

Benati, Luca

Working Paper

Cagan's paradox revisited

Discussion Papers, No. 18-26

Provided in Cooperation with:

Department of Economics, University of Bern

Suggested Citation: Benati, Luca (2018) : Cagan's paradox revisited, Discussion Papers, No. 18-26, University of Bern, Department of Economics, Bern

This Version is available at:

<https://hdl.handle.net/10419/204926>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics
and Social Sciences

Department of Economics

Cagan's Paradox Revisited

Luca Benati

18-26

December 2018

DISCUSSION PAPERS

Schanzeneckstrasse 1
Postfach 8573
CH-3001 Bern, Switzerland
<http://www.vwi.unibe.ch>

Cagan’s Paradox Revisited*

Luca Benati
University of Bern[†]

Abstract

Data from 20 hyperinflations—from the French Revolution to Venezuela’s 2018 episode—provide nearly *no evidence* of a Laffer curve for seignorage. Rather, in nearly all cases, the relationship between the inflation tax and inflation has been either positive at all inflation rates, or initially positive and then flattening out towards the end of the hyperinflation. Consistent with this, econometric evidence shows that the preferred money demand specification at very high inflation rates is not Cagan’s (1956) ‘semi-log’, which automatically imposes a Laffer curve upon the data: rather, it is either Meltzer’s (1963) ‘log-log’—for which the inflation tax is monotonically increasing in inflation—or a more general functional form making log real money balances a linear function of the Box-Cox transformation of expected inflation (of which the ‘log-log’ is a special case), which allows the inflation tax to flatten out at high inflation rates. My results suggest that the paradox first highlighted by Cagan—of policymakers seemingly inflating in excess of the revenue-maximizing rate during hyperinflations—is the product of the literature’s predominant focus on the semi-log specification.

*I wish to thank Harris Dellas, Giovanni Lombardo, and Francisco Ruge-Murcia for useful discussions. Special thanks to Eugene White, Liuyan Zhao, Carlos Gustavo Machicado, and Zorica Mladenovic, for kindly providing data for the French Revolution’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

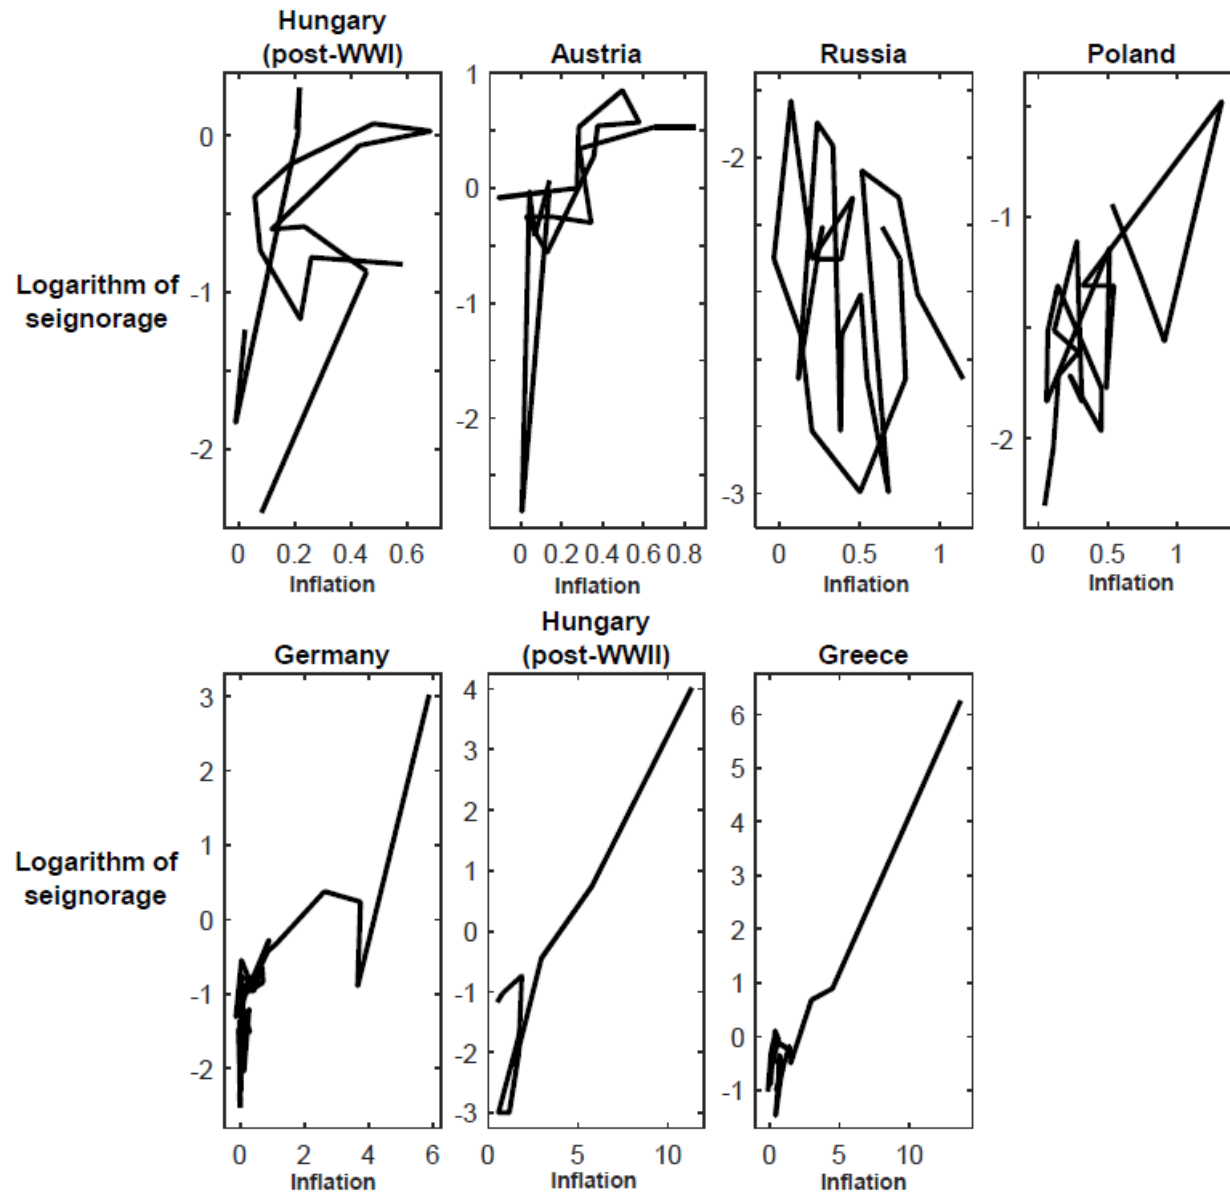


Figure 1 The logarithm of Sargent and Wallace's (1973) measure of seignorage plotted against the inflation rate from Cagan (1956)

1 Introduction

In a classic paper, Phillip Cagan (1956) first documented the puzzling paradox of policymakers seemingly inflating in excess of the revenue-maximizing rate during hyperinflations. Since then, Cagan’s paradox has been broadly confirmed by several subsequent studies, to the point that it has nearly achieved the status of a ‘stylized fact’ in empirical macroeconomics.

In this paper I revisit Cagan’s paradox based on data from 20 hyperinflations, from the French Revolution to Venezuela’s 2018 episode. I report two main findings:

(I) evidence of a Laffer curve for the inflation tax is nearly *non-existent* in the raw data. Rather, in nearly all cases, the relationship between the inflation tax and inflation has been either positive at all inflation rates, or initially positive and then flattening out towards the end of the hyperinflation. As a simple illustration, Figure 1 shows the logarithm¹ of Sargent and Wallace’s (1973) measure of the revenue from money creation² for the seven episodes in Cagan’s (1956) dataset, plotted against the inflation rates from Cagan (1956).³ For the two most extreme episodes—Greece and post-WWII Hungary—the relationship between the two series clearly appears as monotonically increasing. The same appears to hold, at much lower inflation rates, for Poland, whereas for Austria the relationship is initially steeply increasing, and it then appears to flatten out at higher inflation rates, towards the end of the hyperinflation. Evidence for post-WWI Hungary is not clear-cut but, at the very least, it provides little to no support to the notion of a Laffer curve for seignorage. Evidence for Germany is, at first sight, qualitatively in line with that for Greece and post-WWII Hungary, but this crucially hinges on the very last observation: dropping it, a Laffer curve would indeed appear from the data. Finally, evidence for Russia is inconclusive.⁴ As I will discuss in Section 5.1.1, the visual evidence for the remaining episodes in my dataset (which is reported in Figure 3) is qualitatively in line with

¹I show the *logarithm* of Sargent and Wallace’s measure because the dramatic extent of variation exhibited by the level towards the end of the hyperinflations would make it impossible to visually identify its relationship with inflation.

²Sargent and Wallace’s (1973) measure is from their Table 6, page 345, and it had been computed as $(M_t - M_{t-1}) / [1/2(P_t + P_{t-1})]$, where M_t and P_t are the nominal money stock and the price level from Cagan (1956).

³The inflation rates are from Cagan’s (1956) Appendix B (‘Data and Sources’). Cagan measured inflation as $\log_{10}(P_t/P_{t-1})$, and I converted it to natural logarithms.

⁴The evidence for Russia’s episode 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 in output should reasonably thought of as only introducing a minor distortion in the estimates. Barro (1970) also had reservations about Cagan’s money supply data for Greece (and in fact he eschewed them), but, as I discuss in Section 4.1, these data produce, in fact, qualitatively the same results as the money supply data from Agapitides (1945) and Delivanis and Cleveland (1949). So, whatever the shortcomings in Cagan’s Greek money supply data might be, they do not seem to have any material impact on my estimates.

that shown in Figure 1.

(II) Consistent with (I), econometric evidence shows that the preferred money demand specification at very high inflation rates is not Cagan’s (1956) ‘semi-log’, which automatically imposes upon the data a Laffer curve for seignorage: rather, it is either Allan Meltzer’s (1963) ‘log-log’⁵—for which the inflation tax is monotonically increasing in inflation—or a more general functional form making log real money balances a linear function of the Box-Cox transformation of expected inflation,⁶ which allows for the inflation tax to flatten out at high inflation rates. In several cases the data’s preference for the log-log specification is clearly apparent even to the naked eye. A stark illustration is provided by Yugoslavia, for which the *logarithm* of inflation tracks log real money balances remarkably closely,⁷ whereas its *level* exhibits a scant connection with it. The same holds, e.g., for Greece, Germany, and post-WWII Hungary, and to a slightly lesser extent for China and Zimbabwe.

My results therefore suggest that the findings of both (i) a Laffer curve for the inflation tax, and (ii) policymakers having near-uniformly inflated in excess of the revenue-maximizing rate during hyperinflations, are the product of the literature’s predominant focus on the semi-log functional form, which automatically imposes upon the data a Laffer curve.

My overall conclusion is therefore that Cagan’s paradox is simply an illusion, originating from the literature’s predominant focus on the semi-log specification for the demand for real money balances, which automatically imposes upon the data a Laffer curve for the inflation tax. On the other hand, the fact that, at very high inflation rates, the preferred specification is Meltzer’s (1963) log-log (or a specification very close to it), logically rules out the possibility of such paradox.

The paper is structured as follows. In the next section I briefly discuss how alternative functional forms for the demand for real money balances map into different relationships between inflation and seignorage. Section 3 presents a brief overview of the literature, whereas Section 4 discusses in detail the data and their sources. Section 5 presents and discusses the evidence, starting from the simplest—but, I will argue, the most powerful one—i.e. the raw data, and then moving to econometric evidence. Section 6 concludes, and outlines three possible directions for future research.

2 Theory

Figure 2 shows the logarithm⁸ of the revenue from money creation as a function of inflation for alternative functional forms for the demand for real money balances. The

⁵To be precise, Meltzer’s (1963) study did not pertain to hyperinflations.

⁶As I discuss in Section 2, within this framework the semi-log and log-log specifications represent ‘corner solutions’, corresponding to the Box-Cox parameter taking the values of 1 and 0, respectively.

⁷Quite obviously, once appropriately rescaled.

⁸I show the logarithm of seignorage (rather than seignorage itself) in order to allow the reader a direct visual comparison with the evidence shown in Figures 1 and 3.




Figure 2 The logarithm of the revenue from money creation as a function of inflation, for alternative values of the Box-Cox parameter in specification (2) for the demand for real money balances

revenue from money creation is defined as

$$S_t \equiv \frac{dM_t}{dt} \frac{1}{P_t} = \underbrace{\left(\frac{dM_t}{dt} \frac{1}{P_t} \right)}_{\mu_t} \frac{M_t}{P_t} = \mu_t \frac{M_t}{P_t} \quad (1)$$

where M_t and P_t are nominal money balances and the price level, respectively, and μ_t is money growth. The alternative specifications for the demand for real money balances considered in Figure 2 are all particular cases of the functional form

$$\ln \left(\frac{M_t}{P_t} \right) = \beta + \alpha \frac{[\pi_t^e]^\theta - 1}{\theta}, \quad (2)$$

positing a linear relationship between log real money balances and the Box-Cox transformation of expected inflation, π_t^e , with $\alpha < 0$ and the Box-Cox parameter $\theta \in [0, 1]$. A key feature of (2) is that while it nests both Cagan's (1956) semi-log specification

$$\ln \left(\frac{M_t}{P_t} \right) = \beta + \alpha \pi_t^e \quad (3)$$

and Meltzer's (1963) log-log one

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

as special cases—specifically, ‘corner solutions’ corresponding to θ taking the value of either 1 or 0, respectively—it also allows for a continuum of intermediate cases. In particular, starting from $\theta=1$ —for which the Laffer curve for the inflation tax associated with Cagan's semi-log is clearly apparent—smaller and smaller values of θ cause the right hand-side portion of the Laffer curve to become progressively flatter. For values of θ significantly smaller than 0.5—corresponding to which the Laffer curve becomes, asymptotically, exactly flat—the relationship becomes monotonically increasing, exactly as we saw in Figure 1 for Greece and post-WWII Hungary.

In Section 5 I will estimate θ *via* maximum likelihood for each individual episode. To anticipate, in the vast majority of cases the distribution of the MLE estimates generated *via* Random-Walk Metropolis (RWM) is clustered towards 0. Only in *one* case, Russia, $\theta=1$ is a plausible estimate, whereas in nearly all other cases it is very highly implausible based on the distribution produced by RWM. Further, with the single exception of Russia, the implied distributions of the relationship between inflation and the inflation tax provide virtually no evidence of a hump-shaped, Laffer-type relationship.

I now turn to a brief overview of the literature.

3 Related Literature

Cagan’s (1956) paper spawned an enormous 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 functional form for the demand for real money balances

Cagan (1956) did not derive the semi-log specification (3) within a micro-founded framework, but rather postulated it.⁹ In reaction to the empirical shortcoming of the postulated specification for the latest stages of hyperinflations—for which the estimated models’ fit had typically been worse than for the initial stages¹⁰—he speculated, however, that an alternative functional form may be needed in order to meaningfully characterize the dynamics of the data. In particular, he entertained the possibility¹¹

‘[...] that the function that determines the demand for real cash balances does not conform to [the semi-log specification]. 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 functional form Cagan was speculating about should have been either a log-log, or a specification close to it. In the end, Cagan’s solution¹² was neither to use a log-log, nor a specification such as (2), but rather to simply exclude the latest stages of the hyperinflations 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

⁹See Cagan (1956, p. 35, equation 2). In fact, both the semi-log and the log-log specifications can be derived within either a generalized Baumol (1952)-Tobin (1956) framework (see Benati, Lucas, Nicolini, and Weber, 2018), or Sidrauski’s (1967, 1968) ‘money-in-the-utility-function’ framework (see Belongia and Ireland, 2018).

¹⁰This problem has been repeatedly documented by several authors. E.g., for Yugoslavia’s episode, see Petrovic and Mladenovic (2000).

¹¹To 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 to so clearly conform to those implied by a log-log specification.

¹²See Cagan (1956, Section IV.3, ‘Observations that do not fit the regression’, p. 55).

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.’

Even this solution, however, was less than satisfactory, as the residuals in Cagan’s estimated equations were, in fact, strongly serially correlated.

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 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, once 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. As I discuss in Section 5.2.1, 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 unreliable¹³ (in the next sub-section I discuss a specific example, Zhao’s 2018 study of the Chinese hyperinflation).

3.3 Papers 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 which produced a markedly better fit than Cagan’s (1956) semi-log specification.¹⁴ 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 and the real interest rate (plus additional terms). The superior fit of Barro’s specification, compared to Cagan’s (1956), is therefore compatible with the present paper’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 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, as mentioned in the previous sub-section, 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 discuss in detail the data.

4 The Data

Here follows a detailed description of the data and of their sources for the 20 episodes of hyperinflation I consider herein, in chronological order.

¹³E.g. Johansen (2002), with reference to his trace and maximum eigenvalue tests, showed that asymptotic critical values are essentially unreliable in small samples.

¹⁴As 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’.

Table 1 Maximum, mean, and median inflation during hyper-inflationary episodes^a			
	Maximum	Mean	Median
France (January 1794-June 1796)	0.761	0.183	0.154
<i>Based on Cagan's (1956) data</i>			
Austria (January 1921-August 1922)	0.852	0.258	0.212
Germany (September 1920-November 1923)	5.885	0.616	0.157
Hungary (July 1922-February 1924)	0.683	0.248	0.215
Hungary (July 1945-July 1946)	33.670	4.909	1.677
Poland (April 1922-January 1924)	1.322	0.368	0.302
Russia (January 1922-February 1924)	1.142	0.451	0.420
Greece (January 1943-November 1944)	13.659	1.378	0.470
<i>Based on Barro's (1970) data</i>			
Austria (January 1921-December 1922)	0.852	0.220	0.149
Poland (January 1922-January 1924)	1.229	0.349	0.289
Germany (January 1921-August 1923)	2.931	0.339	0.136
Hungary (October 1921-February 1924)	0.683	0.191	0.152
<i>Other datasets</i>			
Germany ^b (June 1918-December 1923)	5.736	0.410	0.071
Greece ^c (January 1941-October 1944)	4.325	0.470	0.249
Greece ^d (January 1941-October 1944)	4.499	0.473	0.284
China (August 1945-May 1949)	4.446	0.552	0.240
Chile (January 1972-November 1974)	0.575	0.120	0.106
Bolivia (February 1983-August 1986)	1.039	0.220	0.155
Argentina (January 1987-April 1991)	0.992	0.185	0.135
Brazil (August 1988-March 1991)	0.592	0.235	0.199
Peru (January 1987-September 1991)	1.512	0.211	0.147
Yugoslavia (January 1991-January 1994)	11.290	1.386	0.496
Congo (May 1991-July 1995)	2.144	0.296	0.241
Angola (December 1995-January 1998)	0.610	0.167	0.061
Bulgaria (January 1996-February 1998)	1.230	0.128	0.039
Zimbabwe (April 2006-June 2008)	3.912	0.706	0.406
Venezuela (January 2017-November 2018)	0.999	0.411	0.355
^a Inflation is computed as the monthly log-difference of the price level.			
^b Based on Graham's (1930) data. ^c Based on Agapitides' (1945) data.			
^d Based on Delivanis and Cleveland's (1949) data.			

4.1 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’,¹⁵ available for the period December 1789-August 1796, 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. On the other hand, 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).

Cagan’s (1956) data are from Tables B1-B14 in his Appendix B.¹⁶ The countries and sample periods are the following: Austria (Jan. 1921-Aug. 1922), Germany (Sep. 1920-Nov. 1923), Hungary, post-WWI (Jul. 1922-Feb. 1924), Hungary post-WWII (Jul. 1945-Jul. 1946), Poland (Apr. 1922-Jan. 1924), Russia (Jan. 1922-Feb. 1924), and Greece (Jan. 1943-Nov. 1944). Barro (1970) eschewed Cagan’s data for Hungary’s post-WWII episode because of the short sample length—just 13 observations. In what follows I will also eschew these data, and I will perform the empirical analysis based on Anderson *et al.*’s (1988) weekly data (described below in Section 4.2), featuring 28 observations. As for Greece, Barro (1970, footnote 36) pointed out that ‘[t]he money-supply data for Greece was unreliable (Cagan, p. 106), and the variation in real income during the war was apparently substantial (*International Labor Review*, December 1945, p. 650).’ In what follows I will instead consider these data because (i) as previously discussed in footnote 4 with reference to Russia’s episode, since we are here dealing with *hyperinflation*s, even an unaccounted-for deep recession in output should reasonably be thought of as only introducing a minor distortion in the estimates; and (ii) as for the ‘unreliability’ of money-supply data, the alternative measures from Agapitides (1945) and Delivanis and Cleveland (1949) (described below) produce, in fact, qualitatively similar results. So, whatever the shortcomings in Cagan’s Greek money supply data might be, they do not seem to have a material impact on the empirical evidence.

Barro’s (1970) data are from Tables A1-A4 in the Appendix.¹⁷ The countries and

¹⁵To 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).

¹⁶Cagan’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 log of real money balances, I recovered an index for log nominal money balances.

¹⁷Barro’s data feature the logarithm of real money balances and inflation, computed as the log-

sample periods are the following: Austria (Jan. 1921-Dec. 1922), Germany (Jan. 1921-Aug. 1923), Hungary, post-WWI (Oct. 1921-Feb. 1924), and Poland (Jan. 1922-Jan. 1924).

Graham’s data for Germany, available for the period Jun. 1918-Dec. 1923, are from Table XII of Graham (1930), and they feature two indices of wholesale prices, and of total monetary circulation, respectively.

Additional data for the Greek hyperinflation, for the period Jan. 1941-Oct. 1944, are from the Appendix of Anderson *et al.* (1988), and they feature money supply measures and price indices from either Agapitides (1945) or Delivanis and Cleveland (1949).

The data for China, which were used in Zhao and Li (2015), have been kindly provided by Liuyan Zhao. As detailed in Zhao and Li (2015), the data for currency are from Wu (1958, p. 92 and p. 122), whereas the price index is from Wu (1958, pp. 160-163). The sample period is August 1945-May 1949.

For Chile, data for M1 and the CPI, available since December 1965 and January 1947, respectively, are both from the *Banco Central do Chile*.

For Bolivia, data for M1, available since January 1952, are from Bolivia’s central bank, and they have been kindly provided by Carlos Gustavo Machicado, whereas data for the consumer price index, available since January 1967, are from the *Instituto Nacional de Estadística*.

For Argentina, data for both M1 and the CPI, available since January 1941 and January 1943, respectively, are from the *Banco Central de la República Argentina*.

For Brazil, a series for the CPI, available since January 1980, is from the *Instituto Brasileiro de Geografia e Estatística* (IBGE), whereas a series for M1, available since July 1988, is from the *Banco Central do Brasil*.

For Peru, a series for the CPI, available since January 1949, is from the *Banco Central de Reserva del Peru*, whereas a monetary aggregate defined as ‘Money plus quasi money’, available since January 1964, is from the International Monetary Fund’s *International Financial Statistics* (henceforth, IMF and *IFS*, respectively).

The data for Yugoslavia, available for the period Dec. 1990-Jan. 1994, are from Petrovic and Mladenovic (2000), and they were kindly provided by Zorica Mladenovic.¹⁸ The dataset features series for retail prices, M1, and the black market exchange rate of the Yugoslav dinar *vis-à-vis* the Deutsche Mark. As extensively discussed by Petrovic and Mladenovic (2000),¹⁹ however, several observations of the retail prices index are unreliable, and in what follows I will therefore exclusively work with the black market exchange rate of the dinar, which, as pointed out by the authors in footnote 10, was collected directly by them from daily newspapers.

The data for Congo are from the IMF’s *IFS*. A series for ‘Money’ is available since

difference of the relevant price index.

¹⁸The data’s original sources are discussed in detail in footnote 10 of Petrovic and Mladenovic (2000).

¹⁹See p. 787, and especially footnote 4.

May 1969, whereas a series for the exchange rate *vis-à-vis* the U.S. dollar is available since September 1983.

The data for Angola are from the IMF’s *IFS*. A series for the CPI is available since Jan. 1993, whereas a series for the central bank’s ‘Reserve Money’ is available since Dec. 1995.

The data for Bulgaria are from the IMF’s *IFS*. A series for the CPI is available since Jan. 1991, whereas a series for M1 is available since Dec. 1995.

The sources of the data for Zimbabwe are as follows. A series for ‘Reserve Money’, available for the period Jan. 1979–Jun. 2008, is from the IMF’s *IFS*. A series for the black-market exchange rate of Zimbabwe’s dollar *vis-à-vis* the U.S. dollar, which was used in McIndoe-Calder (2018), is from Table 4 of McIndoe-Calder (2011), and it is available for the period Apr. 2006–Apr. 2009. As discussed by McIndoe-Calder (2018, Section III, pp. 1661–1663) the official CPI series (which is available from the IMF’s *IFS*) is unreliable, especially for the latest stages of the hyperinflation. Exactly as for Yugoslavia, in what follows I will therefore exclusively focus on the black-market exchange rate.

For Venezuela, data for M1, available since January 2013, are from the *Banco Central de Venezuela*. As for the price index, since the government stopped publishing official CPI figures in December 2015, I was compelled, once again, to resort to the black-market exchange rate—in the present case, for the Bolívar *vis-à-vis* the U.S. dollar, which is available at the daily frequency at the website <https://dolartoday.com>.

The original data for Argentina, Chile, and Peru were seasonally unadjusted, and I have therefore seasonally adjusted them *via* ARIMA X-12.

4.2 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*²⁰—and it is from *Wirtschaft und Statistik* after that. Since 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), I removed it *via* ARIMA X-12. A *daily* series for the spot exchange rate of the German Reichsmark *vis-à-vis* the British Pound is available almost without interruptions from September 7, 1922 to November 15, 1923 from *Wirtschaft und Statistik*.²¹ 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 *exactly* match the dates in which the monetary aggregate had been released with the dates for the exchange rate. A weekly price

²⁰See at: https://www.destatis.de/GPStatistik/receive/DESerie_serien_00000012?list=all

²¹See at: https://www.destatis.de/GPStatistik/receive/DESerie_serien_00000012?list=all The table is labelled as ‘Bewegung der Wechselkurse: Wechselkursset Berlin auf’.

series (labelled as ‘Großhandindexziffer’, i.e. ‘Wholesale price index’), available from August 7, 1922 to November 15, 1923, is from *Wirtschaft und Statistik*.²² 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 often 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.

Two series for a monetary aggregate (labelled as ‘Notes’) and a price index for Hungary’s post-WWII episode, available from December 31, 1945 to July 23, 1946, are from the appendix of Anderson, Bomberger, and Makinen (1988).

I now turn to the evidence.

5 Evidence

‘Good empirical evidence tells its story regardless of the precise way in which it is analyzed. In large part, it is its simplicity that makes it persuasive. [...] No single test is held out as decisive. Many different types of data are examined. Mayer (1972) counts 16 different types of evidence adduced by Friedman in support of the permanent income hypothesis. No single episode in *A Monetary History* was held out as decisive. No single test reported in Fama’s survey proved or disproved anything but a persuasive pattern emerged from the totality.’

—Lawrence Summers (1991)

In building up my argument that

(I) available data for hyperinflationary episodes show no evidence of a Laffer curve for the revenue from money creation, and

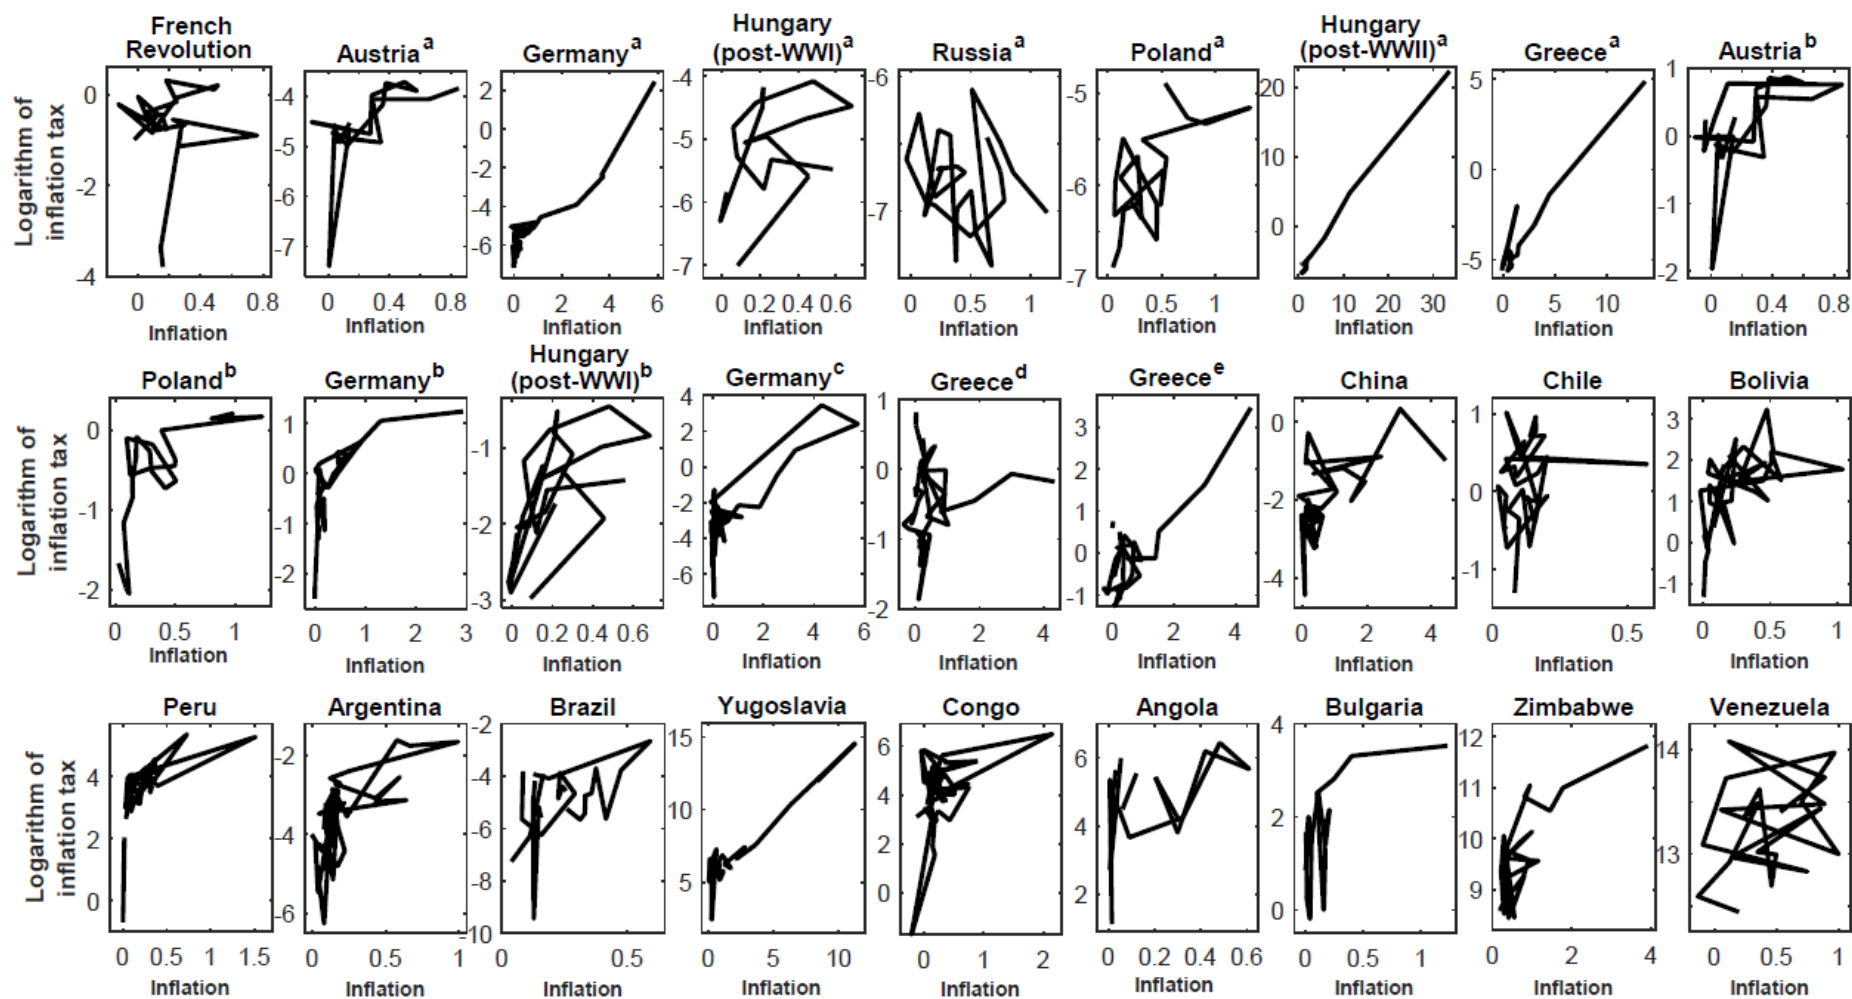
(II) the data, especially at very high inflation rates, show a clear preference for Meltzer’s log-log functional form (or a specification very close to it), rather than the semi-log proposed by Cagan (1956),

I start from the simplest kind of evidence—the raw data—and I then move to maximum likelihood estimation of the Box-Cox parameter in equation (2) at the level of individual countries. In doing so I am motivated by the conviction—forcefully articulated, e.g., by Summers (1991) in the above quotation—that the most convincing type of empirical evidence is the simplest.

5.1 A look at the raw data

Table 1 reports, for each individual episode, the maximum, mean, and median monthly inflation rate, computed as the log-difference (in natural logarithms) of the relevant

²²See at: https://www.destatis.de/GPStatistik/receive/DESerie_serie_00000012?list=all



^aBased on Cagan's data; ^bBased on Barro's data; ^cBased on Graham's data; ^dBased on Agapitides's data; ^eBased on Delivanis and Cleveland's data.

Figure 3 Inflation and the logarithm of the inflation tax for the hyperinflation episodes in the dataset (monthly data)

price index. 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 50 per cent per month.²³ The dataset features a significant extent of variation in inflationary experiences, ranging from post-WWII Hungary’s peak of 33.67;²⁴ 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 least extreme, Chile, with a maximum inflation rate of 0.575. 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.

5.1.1 Inflation and the revenue from money creation during hyperinflations

Figure 3 shows, for the 20 episodes in my dataset, the logarithm of the revenue from money creation plotted against the inflation rate.²⁵ Following Drazen (1985),²⁶ I compute the revenue from money creation as $\mu_t(M_t/P_t)$ in (1), which, based on end-of-period discrete-time observations, I approximate as

$$\theta_t \left[\frac{1}{2} \left(\frac{M_{t-1}}{P_{t-1}} + \frac{M_t}{P_t} \right) \right] \quad (5)$$

where $\theta_t \equiv (M_t - M_{t-1})/M_{t-1}$, and M_t and P_t are the end-of-period figures for nominal money balances and the price level, respectively. As before, the inflation rate is computed as the log-difference (in natural logarithms) of the relevant price index. For the seven episodes in Cagan’s dataset, the evidence in Figure 3 is qualitatively the same as that shown in Figure 1 based on Sargent and Wallace’s (1973) measure of seignorage. As for the remaining episodes and datasets, the crucial point to stress is that, once again, there is essentially *no evidence* of a Laffer curve for the inflation tax. Rather, two broad patterns appear to have characterized hyperinflation episodes:

²³In 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, respectively, when the inflation rate had been equal to 4.44 and 8.35.

²⁴Hungary’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.

²⁵Figure 3 only reports evidence based on monthly data. The corresponding evidence based on weekly data for Germany, and for post-WWII Hungary, is in line with that in Figure 3, and it is available upon request.

²⁶To be precise, 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 $\mu_t(M_t/P_t)$ in my equation (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 μ_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 $\mu_t(M_t/P_t)$.

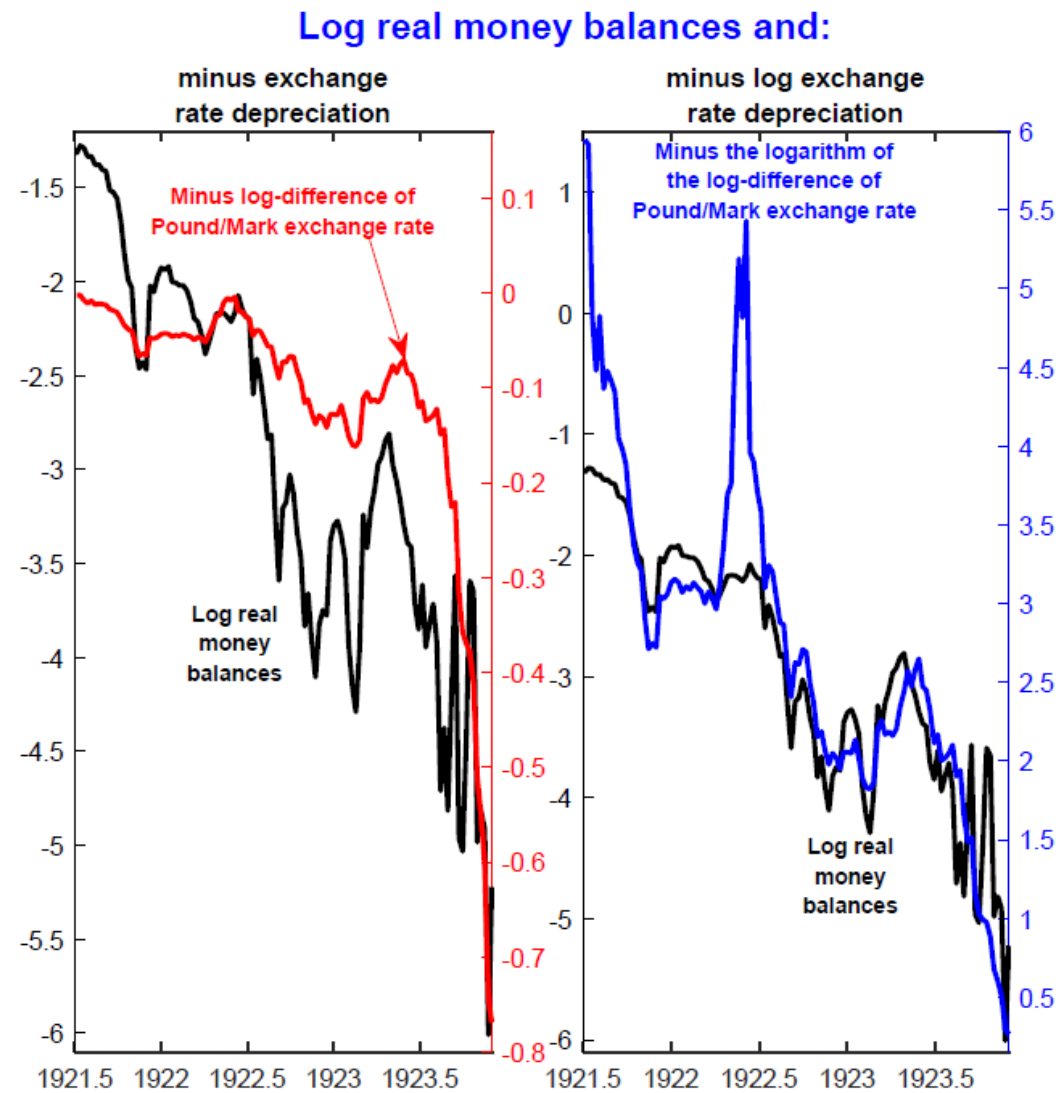


Figure 4a Germany, June 1921-November 1923: the weekly raw data for (log) exchange rate depreciation and log real money balances (12-month rolling averages)

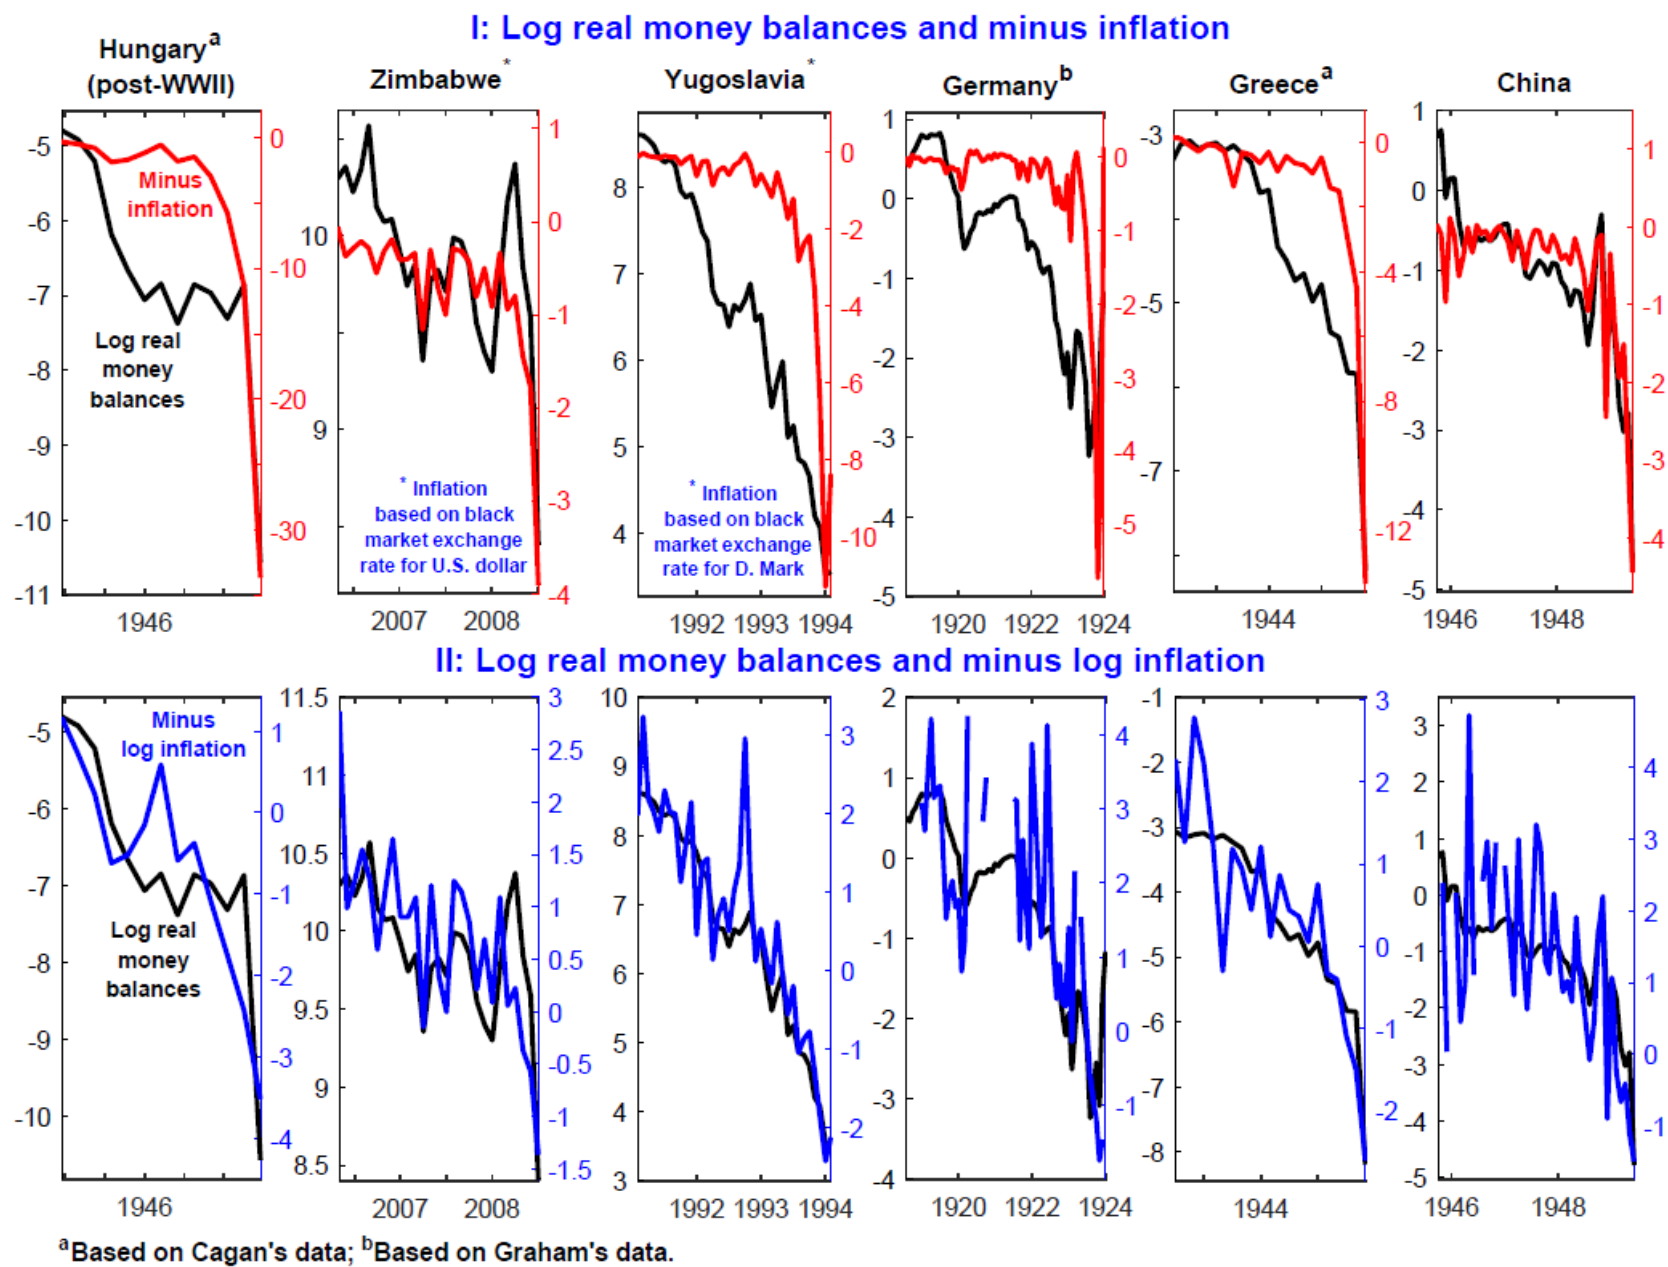


Figure 4b Selected monthly raw data for (log) inflation and log real money balances




Figure 4c Monthly raw data for (log) inflation and log real money balances for Latin-American countries

(i) in several instances—notably, post-WWII Hungary, Greece, Yugoslavia,²⁷ Germany, and Zimbabwe—the revenue from money creation appears to have been monotonically increasing with inflation.

(ii) In the remaining episodes—e.g., Austria, Poland, Bolivia, Argentina, and Bulgaria—the inflation tax appears to have flattened out at comparatively high inflation rates (i.e., towards the latest stages of the hyperinflation).

The implication is that, as long as we are willing to restrict our attention to the family of specifications encoded in equation (2), hyperinflation data are best described by values of θ between 0 and about 0.5.

5.1.2 The dynamics of real money balances and inflation during hyperinflations

Figures 4a-4c provide additional evidence on this, by showing, for some selected episodes, the logarithm of real money balances together with either the *level* of inflation (on the left-hand side panel in Figure 4a, and in the top row in Figures 4b-4c), or its *logarithm*²⁸ (on the right-hand side panel in Figure 4a, and in the bottom row in Figures 4b-4c). The evidence in the two groups of panels 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.

In several instances the log-log specification provides a manifestly more plausible description of the joint dynamics of real money balances and inflation, in the specific sense that log real balances track the *logarithm* of inflation much more closely than its *level*.²⁹ This is especially clear for Germany based on the weekly data for the Pound/Mark exchange rate (see Figure 4a), and for Yugoslavia, Germany, Greece, and post-WWII Hungary, whereas evidence for China and Zimbabwe is slightly weaker (see Figure 4b). On the other hand, evidence for other countries is, most of the times, not clear-cut, which calls for econometric methods in order to be able to discriminate between alternative functional forms.

²⁷An important point to stress is that Hungary’s 1946 episode, Greece, and Yugoslavia are the three most extreme cases of hyperinflation in recorded history (see Table 1). The fact that, as documented in Figure 3, for any of these episodes the relationship between inflation and the inflation tax had consistently been positive at all inflation rates represents therefore a stark refutation of the notion that when inflation crosses a certain threshold, further increases lead to a fall in seigniorage.

²⁸Because of the very high-frequency of the data, even if we are here dealing with hyperinflationary episodes, in a few instances the price level had decreased from one month to the next, so that the inflation rate turned out to be negative. E.g., for Germany this was the case (based on Cagan’s data) not only at the very beginning of the hyperinflation, for all months between December 1920 and May 1921, but also in March 1923, well into the most virulent phase of the hyperinflation. In all of these cases, the corresponding observations for log inflation are plotted as missing.

²⁹Up to a scale factor, which in the figures is accounted for by allowing the right hand-side and the left hand-side scales to differ.

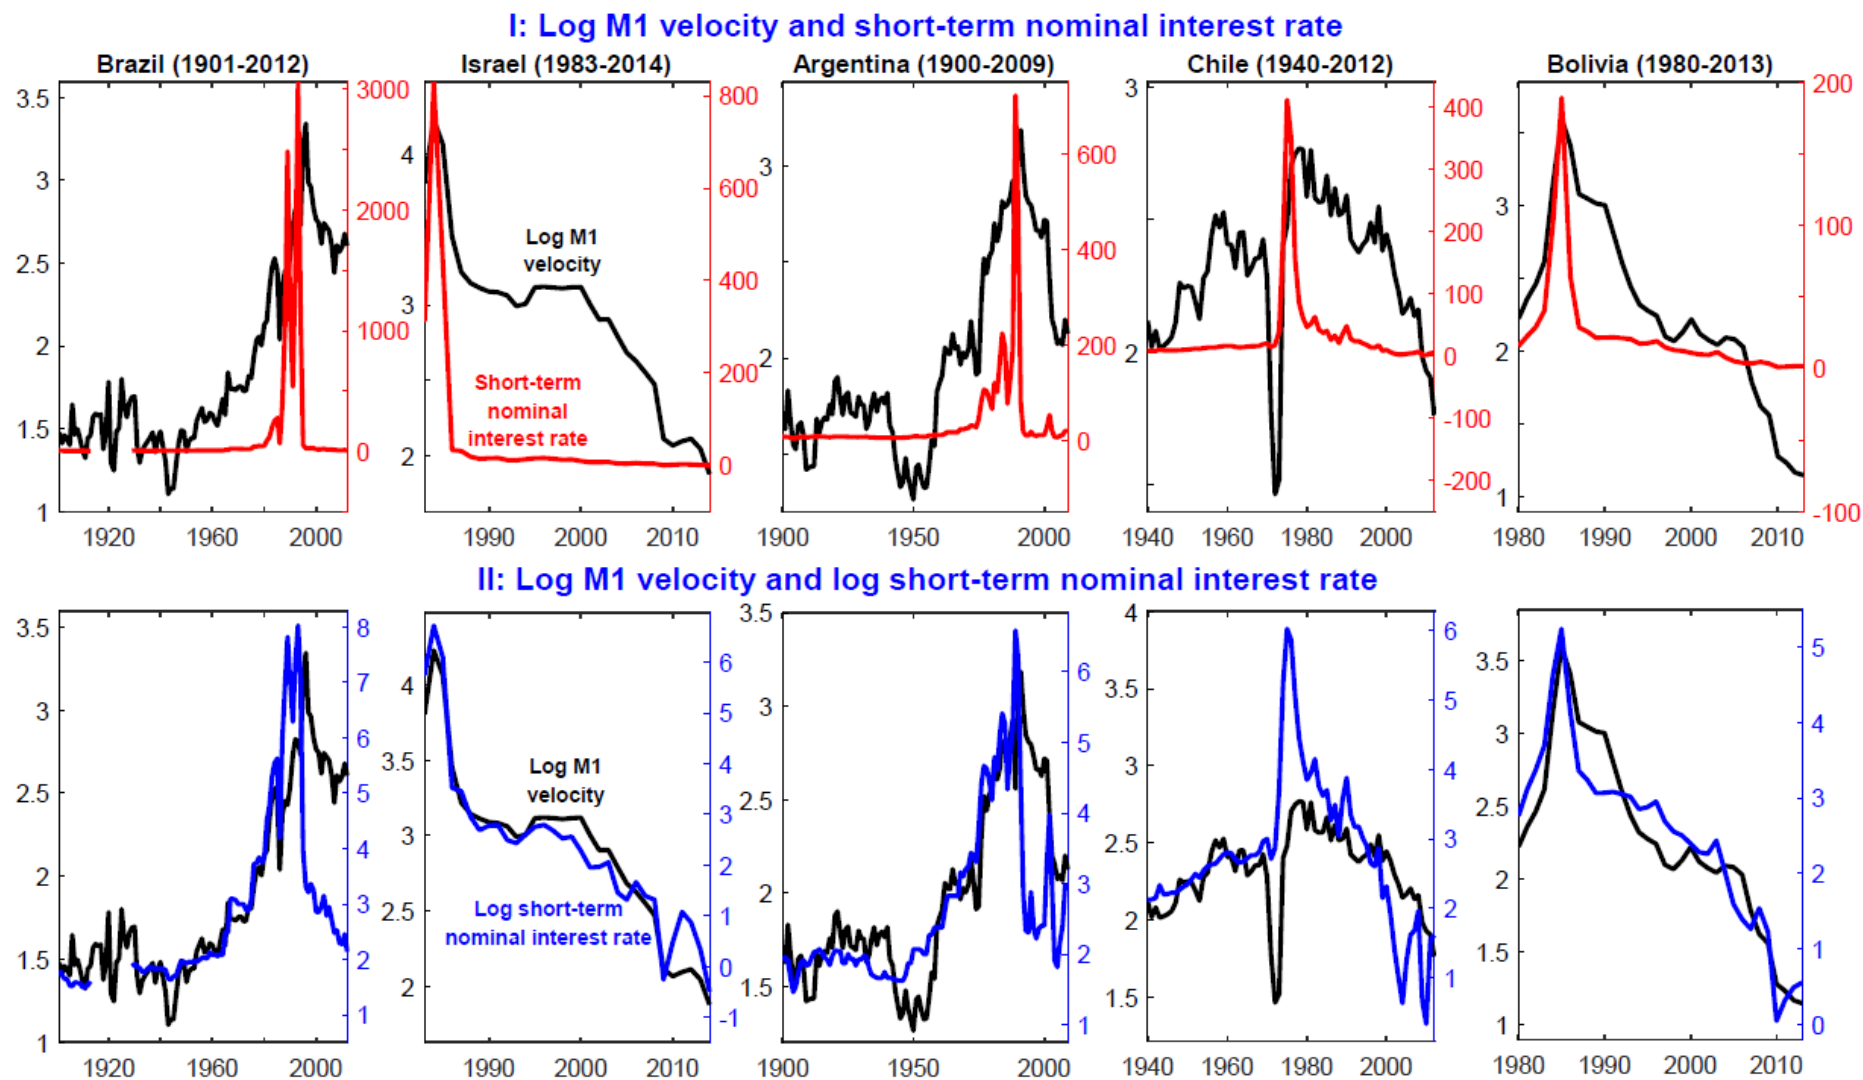


Figure 5a Long-run data for M1 velocity and a short-term nominal interest rate for high-inflation countries

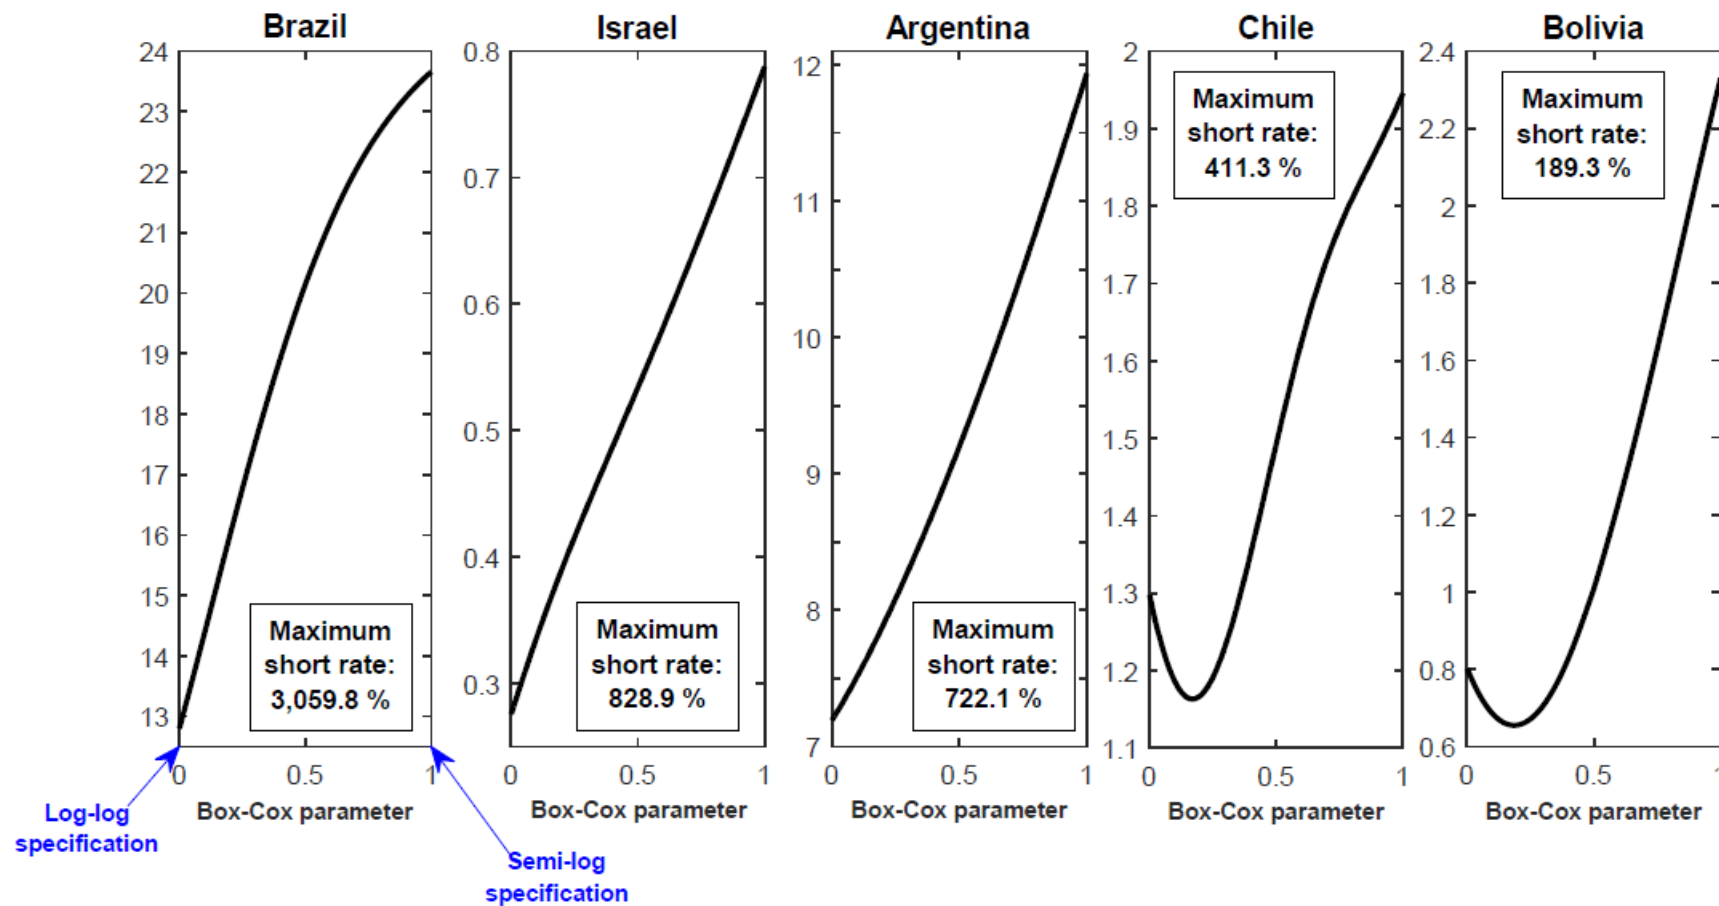


Figure 5b Sum of squared residuals from cointegration-based estimates of long-run money demand specifications, for alternative values of the Box-Cox parameter for the short rate

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

Figures 5a-5b provide additional evidence along these lines based on long-run data for M1 velocity³⁰ and a short-term nominal interest rate for five high-inflation countries from Benati, Lucas, Nicolini, and Weber’s (2018) dataset.³¹ 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.

The evidence in the figure speaks for itself: whereas fluctuations in log M1 velocity bear a uniformly 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 both Israel and Bolivia, for which the logarithms of M1 velocity and 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; it is apparent 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 millennium; and it is very apparent 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.

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 (although, as my econometric results will show, most of the times this is in fact *not* the case). 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.

An important counter-argument to this is that, as the economy approaches, and then enters a full-blown hyperinflation, the functional form of 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

³⁰I 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 which 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 constant. 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.

³¹The sources of the data are described in detail in Appendix B of Benati *et al.* (2018).

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 4a should be regarded as uninformative, and only data pertaining to hyperinflation episodes would be relevant for the issue of discriminating between alternative functional forms for the demand for money during hyperinflations. As already mentioned, however—and as we will see in Section 5.2.4—even when narrowly focusing on hyperinflation episodes, empirical evidence near-uniformly suggests that the most plausible characterization of the data is provided by Meltzer’s log-log (or a specification close to it), rather than by Cagan’s semi-log. This means that even if the previously mentioned objection to the evidence reported in Figure 5a is correct in principle, in fact it appears to be incorrect in practice.

Figure 5b provides evidence on the comparative plausibility of alternative values of θ in specification

$$\ln(V_{1,t}) \equiv \ln\left(\frac{Y_t}{M_{1,t}}\right) = A + B\frac{R_t^\theta - 1}{\theta}, \quad (6)$$

for the logarithm of M1 velocity, $V_{1,t}$, as a function of the Box-Cox transformation of the nominal interest rate, R_t , with A and B being parameters to be estimated. Specifically, the figure shows, for alternative values of θ between 0 and 1, the sum of the squared residuals from estimates of equation (6).³² In line with the previous discussion of the visual evidence in Figure 5a, the evidence in Figure 5b uniformly points towards the specification preferred by the data (in the sense of being associated with the smallest sum of squared residuals) being either the log-log—for Brazil, Israel, and Argentina, i.e., the countries with the highest nominal rates—or a specification very close to it for the two countries with comparatively lower interest rates.

Overall, the evidence in Figures 5a-5b points towards Meltzer’s log-log functional form, or a specification very close to it, as providing the most plausible characterization of the joint dynamics of M1 velocity and nominal interest rates. Sure enough, this is based on just five high-inflation countries. In the spirit of Summers’ (1991) previous quotation, however, my objective here is to pursue ‘simplicity’, to look at ‘different types of evidence’, in the hope that ‘a persuasive pattern emerge[s] from the totality’. Under this respect, the evidence in Figures 5a-5b is clearly valuable.

I now move to maximum likelihood estimation of the Box-Cox parameter in equation (2), by working within a rational expectations framework in the spirit of Salemi and Sargent (1979).

³²The results reported in Figure 4b are based on Stock and Watson’s (1993) ‘dynamic OLS’ estimator, but qualitatively the same evidence is produced by simple OLS estimation of equation (6). This evidence is available upon request.

5.2 Estimating the Box-Cox parameter

5.2.1 Methodological issues

Before describing the approach adopted in the present work, it is worth spending a few words discussing why I have chosen to eschew cointegration methods, which, following Taylor (1991), have been used in several subsequent studies of hyperinflations.³³ At first sight, cointegration presents crucial advantages compared to the rational-expectations approach used by Sargent (1977), and especially by Salemi and Sargent (1979). As discussed by Taylor (1991), 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. This advantage, however, has to be balanced against two crucial shortcomings:

first, with the small samples which are typical of hyperinflationary episodes,³⁴ the reliability of cointegration methods is entirely open to question. For example, even bootstrapping the test statistics³⁵ as in Cavaliere *et al.* (2012), the Monte Carlo evidence in Benati *et al.*'s (2018) Appendix E.3.1 suggests that Johansen's tests will not detect cointegration—if it is there—a large, or even very large fraction of the times.³⁶

Second, as mentioned in footnote 28, because of the very high-frequency of the data inflation typically turns out to be negative for a few months (or weeks) in most episodes.³⁷ This implies that since cointegration methods use the *actual* inflation rate—as opposed to the *expected* inflation rate which enters the theoretical functional

³³See, e.g., Engsted (1994), Michael, Nobay, and Peel (1994), and Petrovic and Vujosevic (1996).

³⁴E.g., based on Cagan's dataset the number of observations is equal to 39 for Germany, 20 for Austria, and just 13 for post-WWII Hungary.

³⁵To the very best of my knowledge, all cointegration-based studies of hyperinflations have been based on asymptotic critical values, which, as illustrated by Johansen (2002) with reference to his trace and maximum eigenvalue tests, are essentially unreliable in small samples.

³⁶So the argument is the same as in Engle and Granger (1987). For example, in spite of the visual evidence of a strong correlation between log real money balances and log inflation for Greece (see Figure 4b) Johansen's trace and maximum eigenvalue tests, bootstrapped as in Cavaliere *et al.* (2012) do not detect cointegration between the two series (with p -values equal to 0.469 and 0.535 respectively). The most likely explanation is the short sample period: with just 21 observations, by Engle and Granger's (1987) argument detecting cointegration is clearly problematic. By the same token, Johansen's tests do not detect cointegration between the two series for (e.g.) China, Poland, and Brazil, in spite of the visual evidence of a strong correlation between them in Figures 4b-4d. On the other hand, cointegration is strongly detected (e.g.) for Yugoslavia, with bootstrapped p -values for Johansen's tests equal to 0.043 and 0.037, respectively.

³⁷Based on monthly data, inflation has consistently been positive for Poland, post-WWII Hungary, and Greece (starting from the third observation in the sample) based on Cagan's data; for Poland and Germany (starting from the sixth observation) based on Barro's data; and for Yugoslavia, Zimbabwe, Brazil, Peru, Chile, Argentina (except for the very last observation), and China (with the exception of four months).

form (2)—they can only be applied to the semi-log specification. This, however, is unduly restrictive: *precisely* because we are dealing with hyperinflations, we can safely assume that expected inflation had consistently been positive throughout the entire episode even if the actual inflation rate had turned out negative in a few months or weeks. Since expected inflation is unobserved, however, this assumption is ultimately not operational, unless the dynamics of inflation expectations is explicitly modelled, which is what I do in the next section.

5.2.2 An approach in the spirit of Salemi and Sargent (1979)

As discussed in Section 3.2, 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 what follows 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 rational expectations 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 log real money balances and the Box-Cox transformation of expected inflation as $\tilde{m}_t \equiv \ln(M_t/P_t)$ and

$$\tilde{\pi}_t \equiv \frac{[\pi_{t+1|t}]^\theta - 1}{\theta} \quad (7)$$

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

$$\tilde{m}_t = \beta + \alpha \tilde{\pi}_t + u_t \quad (8)$$

where u_t is a money demand disturbance. Whereas Sargent (1977) and Salemi and Sargent (1979) postulate that u_t is a random walk,³⁸ in what follows I assume that it evolves according to

$$u_t = \rho u_{t-1} + v_t, \quad |\rho| \leq 1, \quad (9)$$

i.e. a specification nesting the random walk case, but also allowing for stationarity. The key reason³⁹ for doing so is the evidence in Figures 5a-5d, especially for countries such as (e.g.) Yugoslavia. A comparison between the two panels in the third column of Figure 5b naturally suggests (at least) two possible interpretations of the

³⁸As discussed in Section 3, theoretical models of seignorage in which the governments finances, *via* the inflation tax, a constant fraction of GDP, produce a random-walk disturbance in the money demand equation.

³⁹Christiano (1987) produces some evidence against the random walk assumption based on Cagan’s data for Germany’s episode.

joint dynamics of Yugoslavia's 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, as previously discussed, 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 inflation⁴⁰ and real money balances in the bottom panel track each other very closely,⁴¹ to the point that, as mentioned in footnote 36, Johansen's tests detect cointegration between the two series at the 5 per cent level. Finally, by the same token, simple logic suggests that for any value of θ between 0 and 1 in (7), the higher θ , the higher ρ ought to be (and *vice versa*) for specification (8)-(9) to provide a plausible characterization of the relationship between expected inflation and real money balances. This logically implies that imposing $\rho=1$ automatically 'stacks the cards' in favor of the semi-log, and against the log-log, so that estimates of θ 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 (9), 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 \tilde{m}_t and $\tilde{\pi}_t$, conceptually in line with Hamilton (1985) and Burmeister and Wall (1982, 1987) I postulate that it is described by

$$\begin{bmatrix} \tilde{m}_t \\ \tilde{\pi}_t \end{bmatrix} = \begin{bmatrix} \bar{m} \\ \bar{\pi} \end{bmatrix} + \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \tilde{m}_{t-1} \\ \tilde{\pi}_{t-1} \end{bmatrix} + \begin{bmatrix} H_{\tilde{m}} \\ H_{\tilde{\pi}} \end{bmatrix} u_t + \begin{bmatrix} \epsilon_t^{\tilde{m}} \\ \epsilon_t^{\tilde{\pi}} \end{bmatrix} \quad (10)$$

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.

In the spirit of Salemi and Sargent (1979), it can be easily shown that the rational expectations hypothesis, together with equation (8), imposes the following restrictions upon (10):

$$\bar{m} = \beta + \alpha \bar{\pi} \quad (11)$$

$$a(L) = \alpha c(L) \quad (12)$$

$$b(L) = \alpha d(L) \quad (13)$$

and

$$H_{\tilde{m}} = 1 + \alpha H_{\tilde{\pi}} \quad (14)$$

⁴⁰ Although the theoretical relationship in equation (7) 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.

⁴¹ Obviously, once appropriately rescaled, which in Figures 5a-5d is implicitly implemented by allowing for different scales in the left- and right-hand side axes.

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 (8) for the demand for real money balances, and the equation for \tilde{m}_t in the VAR representation (10), produces the following restriction for the error term for \tilde{m}_t in (10):

$$\epsilon_t^{\tilde{m}} = \alpha \epsilon_t^{\tilde{\pi}} \quad (15)$$

By defining the state vector as $\xi_t = [u_t, \tilde{\pi}_t, \tilde{\pi}_{t-1}, \dots, \tilde{\pi}_{t-p+1}]'$, where p is the lag order in the lag polynomials $a(L)$, ..., $d(L)$,⁴² the model can be cast in state-space form, with state equation

$$\begin{aligned} \underbrace{\begin{bmatrix} u_t \\ \tilde{\pi}_t \\ \tilde{\pi}_{t-1} \\ \dots \\ \tilde{\pi}_{t-p+2} \\ \tilde{\pi}_{t-p+1} \end{bmatrix}}_{\xi_t} &= \underbrace{\begin{bmatrix} 0 \\ \tilde{\pi} \\ 0 \\ \dots \\ 0 \\ 0 \end{bmatrix}}_f + \underbrace{\begin{bmatrix} \rho & 0 & 0 & \dots & 0 \\ 0 & d_1 & d_2 & \dots & d_p \\ 0 & 1 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix}}_F \underbrace{\begin{bmatrix} u_{t-1} \\ \tilde{\pi}_{t-1} \\ \tilde{\pi}_{t-2} \\ \dots \\ \tilde{\pi}_{t-p+1} \\ \tilde{\pi}_{t-p+2} \end{bmatrix}}_{\xi_{t-1}} + \\ &+ \underbrace{\begin{bmatrix} 0 & 0 & \dots & 0 \\ c_1 & c_2 & \dots & c_p \\ 0 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \dots & 0 \end{bmatrix}}_G \underbrace{\begin{bmatrix} \tilde{m}_{t-1} \\ \tilde{m}_{t-2} \\ \dots \\ \tilde{m}_{t-p} \end{bmatrix}}_{M_t} + \underbrace{\begin{bmatrix} v_t \\ \epsilon_t^{\tilde{\pi}} \\ 0 \\ \dots \\ 0 \\ 0 \end{bmatrix}}_{w_t} \end{aligned} \quad (16)$$

with

$$Q \equiv E(w_t w_t') = \begin{bmatrix} \sigma_v^2 & 0 & 0 & \dots & 0 \\ 0 & \sigma_{\tilde{\pi}}^2 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} \quad (17)$$

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

As for the observation equations, the first is given by

$$\pi_t = \pi_{t|t-1} + \eta_t = [\theta \tilde{\pi}_{t-1} + 1]^{\frac{1}{\theta}} + \eta_t \quad (18)$$

⁴²As 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 $p=4$ with weekly data, and 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).

where η_t is a rational expectations forecast error, i.e. $\eta_t = \pi_t - \pi_{t|t-1}$, so that $\eta_{t|t-1} = 0$, 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,

$$Z \equiv E(\eta_t w_t') = \begin{bmatrix} \sigma_{\eta v} & \sigma_{\eta \tilde{\pi}} & 0 & \dots & 0 \end{bmatrix} \quad (19)$$

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,⁴³ 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 (10), whereas the second is to simply use equation (8), which can be rewritten as

$$\tilde{m}_t = \beta + [1 \ \alpha \ 0 \ \dots \ 0] \xi_t. \quad (20)$$

Because of restriction (15), the two representations for \tilde{m}_t are equivalent. In what follows I will use equation (20).

The state-space model described by equations (16), (18), and (20) is linear with the single exception of the observation equation for inflation, expression (18). Following (e.g.) Harvey (1989), I therefore compute the log likelihood based on the ‘generalized Kalman filter’. Specifically, I take a first-order Taylor expansion of (18) around $\xi_{t|t-1}$, thus obtaining the following approximate expression for the observation equation for inflation

$$\begin{aligned} \pi_t &\simeq \underbrace{[\theta \tilde{\pi}_{t-1|t-1} + 1]^{\frac{1}{\theta}} - \tilde{\pi}_{t-1|t-1} [\theta \tilde{\pi}_{t-1|t-1} + 1]^{\frac{1-\theta}{\theta}}}_{c_\pi} + [\theta \tilde{\pi}_{t-1|t-1} + 1]^{\frac{1-\theta}{\theta}} \tilde{\pi}_{t-1} + \eta_t = \\ &= c_\pi + [\theta \tilde{\pi}_{t-1|t-1} + 1]^{\frac{1-\theta}{\theta}} \tilde{\pi}_{t-1} + \eta_t \end{aligned} \quad (21)$$

Replacing the original non-linear observation equation (18) with (21) 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).

5.2.3 Maximum likelihood estimation

I maximize the log-likelihood numerically *via* simulated annealing.⁴⁴ Having found the parameter vector which maximizes the likelihood, \hat{B}_{MLE} , rather than relying on

⁴³E.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.

⁴⁴Specifically, 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

asymptotic formulas, I 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.⁴⁵ Specifically, the proposal draw for B, \tilde{B} , is accepted with probability $\min[1, r(B_{s-1}, \tilde{B} | Y)]$ —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

$$r(B_{s-1}, \tilde{B} | Y) = \frac{L(\tilde{B} | Y)}{L(B_{s-1} | Y)} \quad (22)$$

which uniquely involves the likelihood.⁴⁶ All other estimation details are identical to Benati (2008), to which the reader is referred to. I use 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. For all episodes and datasets, the fractions of accepted draws are uniformly very close to the 23 per cent ideal acceptance rate in high dimensions⁴⁷, and the draws exhibit little autocorrelation based on two alternative statistics, the first autocorrelation, and the draws’ inefficiency factors.⁴⁸

5.2.4 Evidence

Figure 6 shows, for each of the 20 episodes I consider, the distribution of the MLE estimate of the Box-Cox parameter, θ , generated by RWM, whereas Figure 7 shows scatterplots of the actual observations for the logarithm of the inflation tax against inflation⁴⁹ (the blue dots), together with the median, and the 16-84 and 5-95 percentiles of the distribution of the theoretical functional relationship between the two

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.

⁴⁵So what I am doing can be interpreted as Bayesian estimation *via* RWM of model (20), (18), and (16) with flat priors for all parameters.

⁴⁶With Bayesian priors it would be

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

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

⁴⁷See Gelman, Carlin, Stern, and Rubin (1995).

⁴⁸The 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.

⁴⁹So, to be clear, these data are the same as those shown in Figure 1. For Germany, Austria, Poland, and post-WWI Hungary Figure 7 shows results based on Cagan’s data: the corresponding

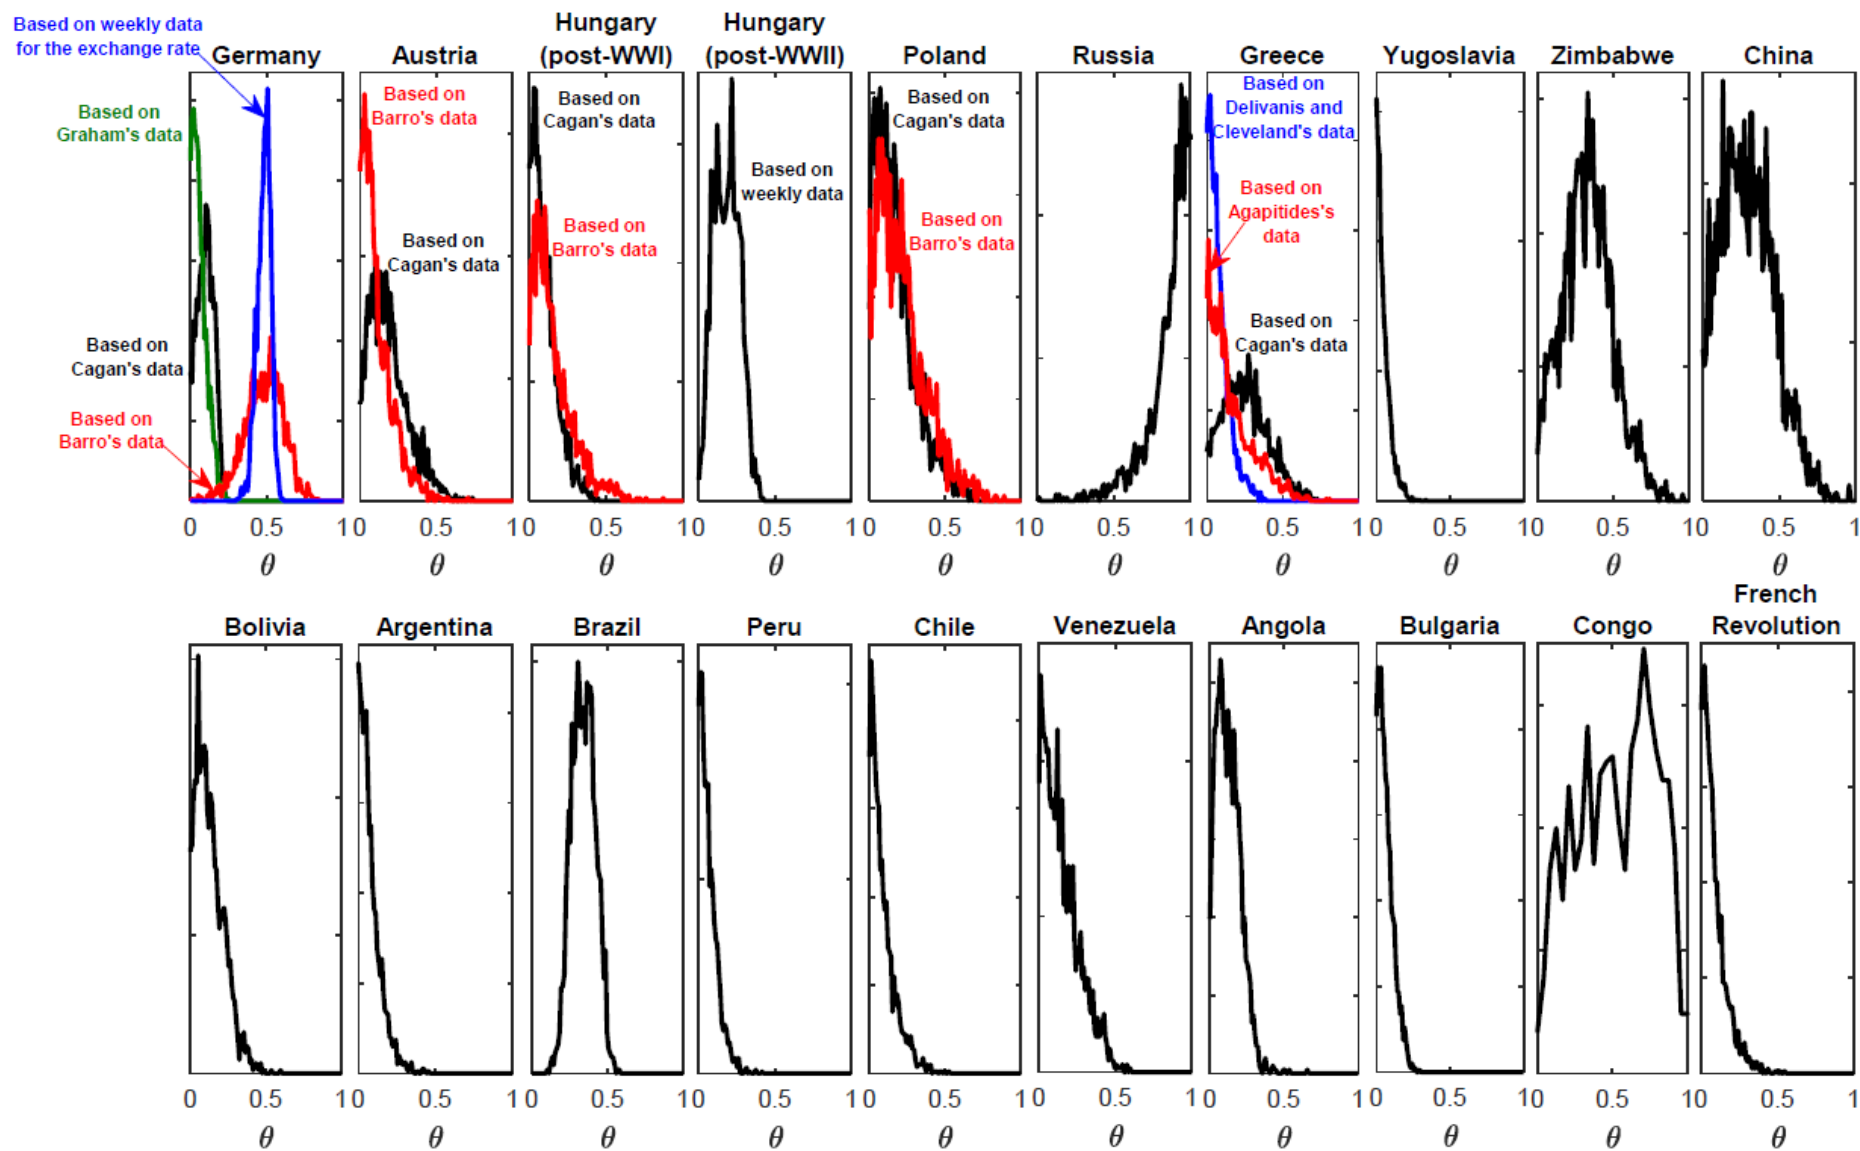


Figure 6 Maximum likelihood estimates: distributions of the Box-Cox parameter's draws generated via Random-Walk Metropolis

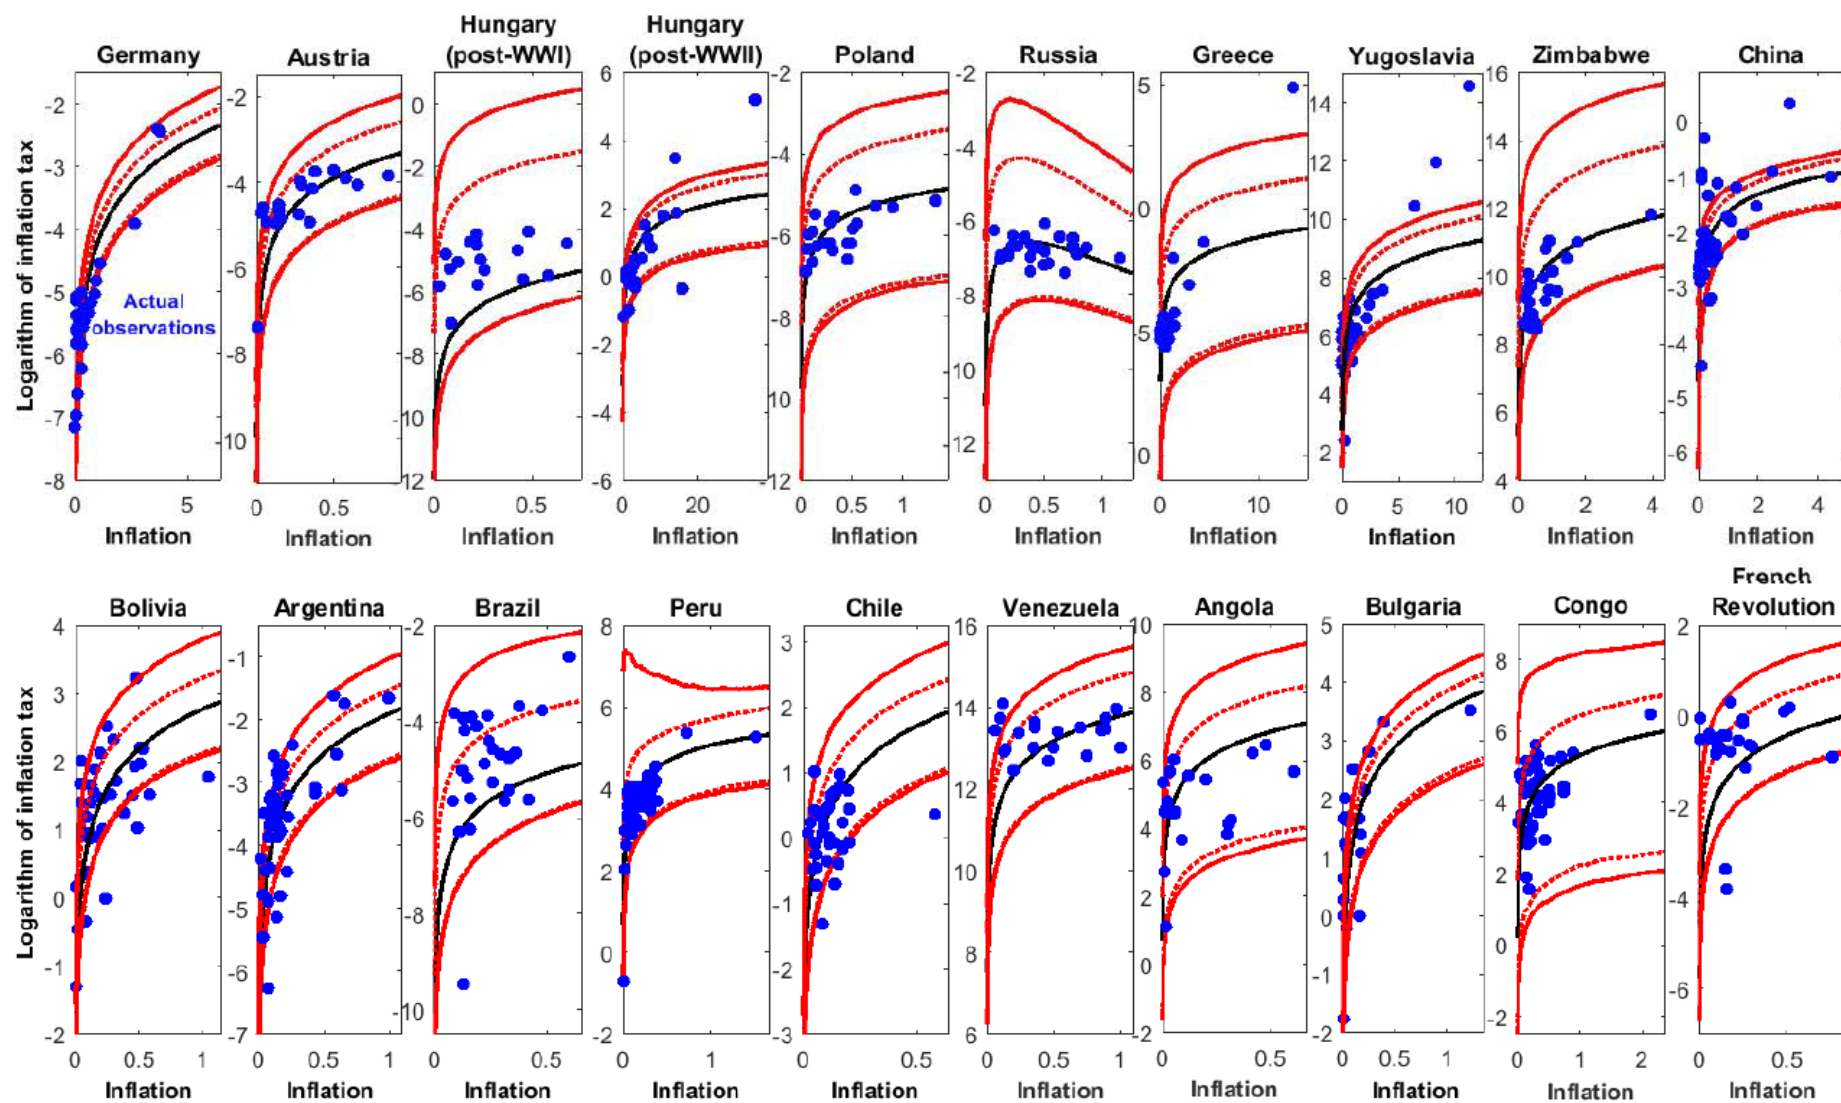


Figure 7 The estimated relationship between inflation and the logarithm of the inflation tax: median, and 16-84 and 5-95 percentiles of the distributions generated via Random-Walk Metropolis

series implied by the MLE estimates generated by RWM. The evidence in the two figures points towards the following facts:

(I) between Cagan’s semi-log specification, and Meltzer’s log-log, the data near-uniformly, and quite clearly prefer the latter, or a specification very close to it.⁵⁰ This is testified by the fact that in just a *single* case (Russia) the distribution of the MLE estimate of θ is clustered towards 1, whereas in *all* other cases⁵¹ it is clearly far away from 1, sometimes starkly so.⁵² Further, in the overwhelming majority of cases, the distribution of the MLE estimate of θ is clustered towards 0. This is the case for Germany based on either Cagan’s or Graham’s data; for Austria, Poland, and post-WWII Hungary based on either Cagan’s or Barro’s data; for Greece based on either Agapitides’, or Delivanis and Cleveland’s data, whereas evidence based on Barro’s data is slightly weaker; for Yugoslavia, Bolivia, Argentina, Peru, Chile, Venezuela, Angola, Bulgaria, and the French Revolution. For Brazil, and for post-WWII Hungary based on weekly data,⁵³ the distribution, although not clustered towards zero, is nonetheless firmly below 0.5, which, based on the results in Figure 2, points towards a uniformly slightly increasing functional relationship between inflation and seignorage. By the same token, for both China and Zimbabwe the bulk of the mass of the distribution is below 0.5, with median estimates around 0.3-0.35. Finally, in one case (Congo) the distribution is so spread out that it is essentially impossible to make any statement about θ with any degree of confidence.

(II) Consistent with (I), the evidence in Figure 7 points towards a *single* case—once again, Russia—in which the relationship between inflation and the logarithm of the inflation tax exhibits a Laffer curve for the latter. In *all* other cases the relationship is monotonically increasing, either strongly so—as for, e.g., Germany, Bolivia, Argentina, and Chile—or more gently, as in the case of, e.g., Venezuela. Additional evidence in Figure A.1 (see the discussion in footnote 49) confirms the broad qualitative features of Figure 7, in terms of the absence of a Laffer curve for the revenue from money creation for all the episodes and datasets reported therein.

My overall conclusion is therefore that Cagan’s paradox is simply an illusion, originating from the literature’s predominant focus on the semi-log specification for the

set of results based on Barro’s data are shown in Figure A.1 in the Appendix. For Germany, results based on either Graham’s data, or weekly data for the nominal exchange rate, are also shown in Figure A.1. For Greece I also show results based on Cagan’s data: the corresponding results based on either Agapitides’, or Delivanis and Cleveland’s data, are shown in Figure A.1.

⁵⁰In the sense of being associated with a value of θ close to 0.

⁵¹With the partial, and weak exception of Congo.

⁵²I do not provide numbers (i.e., specific percentiles of the distributions of θ generated by RWM) because the visual evidence is so stark, but they are available upon request.

⁵³For post-WWII Hungary I estimate the model only based on weekly data because the monthly sample (from Cagan’s dataset) features only 13 observations, with the result that estimation turned out to be problematic. Both Sargent, 1977, and Salemi and Sargent, 1979 estimated *via* maximum likelihood VARs encoding the restrictions imposed by Cagan’s semi-log specification, conditional on $\theta=1$. It appears however that, with just 13 observations, allowing θ to be a free parameter to be estimated is problematic.

demand for real money balances, which automatically imposes upon the data a Laffer curve for the inflation tax. On the other hand, the fact that, at very high inflation rates, the preferred specification is Meltzer’s (1963) log-log—or a specification very close to it—logically rules out the possibility of such paradox.

6 Conclusions, and 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. In this paper I have revisited Cagan’s paradox based on data from 20 hyperinflations—from the French Revolution to Venezuela’s 2018 episode—reporting two main findings. *First*, in the raw data there is nearly no evidence of a Laffer curve for the revenue from money creation. Rather, in the vast majority of cases the relationship between inflation and the inflation tax has been either positive at all inflation rates, or initially positive and then flattening out towards the end of the hyperinflation. *Second*, consistent with this, econometric evidence shows that the preferred money demand specification at very high inflation rates is not Cagan’s (1956) ‘semi-log’, which automatically imposes upon the data a Laffer curve for the inflation tax: rather, it is either Meltzer’s (1963) ‘log-log’—for which seignorage is monotonically increasing in inflation—or a more general functional form making log real money balances a linear function of the Box-Cox transformation of expected inflation (of which the ‘log-log’ is a special case), which allows for the inflation tax to flatten out at high inflation rates. My results therefore suggest that the paradox first highlighted by Cagan—of policymakers seemingly inflating in excess of the revenue-maximizing rate during hyperinflations—is the product of the literature’s predominant focus on the semi-log functional form. Evidence in favor of the log-log specification becomes overwhelming when considering samples materially longer than the narrow window of time which has typically been associated with hyperinflations. This is the case for countries such as Argentina, Brazil, Bolivia, Chile, and Israel.

In terms of directions for future research, the dataset I have here assembled naturally lends itself to (at least) three:

(I) an exploration of the issue of whether, during hyperinflations, the economy may have been operating under *indeterminacy*, so that hyperinflationary 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 empirical 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), but based on either a log-log specification for the demand for real money balances,

or functional form (2) in Section 2 of this paper, for values of θ close to 0. The next step would then be to estimate the model based on the dataset I have been using in the present work, possibly expanded with additional series pertaining to government finances. To my knowledge, such series are available at least for Germany’s episode (from *Wirtschaft und Statistik*, the publication discussed in Section 4.2).

(II) A second natural direction for future research is to exploit this dataset to reconsider the issue of whether hyperinflations may have been characterized by *explosive* behaviour.

(III) Finally, another issue to explore is the following. Both the evidence from Sargent and Wallace’s (1973) measure of the revenue from money creation for the episodes in Cagan’s (1956) dataset, and the broader evidence reported in Figure 3, suggest that for at least some episodes—notably, Greece, and post-WWII Hungary—the relationship between the logarithm of the inflation tax and inflation is essentially linear, with little to no evidence of *convexity*. As Figure 2 shows, however, even Meltzer’s ‘log-log’ specification does indeed feature a convex relationship between the two series, which implies that, by working within the family of functional forms characterized by (2), it is impossible to replicate the apparent lack of convexity found for some episodes. It would therefore be interesting to explore alternative functional forms for the demand for real money balances, which might be able to generate a (near) linear relationship between inflation and the logarithm of the inflation tax at very high inflation rates.

References

- AGAPITIDES, S. (1945): “The Inflation of the Cost of Living and Wages in Greece during German Occupation,” *International Labor Review*, December, 643–651.
- 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.
- BARRO, R. J. (1970): “Inflation, the Payments Period, and the Demand for Money,” *Journal of Political Economy*, 78(6), 1228–1263.
- BAUMOL, W. J. (1952): “The Transactions Demand for Cash: An Inventory Theoretic Model,” *Quarterly Journal of Economics*, 66, 545–372.
- BELONGIA, M., AND P. IRELAND (2018): “The Demand for Divisia Money: Theory and Evidence,” University of Mississippi and Boston College, mimeo.
- 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 (2018): “Long-Run Money Demand Redux,” *University of Bern, University of Chicago, Minneapolis FED, and University of South Carolina, mimeo*.
- 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. —.
- CAVALIERE, G., A. RAHBK, 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.
- DELIVANIS, D., AND W. C. CLEVELAND (1949): *Greek Monetary Developments, 1939-1948*. Bloomington: University of Indiana Press, 1949.

- DRAZEN, A. (1985): “A General Measure of Inflation Tax Revenues,” *Economics Letters*, 17, 327–330.
- ENGLE, R. F., AND C. W. GRANGER (1987): “Cointegration and Error Correction: Representation, Estimation, and Testing,” *Econometrica*, 55(2), 251–276.
- ENGSTED, T. (1994): “The Classic European Hyperinflations Revisited: Testing the Cagan Model Using a Cointegrated VAR Approach,” *Economica*, 61(243), 331–341.
- 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.
- MCINDOE-CALDER, T. (2011): “Hyperinflation in Zimbabwe: Money Demand, Seigniorage and Aid Shocks,” at: <http://www.academia.edu/595168/>.
- (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.
- MICHAEL, P., A. NOBAY, AND D. PEEL (1994): “The German Hyperinflation and the Demand for Money Revisited,” *International Economic Review*, 35(1), 1–22.

- 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.
- PETROVIC, P., AND Z. VUJOSEVIC (1996): "The Monetary Dynamics in the Yugoslav Hyperinflation of 1991-1993: The Cagan Money Demand," *European Journal of Political Economy*, 12, 467–483.
- 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.
- 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, pp. —.
- SELDEN, R. T. (1956): "Monetary Velocity in the United States," in M. Friedman, ed., *Studies in the Quantity Theory of Money*, University of Chicago Press, pp. 405–454.
- SIDRAUSKI, M. (1967a): "Inflation and Economic Growth," *Journal of Political Economy*, 75(6), 796–810.
- (1967b): "Rational Choice and Patterns of Growth in a Monetary Economy," *American Economic Review*, 57(2), 534–544.
- SIMS, C. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.
- STOCK, J. H., AND M. W. WATSON (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61(4), 783–820.
- SUMMERS, L. H. (1991): "The Scientific Illusion in Empirical Macroeconomics," *The Scandinavian Journal of Economics*, 93(2), 129–148.
- 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.

Figures for Online Appendix




Figure A.1 Additional results on the estimated relationship between inflation and the logarithm of the inflation tax: median, and 16-84 and 5-95 percentiles of the distributions generated via Random-Walk Metropolis