr *vis-à-vis* the U.S. dollar), which is available at the daily frequency at the website https://dolartoday.com.[36]

B.1.2 Weekly data

As for Germany, a weekly series for the money stock (labelled as 'Notenumlauf', i.e. 'Banknotes in circulation'), available from December 14, 1918 to November 15, 1923, is from Flood and Garber's (1980) Table B.1 in Appendix B until the end of December 1922—with the original source of the data being *Wirtschaft und Statistik* (*WS*)[37]—and it is from *WS* after that.[38] A *daily* series for the spot exchange rate of the German Reichsmark *vis-à-vis* the British Pound is available nearly without interruptions from September 7, 1922 to November 15, 1923 from *WS*. An important point to stress is that since, with a couple of exceptions, this series is available for each single business day during this period, I can almost always *exactly* match the dates

[34] The data's original sources are discussed in detail in footnote 10 of Petrovic and Mladenovic (2000).

[35] See p. 787, and especially footnote 4.

[36] The key reason for doing so is that for a sizeable portion of the sample period the profile of the official CPI series is materially different from that of the black-market exchange rate, and among the two there are probably good reasons for putting more trust in the latter, rather than in the official statistics produced by the government. (E.g., for Argentina, it is worth recalling that under Christina Kirchner's government *The Economist* stopped reporting the official CPI statistics, on the grounds that, as it was widely known, they were being manipulated.)

[37] See at: https://www.destatis.de/GPStatistik/receive/DESerie_serie_00000012?list=all

[38] The original series contains a periodic pattern at the monthly frequency (so that in the last day of the month the series temporarily increases compared to adjacent observations), which I removed it *via* ARIMA X-12.

in which the monetary aggregate had been released with the dates for the exchange rate. A weekly price series (labelled as 'Großhandelindexziffer', i.e. 'Wholesale price index'), available from August 7, 1922 to November 15, 1923, is from *WS*. A limitation of this series, compared to the series for the Reichsmark/Pound exchange rate, is that it had been released at the weekly frequency on dates which sometimes did not match the release dates for the monetary aggregates. As a result, a dataset comprising this series and the 'Notenumlauf' monetary aggregate suffers from the shortcoming that the two series are not exactly matched on a day-by-day basis. Because of this, in what follows I will almost exclusively focus on the results based on the exchange rate, and I will largely eschew those based on the wholesale price index.

As for post-WWII Hungary, two series for a monetary aggregate (labelled as 'Notes') and a price index are from the appendix of Anderson, Bomberger, and Makinen (1988).

B.2 Long-run data

B.2.1 Annual data

All of the annual data for Brazil, Argentina, Chile, and Bolivia shown in Figure 9 are from the dataset assembled by Benati, Lucas, Nicolini, and Weber (2021).

B.2.2 Quarterly data

The sources of the quarterly data for Israel and Mexico used for Figure 3 are as follows.

For Israel, a seasonally adjusted series for nominal GDP is from the International Monetary Fund's International Financial Statistics. A seasonally adjusted monthly series for M1 is from the Bank of Israel, and it has been converted to the quarterly frequency by taking averages within the quarter. A monthly series for the central bank's discount rate is from the International Monetary Fund's International Financial Statistics, and it has been converted to the quarterly frequency by taking averages within the quarter.

For Mexico, a seasonally adjusted series for nominal GDP is from INEGI. A seasonally unadjusted series for M1 is from Banco de México (Mexicos' central bank), and it has been seasonally adjusted via ARIMA X-12. A monthly series for a short-term rate (91 day Cetes) is from Banco de México, and it has been converted to the quarterly frequency by taking averages within the quarter.

C Computing Seignorage Based on Discretely Sampled Observations for Money and Prices[39]

What is the most reliable way of computing seignorage based on discretely sampled observations for money and the price level? In continuous time there is no ambiguity about how to compute it: by defining the money stock and the price level as M_t and P_t, the instantaneous revenue from money creation is given by expression (1) in the main text.[40] Working with observations for M_t and P_t sampled at discrete intervals, on the other hand, things are not clear-cut, since the explosive dynamics that typically characterize hyperinflations causes alternative, and apparently equally sensible ways of computing seignorage to often produce materially different estimates. This is the case, in particular, for the latest—and typically most extreme—stages of hyperinflations, i.e. precisely those which are most informative about the presence, or absence, of a Laffer curve for seignorage.

As a simple illustration, consider the following discrete-time version of (1), which as I will show works well under a wide range of empirically plausible circumstances:

$$\xi_t^* \equiv \theta_t^* \left[\left(\frac{M_{t-1}}{P_{t-1}} \right)^{\omega} \left(\frac{M_t}{P_t} \right)^{1-\omega} \right] \tag{C.1}$$

where $\theta_t^* \equiv \ln(M_t)$-$\ln(M_{t-1})$=$m_t$-$m_{t-1}$, and $\omega \in [0,1]$. I what follows I will work with ω=0.5, although, as I discuss in footnote 42 below, I have also experimented with alternative values. Table C.1 reports, for Yugoslavia and Zimbabwe, empirical measures of seignorage based on (C.1) for the month corresponding to the inflation peak, as well as for 6 and 12 months before the peak, together with the inflation rate, computed as the log-difference of the price level, i.e. $\pi_t \equiv \ln(P_t)$-$\ln(P_{t-1})$=p_t-p_{t-1}. The table also reports the corresponding measures based on an alternative expression found in the literature:[41]

$$\tilde{\xi}_t \equiv \frac{M_t - M_{t-1}}{\frac{1}{2}\left(P_t + P_{t-1}\right)}. \tag{C.2}$$

In either case, seignorage measures have been rescaled so that, 12 months before the inflation peak, they are equal to 1.

[39] This Appendix is conceptually in line with Bresciani-Turroni's (1937, pp. 146-154) analysis of the superiority of geometric averages, compared to their arithmetic counterparts, within the context of Germany's hyperinflation. A key difference is that whereas Bresciani-Turroni considered numerical examples based on actual data, I work based on continuous-time theoretical models.

[40] Working in continuous time, Drazen (1985) presents a measure of the revenue from money creation which is conceptually correct across alternative models. Drazen's measure—see his equation (5)—is equal to (in my notation) the sum of $\theta_t(M_t/P_t)$ in (1), and $(r_t$-$n_t)a_t$, where r_t, n_t, and a_t are the real interest rate, population growth, and 'the (per capita real) value of assets held by government by virtue of people holding real balances' (see Drazen, 1985, p. 328). Since $(r_t$-$n_t)$ is negligible compared to θ_t, and a_t is of the same order of magnitude of M_t/P_t, it logically follows that Drazen's measure is, for all practical purposes, near-identical to $\theta_t(M_t/P_t)$.

[41] See e.g. Sargent and Velde (1995, p. 506).

Table C.1 Empirical measures of seignorage for Yugoslavia and Zimbabwe based on alternative discrete-time approximations						
	Yugoslavia			*Zimbabwe*		
	π_t	ξ_t^*	$\tilde{\xi}_t$	π_t	ξ_t^*	$\tilde{\xi}_t$
12 months before inflation peak	0.59	1	1	1.00	1	1
6 months before inflation peak	1.22	0.56	0.52	0.92	0.60	0.61
Inflation peak	11.29	1.11	0.16	3.91	1.41	0.57

The table clearly illustrates the problem: whereas ξ_t^* and $\tilde{\xi}_t$ are very close 6 months before the inflation peak, for the month associated with the peak they diverge quite significantly. Further, for the most extreme episode, Yugoslavia (and to a much lesser extent for Zimbabwe), ξ_t^* points towards a *monotonically increasing* relationship between inflation and seignorage, whereas $\tilde{\xi}_t$ points towards a *Laffer curve*, with the inflation explosion over the last 6 months leading to the peak being associated with a decrease in seignorage. These results, which are representative of the overall set of 20 episodes,[42] suggest that whereas the two measures tend to produce similar results for the comparatively milder stages of hyperinflations, for the most extreme, and therefore most informative stages they sometimes produce significantly different estimates. This raises the obvious question of which, between ξ_t^* and $\tilde{\xi}_t$, should be regarded as the most reliable measure of the revenue from money creation. Or is it the case that, in fact, neither of them provides a reliable approximation to the true amount of seignorage collected by the government?

C.1 Approximations based on discretely sampled observations

In order to address this issue, in this appendix I explore how to best estimate the amount of seignorage raised by the government between two specific points in time, t_1 and t_2, based on observations for money and the price level corresponding to those specific dates, i.e. $[M_{t_1}\ M_{t_2}]'$ and $[P_{t_1}\ P_{t_2}]'$. The *true* amount of seignorage collected between t_1 and t_2 is defined as the integral of (1) over this interval:

$$\xi_{t_1 \to t_2}^{\text{TRUE}} \equiv \int_{t_1}^{t_2} \theta_t \frac{M_t}{P_t} dt \tag{C.3}$$

As for the evolution of real money balances as a function of expected inflation, I consider both Cagan's (1956) semi-log money demand specification,

$$\ln\left(\frac{M_t}{P_t}\right) = \beta + \alpha \pi_t^e, \tag{C.4}$$

[42] The entire set of results is available upon request.

and Meltzer's (1963) log-log,

$$\ln\left(\frac{M_t}{P_t}\right) = \beta + \alpha \ln(\pi_t^e) \tag{C.5}$$

where π_t^e is expected inflation, and $\alpha < 0$ and β are constants. In order to explore which discrete-time approximation works best across a range of possible empirical scenarios, I consider three alternative assumptions about the evolution of inflation, i.e. either that it is constant, or that it evolves according to either a linear or an exponential time trend.[43]

C.1.1 A steady-state with constant inflation, seignorage, and money growth

With $\pi_t \equiv d\ln P_t/dt = dp_t/dt = \bar{\pi}$, the logarithm of the price level evolves according to $p_t = p_0 + \bar{\pi}t$, where p_0 is an integration constant, so that its level is given by $P_t = P_0 \exp(\bar{\pi}t)$. With π_t^e=π_t=$\bar{\pi}$, from either (C.1) or (C.2) real money balances, and therefore money growth, are also constant, and they are equal to $(M_t/P_t) = (M_0/P_0) \equiv \exp(\bar{\mu})$ and $\theta_t = \bar{\pi}$, respectively. Finally, from $\theta_t = \bar{\pi}$ the logarithm and the level of the money stock evolve according to $m_t = m_0 + \bar{\pi}t$ and $M_t = M_0 \exp(\bar{\pi}t)$. The instantaneous value of seignorage is therefore equal to $\xi_t = \bar{\pi}\exp(\bar{\mu})$, so that the amount of seignorage raised by the government between t-1 and t is equal to

$$\xi_{t\text{-}1\to t}^{\text{TRUE}} \equiv \int_{t\text{-}1}^{t} \theta_t \frac{M_t}{P_t} dt = \int_{t\text{-}1}^{t} \bar{\pi}\exp(\bar{\mu})dt = \exp(\bar{\mu})\bar{\pi} \tag{C.6}$$

It can be trivially shown that ξ_t^* is indeed equal to $\xi_{t\text{-}1\to t}^{\text{TRUE}}$, whereas

$$\tilde{\xi}_t \equiv \exp(\bar{\mu}) \underbrace{\frac{\exp(\bar{\pi}) - 1}{\frac{1}{2}[\exp(\bar{\pi}) + 1]}}_{\phi(\bar{\pi})} = \exp(\bar{\mu})\phi(\bar{\pi}) \tag{C.7}$$

A comparison between (C.6) and (C.7) shows that the ability of the latter to reliably approximate the former hinges on the extent to which $\phi(\bar{\pi})$ is, or is not, sufficiently close to $\bar{\pi}$. Whereas for $\bar{\pi}$=0.5—Cagan's threshold at the monthly frequency for an episode to be classified as a hyperinflation—$\phi(\bar{\pi})$=0.490, progressively higher values of $\bar{\pi}$ quickly lead to large and increasing (in absolute value) negative deviations of $\phi(\bar{\pi})$ from $\bar{\pi}$. E.g., for $\bar{\pi}$=1.313 (Venezuela's maximum inflation rate so far) $\phi(\bar{\pi})$=1.152,

[43] I define the alternative scenarios in terms of the evolution of inflation—rather than seignorage, or money growth—for the following reasons. Since seignorage is what we aim to estimate, it is not possible to state that (e.g.) a scenario with constant seignorage is more realistic, from the perspective of the 20 episodes I analyze, than one in which seignorage evolves according to an exponential trend. To put it differently, in order to be able to assess which scenarios are more, and which are less realistic, we ought to focus on a variable which is unambiguously observed. Under this respect, the key advantage of inflation over money growth is that the analysis is significantly simpler from a mathematical point of view.

corresponding to a 12.3 per cent underestimation of seignorage compared to its true value. For $\bar{\pi}$ equal to Zimbabwe's maximum inflation rate of 3.912, however, the extent of underestimation already reaches 50.9 per cent, whereas for $\bar{\pi}$=5.885 (Germany's inflation peak) $\phi(\bar{\pi})$=1.989, with an underestimation of 66.2 per cent. For $\bar{\pi} \rightarrow \infty$, $\phi(\bar{\pi}) \rightarrow 2$, so that underestimation tends to 100 per cent: in fact, such an extent of underestimation already obtains for inflation rates corresponding, e.g., to Yugoslavia's peak of 11.29. More generally, underestimation is non-negligible to substantial (or much worse) for any inflation rate greater than about 1.5. The implication is that even if most (or all) episodes could be characterized as fluctuations around a steady-state with constant inflation, in the vast majority of cases $\tilde{\xi}_t$ would provide a poor, or very poor estimate of the true amount of seignorage raised by the government.

For many episodes, however, the notion that inflation may have been fluctuating around a constant value appears as hardly reasonable. For the countries in Cagan's dataset, for example, this is the case only for the three comparatively milder episodes (post-WWI Hungary, Austria, and Russia), whereas for the remaining four cases inflation exhibits a clear upward trend. Evidence for the remaining episodes is qualitatively the same.

In the next two sub-sections I therefore extend the previous analysis by considering two alternative specifications for the inflation trend, first under the assumption of perfect foresight, π_t^e=π_t (i.e. the equivalent, within the present context, to rational expectations), and then under the assumption of adaptive expectations.

C.2 The case of perfect foresight

C.2.1 A linear time trend for inflation

With inflation following the linear time trend $\pi_t \equiv \pi_0 + \gamma t$—where π_0 and $\gamma > 0$ are constants—p_t and P_t evolve according to $p_t = p_0 + \pi_0 t + \gamma t^2/2$ and $P_t = P_0 \exp(\pi_0 t + \gamma t^2/2)$. The fact that inflation is monotonically increasing implies that, based on either (C.4) or (C.5), real money balances are progressively decreasing, so that different from the previous sub-section the analysis has to be performed for either of the two money demand specifications.

Working as before it can be shown, after some math, that for the semi-log specification θ_t=dm_t/dt=π_0+$\alpha\gamma$+γt, and (M_t/P_t)=exp(β+$\alpha\pi_0$+$\alpha\gamma t$), so that the amount of seignorage raised by the government between t-1 and t is equal to

$$\xi_{t\text{-}1\rightarrow t}^{\text{TRUE}} = \left(1+\frac{\pi_0}{\alpha\gamma}\right)[\exp(\alpha\gamma)\text{-}1]\exp(\beta+\alpha\pi_0)\exp[\alpha\gamma(t\text{-}1)]+$$

$$+\frac{1}{\alpha}\exp(\beta+\alpha\pi_0)\exp[\alpha\gamma(t\text{-}1)]\left[t\exp(\alpha\gamma)\text{-}(t\text{-}1)\text{-}\frac{\exp(\alpha\gamma)\text{-}1}{\alpha\gamma}\right] \qquad \text{(C.7)}$$

whereas for the log-log it is equal to

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \exp(\beta)\frac{(\pi_0+\gamma t)^{\alpha+2}\text{-}[\pi_0+\gamma(t\text{-}1)]^{\alpha+2}}{\gamma(\alpha+2)}+$$

$$+\exp(\beta)\left\{(\pi_0+\gamma t)^{\alpha}\text{-}[\pi_0+\gamma(t\text{-}1)]^{\alpha}\right\} \tag{C.8}$$

I now turn to the case of an exponential time trend for inflation.

C.2.2 An exponential time trend for inflation

With inflation following the exponential trend $\pi_t \equiv \pi_0 \exp(\gamma t)$—where, as before, π_0 and $\gamma > 0$ are constants—p_t and P_t evolve according to $p_t = p_0 + (\pi_0/\gamma)\exp(\gamma t)$ and $P_t = P_0 \exp[(\pi_0/\gamma)\exp(\gamma t)]$. Working as as before it can be shown that for the semi-log specification the amount of seignorage raised by the government between t-1 and t is equal to

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \exp(\beta)\pi_0(1+\alpha\gamma)\int_{t\text{-}1}^{t}\exp[\gamma\tau+\alpha\pi_0\exp(\gamma\tau)]d\tau \tag{C.9}$$

whereas for the log-log it is equal to

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \exp(\beta)\pi_0^{\alpha}[\exp(\alpha\gamma)-1]\exp[\alpha\gamma(t\text{-}1)]+$$

$$+\exp(\beta)\pi_0^{1+\alpha}\frac{\exp[(1+\alpha)\gamma]\text{-}1}{(1+\alpha)\gamma}\exp[(1+\alpha)\gamma(t\text{-}1)] \tag{C.10}$$

For given values of the structural parameters α, β, γ, and π_0 the integral in (C.9) can easily be computed numerically.

C.2.3 Evidence

Given how convoluted expressions (C.7)-(C.10) are, different form appendix C.1.1 it is not possible to assess in a straightforward manner the extent to which either (C.1) or (C.2) may or may not provide a good approximation to the true amount of seignorage raised by the government. Conditional on specific parametes values, however, such an assessment can be easily performed.

Figure A.1 reports results from the following exercise. I set π_0=0.5 and β=1.[44] As for γ, I set it to 0.1 for the linear trend, and to 0.07 for the exponential one, so that in either case inflation reaches a value of about 3 after 25 periods.[45] Finally, as for α I consider the following values. For the semi-log specification, based on the

[44] As mentioned, 0.5 is Cagan's threshold at the monthly frequency for an episode to be classified as a hyperinflation. As for β, I experimented with several alternative values, and they all produced qualitatively the same results reported herein.

[45] This appears as broadly plausible based on the values for the maximum inflation rates reported in Table 1 in the main text.

range of estimates reported in Sargent (1977), I set α to either -0.5, -1, -2, -4, or -6. As for the log-log I consider Baumol and Tobin's benchmark value of -0.5, as well as either smaller or larger values, specifically -0.25, -0.75, -1, and -1.25. An elasticity of -1 is of particular interest because, for α<-1, an increase in inflation is associated with a decrease in seignorage for *all* inflation levels. Finally, I also consider α=-1.25 because, as I discussed in Section 8 in the main text, for the most extreme episodes there is some evidence that the elasticity of money demand might have been below -1. Figure A.1 reports, for either inflation trend, either money demand functional form, and either of the two expressions for computing seignorage based on discretely sampled data—i.e., either (C.1) or (C.2)—the approximation error, defined as the percentage deviation of either ξ_t^* or $\tilde{\xi}_t$ from $\xi_{t\text{-}1\to t}^{\text{TRUE}}$,[46] with, e.g., -0.01 meaning minus 1 per cent. In the figure I label ξ_t^* as the 'geometric average', and $\tilde{\xi}_t$ as the 'arithmetic average'.

The following main results emerge from Figure A.1:

(1) consistent with appendix C.1.1, the approximation based on $\tilde{\xi}_t$ is uniformly poor, or very poor for any inflation rate beyond about 1.5.

(2) Based on ξ_t^* the approximation errors are negligible in three cases out of four. In particular, based on the log-log functional form this is the case for either specification for the inflation trend. This is crucial because, as shown in Section 5 in the main text, evidence clearly suggests that the log-log provides the most plausible description of the data.

(3) For the single instance in which ξ_t^* provides a mostly poor approximation (i.e., the case in which inflation follows an exponential trend, and money demand takes the semi-log form) it is to be noticed that (*i*) seignorage gets uniformly *under*estimated, and (*ii*) for any value of α, the magnitude of the approximation error is (in absolute value) monotonically *increasing* with inflation. The implication is that the presence of such approximation error tends to produce spurious evidence of a Laffer curve (to be precise, since we are here dealing with the semi-log, the approximation error tends to magnify the evidence of a Laffer curve which is a mathematical property of such specification). This means that *in no way* the absence of a Laffer curve in the data for π_t and ξ_t^* documented in Section 4 in the main text can be ascribed to the possible presence of an approximation error.[47]

As shown in Section 5 in the mai text the log-log functional form provides a significantly more plausible description of the data than the semi-log. This implies that, as the second column of Figure A.1 shows, ξ_t^* provides an excellent approximation to the true amount of seignorage for either of the two specifications for the inflation trend. Since, as previously shown, this is also the case for a constant-inflation steady

[46] I.e., $(\xi_t^*\text{-}\xi_{t\text{-}1\to t}^{\text{TRUE}})/\xi_{t\text{-}1\to t}^{\text{TRUE}}$ and $(\tilde{\xi}_t\text{-}\xi_{t\text{-}1\to t}^{\text{TRUE}})/\xi_{t\text{-}1\to t}^{\text{TRUE}}$, respectively.

[47] In an attempt to improve the quality of the approximation within this specific instance, I have also experimented with values of ω different from 0.5. Values around 0.35-0.4 lead to non-negligible improvements for values of α greater than about -0.8, whereas the improvement for smaller values of α is essentially negligible. These results are available upon request.

state, the implication is that under perfect foresight ξ_t^* provides a reliable estimate of the true amount of seignorage for any empirically plausible scenario.

C.3 The case of adaptive expectations

As pointed out by several authors,[48] within the context of hyperinflations the assumption of rational expectations (here, perfect foresight) may be too extreme, for two reasons. First, the often explosive nature of the process—i.e., its extraordinary speed, with inflation typically moving from near price stability to hyperinflation within a matter of months—should logically be expected to have caught agents at least partially by surprise. Second, in most episodes the hyperinflation had been historically unprecedented for that country. A stark illustration is represented by the Weimar Republic episode: up until WWI, Germany had been operating for centuries under commodity standards—over the most recent decades, under a Gold Standard—and it had therefore experienced the extreme price stability that is a hallmark of such regimes. It is therefore highly implausible that the German public might have had rational expectations about the Weimar Republic episode right from the beginning of the hyperinflation.

In this appendix I therefore extend the previous analysis to the case in which inflation expectations evolve according to the adaptive expectations mechanism

$$\frac{d\pi_t^e}{dt} = \delta(\pi_t^e - \pi_t) \tag{C.11}$$

The case of perfect foresight is the limit of (C.11) for $\delta \rightarrow \infty$.

C.3.1 A linear time trend for inflation

As in appendix C.2.1 inflation evolves according to $\pi_t = \pi_0 + \gamma t$, and p_t and P_t evolve according to $p_t = p_0 + \pi_0 t + \gamma t^2/2$ and $P_t = P_0 \exp(\pi_0 t + \gamma t^2/2)$. I conjecture that π_t^e also evolves according to a linear time trend, $\pi_t^e = A + Bt$, with coefficients A and B to be determined. It can easily be shown that A and B are equal to A=π_0-γ/δ and B=γ, so that $\pi_t^e = \pi_0 - \gamma/\delta + \gamma t$. Working as before, it can be shown that for the semi-log m_t and M_t evolve according to m_t=m_0+$(\pi_0$+$\alpha\gamma)t$+$\gamma t^2/2$, with m_0=β+$\alpha\pi_0$+p_0-$\alpha\gamma/\delta$, and M_t=$M_0 \exp[(\pi_0 + \alpha\gamma)t + \gamma t^2/2]$, so that θ_t=dm_t/dt=π_0+$\alpha\gamma$+γt, and (M_t/P_t)=$\exp(\beta$+$\alpha\pi_0$-$\alpha\gamma/\delta$+$\alpha\gamma t)$. A comparison with appendix C.2.1 shows that the expression for money growth is identical, whereas the only difference between the two expressions for real money balances is the term -$\alpha\gamma/\delta$ within the exponential function. Since α<0, this implies that, as long as δ<∞, real money balances are greater than the corresponding level under perfect foresight, so that for any level of inflation the government collects a strictly greater amount of seignorage than under

[48] See e.g. Fischer's (1987) discussion of Sargent and Wallace (1987).

perfect foresight. Working as before it can be shown that for the semi-log specification the amount of seignorage raised by the government between t-1 and t is equal to

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \left(1+\frac{\pi_0}{\alpha\gamma}\right)[\exp(\alpha\gamma)\text{-}1]\exp(\beta+\alpha\pi_0\text{-}\alpha\gamma/\delta)\exp[\alpha\gamma(t\text{-}1)]+$$

$$+\frac{1}{\alpha}\exp(\beta+\alpha\pi_0\text{-}\alpha\gamma/\delta)\exp[\alpha\gamma(t\text{-}1)]\left[t\exp(\alpha\gamma)\text{-}(t\text{-}1)\text{-}\frac{\exp(\alpha\gamma)\text{-}1}{\alpha\gamma}\right] \tag{C.12}$$

By the same token, it can be shown that for the log-log it is equal to

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \exp(\beta)\left\{(\tilde{\pi}_0+\gamma t)^{\alpha}\text{-}[\tilde{\pi}_0+\gamma(t\text{-}1)]^{\alpha}\right\}+$$

$$+\frac{\exp(\beta)}{\gamma(1+\alpha)}\left\{(\pi_0+\gamma t)(\tilde{\pi}_0+\gamma t)^{1+\alpha}\text{-}[\pi_0+\gamma(t\text{-}1)][\tilde{\pi}_0+\gamma(t\text{-}1)]^{1+\alpha}\right\}-$$

$$-\frac{\exp(\beta)}{\gamma(1+\alpha)(2+\alpha)}\left\{(\tilde{\pi}_0+\gamma t)^{2+\alpha}\text{-}[\tilde{\pi}_0+\gamma(t\text{-}1)]^{2+\alpha}\right\} \tag{C.13}$$

where $\tilde{\pi}_0 = \pi_0\text{-}\gamma/\delta$. It can be trivially checked that for $\delta \to \infty$ expressions (C.12) and (C.13) converge to the corresponding expressions for the case of perfect foresight, (C.7) and (C.8).

I now turn to the case of an exponential time trend for inflation.

C.3.2 An exponential time trend for inflation

As in appendix C.2.2 inflation evolves according to $\pi_t \equiv \pi_0\exp(\gamma t)$—where, as before, π_0 and $\gamma > 0$ are constants—p_t and P_t evolve according to $p_t = p_0 + (\pi_0/\gamma)\exp(\gamma t)$ and $P_t = P_0\exp[(\pi_0/\gamma)\exp(\gamma t)]$. I conjecture that π^e_t also evolves according to an exponential time trend, $\pi^e_t = A\exp(Bt)$, with coefficients A and B to be determined. It can be shown that A and B are equal to A=$\delta\pi_0/(\delta+\gamma)$ and B=γ, so that $\pi^e_t = [\delta\pi_0/(\delta+\gamma)]\exp(\gamma t)$. Working as before, it can be shown that for the semi-log

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \exp(\beta)\gamma\pi_0\left(\frac{\alpha\delta}{\delta+\gamma}+\frac{1}{\gamma}\right)\int_{t\text{-}1}^{t}\exp\left[\gamma\tau+\pi_0\frac{\alpha\delta}{\delta+\gamma}\exp(\gamma\tau)\right]d\tau \tag{C.14}$$

whereas for the log-log

$$\xi^{\text{TRUE}}_{t\text{-}1\to t} = \exp(\beta)\left(\frac{\delta\pi_0}{\delta+\gamma}\right)^{\alpha}[\exp(\alpha\gamma)-1]\exp[\alpha\gamma(t\text{-}1)]+$$

$$+\exp(\beta)\pi_0\left(\frac{\delta\pi_0}{\delta+\gamma}\right)^{\alpha}\frac{\exp[(1+\alpha)\gamma]\text{-}1}{(1+\alpha)\gamma}\exp[(1+\alpha)\gamma(t\text{-}1)] \tag{C.15}$$

It can be trivially checked that for $\delta \to \infty$ expressions (C.14) and (C.15) converge to the corresponding expressions for the case of perfect foresight, (C.9) and (C.10).

C.3.3 Evidence

Figure A.2 shows the results from an exercise analogous to that performed in appendix C.2.3. I only report results for the log-log which, as shown in Section 5 in the main text, provides a significantly more plausible description of the data than the semi-log. Results for the semi-log are qualitatively the same as those shown in Figure A.1 for the case of perfect foresight, and they are available upon request.

Figure A.2 reports the approximation errors for ξ_t^*, defined as in appendix C.2.3 as the percentage deviation of ξ_t^* from $\xi_{t\text{-}1\rightarrow t}^{\text{TRUE}}$. I do not report the corresponding evidence for the arithmetic average because, once again, $\tilde{\xi}_t$ provides an extremely poor approximation to the true amount of seignorage, but this evidence is available upon request. I consider five values for δ, 0.5, 1, 2, 5, and 10. The evidence in Figure A.2 is qualitatively in line with that shown in Figure A.1. Even focusing on the case of adaptive expectations, as opposed to that of perfect foresight, the geometric average ξ_t^* still provides an excellent approximation to the true amount of seignorage collected by the goverment.

D Two Alternative Approaches to Estimating the Elasticity of Money Demand

D.1 A semi-structural approach in the spirit of Salemi and Sargent (1979)

Salemi and Sargent (1979) postulated a VAR representation for the joint dynamics of inflation and money growth, and estimated it *via* maximum likelihood conditional on the rational expectations restrictions implied by Cagan's semi-log functional form. In this appendix I adopt an approach combining Salemi and Sargent's (1979) insight of postulating a time-series representation for the series of interest, and imposing upon it the restrictions implied by a theoretical specification for the demand for real money balances, with elements borrowed from Hamilton (1985), and from Burmeister and Wall (1982, 1987).

I define the logarithms of real money balances and the expected inflation as $\tilde{m}_t \equiv \ln(M_t/P_t)$ and $\tilde{\pi}_t \equiv \ln(\pi_{t+1|t})$ respectively, with $\pi_{t+1|t}$ being the rational expectation of inflation at time t+1, conditional on information at time t. Being conditional on information at time t, $\tilde{\pi}_t$ is, by its very nature, a dated-t object.

Based on this notation, the demand for real money balances is given by

$$\tilde{m}_t = \beta + \alpha\tilde{\pi}_t + u_t \tag{D.1}$$

where u_t is a money demand disturbance. Whereas Sargent (1977) and Salemi and Sargent (1979) postulate that u_t is a random walk,[49] in what follows I assume that

[49] Theoretical models of seignorage in which the governments finances, *via* the inflation tax, a

it evolves according to

$$u_t = \rho u_{t-1} + v_t, \qquad |\rho| \leq 1, \tag{D.2}$$

i.e. a specification nesting the random walk case, but also allowing for stationarity. The key reason[50] for doing so is the evidence in Figure 7, especially for episodes such as post-WWII Hungary, Zimbabwe, Yugoslavia, Germany, and Greece. Focusing, e.g., on Yugoslavia, a comparison between the two panels in the third column of Figure 7 naturally suggests (at least) two possible interpretations of the joint dynamics of log real money balances and expected inflation. One possibility is that the true money demand specification is Cagan's semi-log and that the disturbance is very highly persistent, possibly a random walk, which is suggested by the very persistent divergence between the two series in the top panel. An alternative interpretation—which appears (at least, to me) as distinctly more appealing—is that the true functional form is Meltzer's log-log and that the disturbance has very little persistence. This is suggested by the fact that the logarithms of inflation[51] and real money balances in the bottom panel track each other very closely.[52]. This logically implies that imposing ρ=1 automatically 'stacks the cards' in favor of the semi-log, and against the log-log, so that estimates obtained conditional on the assumption that u_t follows a random walk should not be regarded as reliable. In what follows I will therefore assume that u_t evolves according to (D.2), and I will estimate ρ *via* maximum likelihood together with the other parameters of the model.

Turning to the time-series characterization of the joint dynamics of the series of interest, conceptually in line with Hamilton (1985) and Burmeister and Wall (1982, 1987) I postulate that it is described by

$$\begin{bmatrix} \tilde{m}_t \\ \tilde{\pi}_t \end{bmatrix} = \begin{bmatrix} \bar{m} \\ \bar{\pi} \end{bmatrix} + \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \tilde{m}_{t-1} \\ \tilde{\pi}_{t-1} \end{bmatrix} + \begin{bmatrix} H_{\tilde{m}} \\ H_{\tilde{\pi}} \end{bmatrix} u_t + \begin{bmatrix} \epsilon_t^{\tilde{m}} \\ \epsilon_t^{\tilde{\pi}} \end{bmatrix} \tag{D.3}$$

where $a(L)$, ..., $d(L)$ are polynomials in the lag operator, $\bar{m}$ and $\bar{\pi}$, and $H_{\tilde{m}}$ and $H_{\tilde{\pi}}$, are constants, and $\epsilon_t^{\tilde{m}}$ and $\epsilon_t^{\tilde{\pi}}$ are shocks. Before turning to the restrictions imposed by (D.1) on (D.3), it is worth spending a few words discussing why, within the present context, it is necessary to work with a model such as (D.3) postulating a linear VARX representation for the joint dynamics of log real money balances and the *logarithm* of expected inflation, rather than either inflation or expected inflation. The key reason

constant fraction of GDP, produce a random-walk disturbance in the money demand equation.

[50] Christiano (1987) produces some evidence against the random walk assumption based on Cagan's data for Germany's episode.

[51] Although the theoretical relationship is between log real money balances and *expected*—rather than *actual*—inflation, by the rational expectation hypothesis the difference between them should be white noise. This logically implies that the fact that the *logarithm* of *actual* inflation tracks log real money balances much more closely than its *level* should be taken as an indication that the same holds for *expected* inflation.

[52] Obviously, once appropriately rescaled, which is implicitly implemented by allowing for different scales in the left- and right-hand side axes.

for this is that if we specified a linear VARX (or VAR) representation for the joint dynamics of $\tilde{m}_t$, and either $\pi_{t+1|t}$ or π_t, this would be inconsistent with the *non-linear* relationship between $\tilde{m}_t$ and $\pi_{t+1|t}$ in equation (D.1).[53] To the extent that we want to work with a *linear* time-series representation for the joint dynamics of the series of interest, the linear relationship between $\tilde{m}_t$ and $\tilde{\pi}_t$ in equation (D.1) logically implies that such time-series representation ought to be for those two series.

Equation (D.1) imposes the following restrictions upon (D.3):

$$\bar{m} = \beta + \alpha\bar{\pi} \tag{D.4}$$

$$a(L) = \alpha c(L) \tag{D.5}$$

$$b(L) = \alpha d(L) \tag{D.6}$$

$$\epsilon_t^{\tilde{m}} = \alpha\epsilon_t^{\tilde{\pi}} \tag{D.7}$$

and

$$H_{\tilde{m}} = 1 + \alpha H_{\tilde{\pi}} \tag{D.8}$$

In what follows I normalize $H_{\tilde{m}}$ to be equal to 1, which implies that $H_{\tilde{\pi}}$=0.

Finally, imposing equality between specification (D.1) for the demand for real money balances, and the equation for $\tilde{m}_t$ in the VAR representation (D.3), produces the following restriction for the error term for $\tilde{m}_t$ in (D.3):

$$\epsilon_t^{\tilde{m}} = \alpha\epsilon_t^{\tilde{\pi}} \tag{D.9}$$

By defining the state vector as $\xi_t = [u_t, \tilde{\pi}_t, \tilde{\pi}_{t-1}, ..., \tilde{\pi}_{t-p+1}]'$, where p is the lag order in the lag polynomials $a(L)$, ..., $d(L)$,[54] the model can be cast in state-space form, with state equation

$$\underbrace{\begin{bmatrix} u_t \\ \tilde{\pi}_t \\ \tilde{\pi}_{t-1} \\ ... \\ \tilde{\pi}_{t-p+2} \\ \tilde{\pi}_{t-p+1} \end{bmatrix}}_{\xi_t} = \underbrace{\begin{bmatrix} 0 \\ \bar{\pi} \\ 0 \\ ... \\ 0 \\ 0 \end{bmatrix}}_{f} + \underbrace{\begin{bmatrix} \rho & 0 & 0 & ... & 0 \\ 0 & d_1 & d_2 & ... & d_p \\ 0 & 1 & 0 & ... & 0 \\ ... & ... & ... & ... & ... \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & ... & 0 \end{bmatrix}}_{F} \underbrace{\begin{bmatrix} u_{t-1} \\ \tilde{\pi}_{t-1} \\ \tilde{\pi}_{t-2} \\ ... \\ \tilde{\pi}_{t-p+1} \\ \tilde{\pi}_{t-p+2} \end{bmatrix}}_{\xi_{t-1}} + \tag{D.10}$$

[53] The same logic would hold if, e.g., following Salemi and Sargent (1979), we postulated a linear representation for the joint dynamics of inflation and money growth.

[54] As discussed by Salemi and Sargent (1979, p. 746), because of the short sample length which is typical of nearly all hyperinflationary episodes, the lag order ought necessarily to be set to a comparatively small value. I set it to p to either 1 or 2 with monthly data (in what follows I only present results based on p=2, but the alternative set of results based on p=1, which is qualitatively the same, is available upon request).

$$+\underbrace{\begin{bmatrix} 0 & 0 & ... & 0 \\ c_1 & c_2 & ... & c_p \\ 0 & 0 & ... & 0 \\ ... & ... & ... & ... \\ 0 & 0 & 0 & 0 \\ 0 & 0 & ... & 0 \end{bmatrix}}_{G}\underbrace{\begin{bmatrix} \tilde{m}_{t-1} \\ \tilde{m}_{t-2} \\ ... \\ \tilde{m}_{t-p} \end{bmatrix}}_{M_t}+\underbrace{\begin{bmatrix} v_t \\ \epsilon_t^{\tilde{\pi}} \\ 0 \\ ... \\ 0 \\ 0 \end{bmatrix}}_{w_t}$$

with

$$Q \equiv E(w_t w_t') = \begin{bmatrix} \sigma_v^2 & \sigma_{v,\tilde{\pi}} & 0 & ... & 0 \\ \sigma_{v,\tilde{\pi}} & \sigma_{\tilde{\pi}}^2 & 0 & ... & 0 \\ 0 & 0 & 0 & ... & 0 \\ ... & ... & ... & ... & ... \\ 0 & 0 & 0 & ... & 0 \end{bmatrix} \tag{D.11}$$

where $E(\cdot)$ is the unconditional expectation operator.

As for the observation equations, the first is given by

$$\pi_t = \pi_{t|t-1} + \eta_t = [\theta\tilde{\pi}_{t-1} + 1]^{\frac{1}{\theta}} + \eta_t \tag{D.12}$$

where η_t is a rational expectations forecast error, i.e. η_t=π_t-$\pi_{t|t-1}$, with $\eta_{t|t-1}$=0, and the second equality comes from the very definition of $\tilde{\pi}_t$ as the Box-Cox transformation of $\pi_{t+1|t}$. I allow for η_t to be correlated with the disturbances in the state equation, that is,

$$Z \equiv E(\eta_t w_t') = \begin{bmatrix} \sigma_{\eta v} & \sigma_{\eta\tilde{\pi}} & 0 & ... & 0 \end{bmatrix} \tag{D.13}$$

where $\sigma_{\eta v}$ and $\sigma_{\eta\tilde{\pi}}$ are covariances which I estimate together with the other parameters of the model. The key reason for this is that, from a general equilibrium perspective, rational expectations forecast errors do originate from the shocks hitting the system,[55] so that they ought to be allowed to be correlated to them.

As for the observation equation for $\tilde{m}_t$, there are two equivalent ways to proceed. The first is to use the equation for $\tilde{m}_t$ in the VAR representation (D.3), whereas the second is to simply use equation (D.1), which can be rewritten as

$$\tilde{m}_t = \beta + [1\ \alpha\ 0\ ...\ 0]\xi_t. \tag{D.14}$$

Because of restriction (D.9), the two representations for $\tilde{m}_t$ are equivalent. In what follows I will use equation (D.14).

The state-space model described by equations (D.10), (D.12), and (D.14) is linear with the single exception of the observation equation for inflation, expression (D.12). Following (e.g.) Harvey (1989), I therefore take a first-order Taylor expansion of

[55] E.g., any solution method for linear rational expectations models, such as Sims' (2000) produces a linear mapping between the model's structural shocks and the rational expectations forecast errors.

(D.12) around $\xi_{t|t-1}$, thus obtaining the following approximate expression for the observation equation for inflation

$$\pi_t \simeq \underbrace{[\theta\tilde{\pi}_{t-1|t-1}+1]^{\frac{1}{\theta}} - \tilde{\pi}_{t-1|t-1}[\theta\tilde{\pi}_{t-1|t-1}+1]^{\frac{1-\theta}{\theta}}}_{c_\pi} + [\theta\tilde{\pi}_{t-1|t-1}+1]^{\frac{1-\theta}{\theta}}\tilde{\pi}_{t-1} + \eta_t =$$

$$= c_\pi + [\theta\tilde{\pi}_{t-1|t-1}+1]^{\frac{1-\theta}{\theta}}\tilde{\pi}_{t-1} + \eta_t \tag{D.15}$$

Replacing the original non-linear observation equation (3.12) with (3.15) results in a fully linear system, which allows to compute the log-likelihood *via* the standard 'prediction error decomposition' formula—see e.g. Harvey (1989), or Hamilton (1994).

D.1.1 Maximum likelihood estimation

I maximize the log-likelihood *via* simulated annealing.[56] Having found the parameter vector which maximizes the likelihood, $\hat{B}_{MLE}$, I stochastically map the log-likelihood's surface *via* Random-Walk Metropolis (RWM). The only difference between the 'standard' RWM algorithm that is routinely used for Bayesian estimation and what I am doing here is that the 'jump' to the new position in the Markov chain is accepted or rejected based on a rule which does not involve any Bayesian priors, as it uniquely involves the likelihood of the data.[57] Specifically, the proposal draw for B, $\tilde{B}$, is accepted with probability min[1, $r(B_{s-1}, \tilde{B} \mid Y)$]—where Y is the matrix of the data for $\tilde{m}_t$ and π_t—and rejected otherwise, where B_{s-1} is the current position in the Markov chain, and

$$r(B_{s-1}, \tilde{B}|Y) = \frac{L(\tilde{B} \mid Y)}{L(B_{s-1} \mid Y)} \tag{D.16}$$

which uniquely involves the likelihood.[58] All other estimation details—e.g., computing the Hessian at the mode of the log-likelihood Berndt, Hall, Hall, and Hausman's methodology—are identical to Benati (2008), to which the reader is referred to. I use

[56] Specifically, following Goffe *et al.* (1994), I implement simulated annealing *via* the algorithm proposed by Corana *et al.* (1987), setting the key parameters to T_0 =100,000, $r_T = 0.9$, $N_t = 5$, $N_s = 20$, $\epsilon = 10^{-6}$, and $N_\epsilon = 4$, where T_0 is the initial temperature, r_T is the temperature reduction factor, N_t is the number of times the algorithm goes through the N_s loops before the temperature starts being reduced, N_s is the number of times the algorithm goes through the function before adjusting the step size, ϵ is the convergence (tolerance) criterion, and N_ϵ is the number of times convergence is achieved before the algorithm stops. Finally, initial conditions were chosen stochastically by the algorithm itself, whereas the maximum number of functions evaluations, set to 1,000,000, was never achieved.

[57] So what I am doing can be interpreted as Bayesian estimation with flat priors for all parameters.

[58] With Bayesian priors it would be

$$r(B_{s-1}, \tilde{B}|Y) = \frac{L(\tilde{B} \mid Y)P(\tilde{B})}{L(B_{s-1} \mid Y)P(B_{s-1})}$$

where $P(\cdot)$ encodes the priors about B.

1,000,000 draws for the burn-in period, and 1,000,000 draws for the ergodic distribution, which I 'thin' by sampling every 1,000 draws in order to reduce the draws' autocorrelation. The fractions of accepted draws are uniformly very close to the 23 per cent ideal acceptance rate in high dimensions[59], and the draws exhibit little autocorrelation based on the draws' inefficiency factors.[60]

D.1.2 Evidence

Once again I consider three panels, featuring both all of 20 the episodes considered jointly, and the 10 episodes with either the highest or the lowest median inflation rates, respectively. Consistent with the use of country-specific dummies in the main text, I allow both $\bar{m}$ and $\bar{\pi}$ (and therefore β) to be country-specific, whereas I impose that all of the remaining parameters be the same for all of the countries in the panel. The first panel of Figure A.10 shows the distributions of the draws from the ergodic distributions generated *via* Random-Walk Metropolis for the elasticity of real money balances. Evidence is qualitatively in line with, and quantitatively very close to, that based on panel VARs discussed in Section 8 in the main text. In particular, as discussed in Section 8.1, (i) only for the panel featuring the 10 most extreme episodes there is some weak evidence that the elasticity may have been greater than one in absolute value, whereas for either the 10 least extreme episodes, or all of the 20 episodes considered jointly, evidence strongly rejects such notion; and (ii) evidence strongly suggests that higher inflation has consistently been associated with a larger (in absolute value) elasticity. In particular, whereas based on all of the 20 episodes jointly the median and the 5th and 95th percentiles of the distribution of the elasticity are equal to -0.573 [-0.669; -0.486], the corresponding figures for the 10 most and least extreme episodes are -0.932 [-1.094; -0.759] and -0.306 [-0.135; -0.468], respectively.

D.2 A model with recursive least-squares learning about inflation

The second panel of Figure A.10 shows the distributions of the draws from the ergodic distributions generated *via* Random-Walk Metropolis for the elasticity of real money balances, based on a model in which, at each point in time, expected inflation next period is produced based on either an AR(p) model for inflation, or a VAR(p) model for inflation and money growth, which are estimated *via* recursive least squares.[61]

[59] See Gelman, Carlin, Stern, and Rubin (1995).

[60] The inefficiency factors are defined as the inverse of the relative numerical efficiency measure of Geweke (1992), $RNE = (2\pi)^{-1} \frac{1}{S(0)} \int_{-\pi}^{\pi} S(\omega) d\omega$, where $S(\omega)$ is the spectral density of the sequence of draws from RWM for the quantity of interest at the frequency ω. I estimate the spectral densities as before, based on the FFT transform.

[61] The figure shows results based on the univariate scheme for inflation forecasting, but those based on the VAR for money growth and inflation are numerically very close (these results are available upon request). This is in line with Sargent and Wallace's (1973) classic result that, within a bivariate

Before discussing the details of the model and of the estimation, which are entirely standard, it is worth stressing, how the results are qualitatively in line with those in Section 8 in the main text, and in the previous subsection based on the semi-structural model in the spirit of Salemi and Sargent (1979). The only difference is that, for any of the three sets of episodes, the estimated elasticities are here uniformly greater in absolute value. In particular, for the panel of 10 most extreme episodes, there is strong evidence that the elasticity may have been slightly below -1, with the median and the 5th and 95th percentiles of the distribution being equal to -1.034 [-1.055; -1.016], and the fractions of draws below -1 being equal to 0.999.

Entering into details, the demand for real money balances is given by

$$\tilde{m}_t = \beta + \alpha \ln \pi^e_{t+1} + u_t \tag{D.17}$$

where u_t is, again, a money demand disturbance, and π^e_{t+1} is the expectation of inflation at t+1. At each point in time t, agents estimate, based on data up to time t-1,[62] either an AR(p) model for inflation, or a VAR(p) model for inflation and money growth (computed as the log-difference of the nominal money stock), based on recursive least squares (RLS) with a constant gain γ. Based on data up to time t-1, and on the estimated model, agents then produce the inflation forecast at t+1, π^e_{t+1}, by projecting the model two periods into the future. Given the sequence of inflation forecasts π^e_{t+1}, model (D.17) can then be estimated *via* MLE as in the previous sub-section. Specifically, details about (i) computing the mode of the log-likelihood *via* simulated annealing, and (ii) performing RWM in order to stochastically map the log-likelihood's surface, as well as all other estimation details, are the same as in the previous sub-section. The only issue deserving a detailed discussion is the inizialization of the RLS algorithm for estimating the two models used to produce the inflation forecasts. Based on the AR(p) model for inflation

$$\pi_t = \phi_0 + \phi_1 \pi_{t-1} + ... + \phi_p \pi_{t-p} + v_t \tag{D.18}$$

where the notation is obvious, I initialize $\hat{\phi}_1^{RLS}$=1, and $\hat{\phi}_0^{RLS}$=$\hat{\phi}_2^{RLS}$=...$\hat{\phi}_p^{RLS}$=0, corresponding to the 'Minnesota prior' assumption that inflation is a random-walk. I regard this assumption as a plausible one for hyperinflations. For a given value of the constant gain γ, which I estimate *via* MLE together with the other parameters of the model, the RLS estimates $\hat{\phi}^{RLS}$=[$\hat{\phi}_0^{RLS}$, $\hat{\phi}_2^{RLS}$, ..., $\hat{\phi}_p^{RLS}$]′ get updated according to

$$\hat{\phi}_t^{RLS} = \hat{\phi}_{t-1}^{RLS} + \gamma \hat{R}_{t-1}^{-1} z_{t-1} [\pi_{t-1} - \hat{\phi}_{t-1}^{RLS} z_{t-1}] \tag{D.19}$$

$$\hat{R}_t = \hat{R}_{t-1} + \gamma [z_{t-1} z'_{t-1} - \hat{R}_{t-1}] \tag{D.20}$$

VAR representation for money growth and inflation, the former does not Granger-cause the latter.

[62] The assumption that agents use data up to t-1 in order to estimate the relevant model is standard in the literature on recursive least squares.

where $z_t = [1\ \pi_{t-1}\ ...\ \pi_{t-p}]'$, and with $\hat{R}_t$ being initialized as I_{p+1}, i.e. a (p+1) identity matrix.

By the same token, based on the VAR(p) model for inflation and money growth

$$\begin{bmatrix} \pi_t \\ \theta_t \end{bmatrix} = \begin{bmatrix} \pi_0 \\ \theta_0 \end{bmatrix} + \begin{bmatrix} A(L) & B(L) \\ C(L) & D(L) \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ \theta_{t-1} \end{bmatrix} + \begin{bmatrix} v_t^{\pi} \\ v_t^{\theta} \end{bmatrix} \tag{D.21}$$

—where, again, $\theta_t = \ln(M_t)$-$\ln(M_{t-1})$; $A(L) = A_1 L + A_2 L^2 + ... + A_2 L^2$; and $B(L)$, $C(L)$, and $D(L)$ defined in the same way—I initialize A_1=D_1=1; and I initialize π_0, θ_0, and all of the other parameters in the lag polynomials $A(L)$, $B(L)$, $C(L)$, and $D(L)$ (except A_1 and D_1) to zero, corresponding once again to 'Minnesota prior' assumptions for either inflation or money growth. In each period t, the VAR (D.21) is estimated equation by equation *via* RLS based on information at time t-1. Specifically, for a given value of γ, the RLS estimates of the VAR's parameters are updated according to equations equivalent to (D.19) and (D.20), where $\hat{R}_t$ is still initialized as an identity matrix. Finally, based on either (D.18) or (D.21) I set the lag order to 2.