Estimating a Cagan-type Demand Function for Gold: 1561–1913

Alexei Deviatov, *Russian Presidential Academy of National Economy and Public Administration*
deviatov@list.ru

Long time series on gold production and the value of gold, taken from Jastram's book *The Golden Constant*, are used to estimate a Cagan-type demand function that relates the total real value of gold to its expected rate of return. The model assumes that gold production and a latent scale variable (income or consumption) are jointly exogenous and that the data are measured with error. The data reject the model: the estimates imply that the real value of gold varies a great deal relative to the expected return and depends on it *negatively*, rather than positively.

Keywords: *gold, Cagan demand function, estimation*

JEL Codes: *E41*

Citation: Deviatov, A. (2019). Estimating a Cagan-type Demand Function for Gold: 1561–1913. *Russian Journal of Money and Finance*, 78(3), pp. 122–136.
doi: 10.31477/rjmf.201903.122

1. Introduction

A Cagan demand function for money (Cagan, 1956) posits that its real value depends positively on its expected return, as measured by the inverse of the inflation rate. Here, we attempt to estimate such a function for gold using data for 1561–1913. The initial date is determined by the available data, while the terminal date is chosen in accordance with our view that World War I marks the beginning of a period of turmoil regarding the role of gold in private portfolios and in the world's monetary system. The data we use are from *The Golden Constant* by Roy Jastram (1977):¹ a time series on the price of gold in terms of consumption (Jastram, 1977, pp. 34–37, Table 3: The Index of Purchasing Power of Gold, England, 1560–1976), which we interpret as the inverse of the price level, and a time series on gold production (Jastram, 1977, pp. 221–225,

¹ I am grateful to Neil Wallace for bringing my attention to the data. All errors are mine.

Appendix C: The Index of World Production of Gold, 1493–1972).² Because the data describe world production of gold, we view our demand function as applying to the world as a whole.

The main challenge, of course, is modeling the expected return. Our approach is simple – perhaps too simple. Measurement error aside, we assume that there are two exogenous random processes. One is a process for gold production and the other is a latent (unobserved by us) process for a scale variable such as total income or consumption. We assume that the processes and the current state are known to the agents and that an expected return is formed based on that information and on the true demand function.³

It turns out that our model is rejected by the data in the following senses: our estimates imply that the real value of gold varies a great deal relative to the expected return and depends on it *negatively*, rather than positively. We hope that our effort will inspire others to consider alternative models of these time series.

2. The data

The gold production data are shown in Figure 1. The data prior to 1851 are averages: 20-year averages before 1811 and 10-year averages during 1811–1850. As described below, we use those averages and part of the model to interpolate the missing annual data on gold production.

Figure 1. The index of world production of gold, 1492–1972

In Figure 2 we plot the logarithm of the purchasing power of gold; that is, the logarithm of the price of gold. Notice that, in contrast to gold production, there does

² Jastram describes how he came to produce the volume from which we take the data: “My interest in gold began in 1936 for a pragmatic reason. As the most junior member of the Stanford University Department of Economics, I was chosen to volunteer to do some research commissioned by Mr. C.O.G. Miller, an industrialist and gentleman scholar” (Jastram, 1977, p. vii).

³ Our model is therefore part of a large literature on the “rational expectations” specifications of the Cagan model. See, for example, Christiano (1987), Engsted (1993), Goodfriend (1982), Salemi and Sargent (1979), Sargent (1977) and Taylor (1991).

not seem to be a trend in the purchasing power of gold. As Jastram says, the title of his book refers to the absence of such a trend.

Figure 2. Logarithm of the purchasing power of gold, 1561–1913

3. The model

Previous attempts to estimate Cagan demand functions use data from observed episodes of extremely high inflation (mainly post-WWI and postWWII European hyperinflations). That choice of data is largely motivated by Cagan's view that during periods of high inflation, the rate of return on money becomes the primary determinant of money demand. This is a convenient view, because it permits simple ordinary least squares (OLS) estimation of the rate-of-return elasticity, even if data on other money demand determinants such as real income or rates of return on other assets are unavailable. However, as we now explain, we cannot adopt such a simple estimation procedure.

Our Cagan demand function, which is essentially the same as that of Goodfriend (1982) and others, is

$$\ln G_t + \ln p_t = \ln y_t + \alpha[E_t \ln p_{t+1} - \ln p_t], \quad (1)$$

where G_t is the stock of gold, p_t is the price of gold, y_t is a latent scale variable (income), α is the return elasticity, and E_t denotes mathematical expectation conditional on information through date t .

The rational expectations hypothesis implies that the future price p_{t+1} can be represented as

$$\ln p_{t+1} = E_t \ln p_{t+1} + v_{t+1}, \quad (2)$$

where v_{t+1} is uncorrelated with contemporaneous and past gold quantities, prices, and income.

Substitution of (2) into (1) yields

$$\ln p_{t+1} = \frac{1}{\alpha} \ln G_t + \frac{1 + \alpha}{\alpha} \ln p_t + u_{t+1}, \quad (3)$$

where

$$u_{t+1} \equiv -\frac{1}{\alpha} \ln y_t + v_{t+1}. \quad (4)$$

It follows that (3) can be consistently estimated using OLS, provided that $\ln p_t$ is uncorrelated with u_{t+1} . Given the processes for G_t and y_t , the fundamental solution to (1) implies that

$$\ln p_t = \frac{1}{1+\alpha} \sum_{i=0}^{\infty} \left(\frac{\alpha}{1+\alpha} \right)^i E_t[\ln y_{t+i} - \ln G_{t+i}]. \quad (5)$$

OLS estimation of (3) is therefore justified if G_t is an exogenous process and if y_t is nonrandom – for example, if y_t is constant. Both assumptions may seem reasonable for data from post-WWI hyperinflations in Austria, Germany, Greece, Hungary, Poland, and Russia, but the second assumption is not justified in the context of our long time series.

In Jastram's data, the stock of gold exhibits an increasing trend, but the price of gold does not. The obvious way to reconcile the two is to account for the trend in y_t . In particular, we assume that p_t is stationary.⁴ Then, because the processes for the gold stock G_t and the scale variable y_t appear to be nonstationary, equation (5) implies that there should be a cointegrating relationship between G_t and y_t .

We assume that G_t is an exogenous process of the form

$$\ln G_t - \ln \gamma_t = \ln \mu_t + \ln G_{t-1} - \ln \gamma_{t-1}, \quad (6)$$

where γ_t (the transitory shock) and μ_t (the permanent shock) are realizations of independent finite-state Markov processes with positive supports. Our model for the cointegrating relationship between G_t and y_t is

$$\ln G_t - \ln \gamma_t - \ln y_t = \ln \zeta_t, \quad (7)$$

where ζ_t is a realization of an independent finite-state Markov process with a positive support. That is,

Random variable	Support	Transition probability
γ_t	$(\gamma^1, \gamma^2, \dots, \gamma^{M_\gamma})$	π_{ij}^γ
μ_t	$(\mu^1, \mu^2, \dots, \mu^{M_\mu})$	π_{ij}^μ
ζ_t	$(\zeta^1, \zeta^2, \dots, \zeta^{M_\zeta})$	π_{ij}^ζ

(8)

where the supports for μ_t , γ_t and ζ_t are equally spaced and the support for γ_t is also symmetric around unity. We assume that the agents know these processes and observe $(\gamma_t, \mu_t, \zeta_t)$ before deciding on their gold holdings.

⁴ The augmented Dickey-Fuller test performed for the purchasing power of gold logarithm over the 1561–1913 sample rejects the null of the unit root at 0.1% significance level.

The process for G_t and the Markov property of γ_t imply that the effect of γ_t on G_t dies out with time (provided that γ_t is asymptotically stable). In contrast, the effect of μ_t on G_t is permanent. This is what motivates our use of the labels 'transitory' and 'permanent' for γ_t and μ_t , respectively. An alternative interpretation is that the process for gold stock G_t has a stochastic trend μ_t and a stochastic cyclical component γ_t around that trend.

Substitution of (6) and (7) in (1) leads to an equation of the same form as (3); namely,

$$\ln p_{t+1} = \frac{1 + \alpha}{\alpha} \ln p_t + u'_{t+1}, \quad (9)$$

where

$$u'_{t+1} = \frac{1}{\alpha} (\ln \gamma_t + \ln \varsigma_t) + v_{t+1}. \quad (10)$$

Therefore, OLS estimation of (9) is justified if γ_t and ς_t are not random, an assumption that is too extreme in our context. We therefore proceed as follows.

Using (6)–(7), we can rewrite (1) as

$$\ln p_t = -\frac{1}{1 + \alpha} (\ln \varsigma_t + \ln \gamma_t) + \frac{\alpha}{1 + \alpha} E_t \ln p_{t+1}. \quad (11)$$

Then, our analogue of (3) is the equilibrium condition in the following definition.

Definition 1. A function $p : R^{M_\gamma} \times R^{M_\varsigma} \rightarrow R_{++}$ (with generic element p_{hl} , which denotes the price when $(\gamma_t, \varsigma_t) = (\gamma^h, \varsigma^l)$) is a stationary equilibrium if it satisfies

$$\ln p_{hl} = -\frac{1}{1 + \alpha} (\ln \varsigma^l + \ln \gamma^h) + \frac{\alpha}{1 + \alpha} \sum_{i=1}^{M_\gamma} \sum_{j=1}^{M_\varsigma} \pi_{hi}^\gamma \pi_{lj}^\varsigma \ln p_{ij} \quad (12)$$

for each $(h, l) \in \{1, 2, \dots, M_\gamma\} \times \{1, 2, \dots, M_\varsigma\}$.

This system of linear equations can be written as $x = a_0 + a_1 \pi x$, where x is an $M_\gamma M_\varsigma \times 1$ vector, $a_1 = \frac{\alpha}{1 + \alpha}$ and π is the Kronecker product of π^γ and π^ς , the transition matrices for γ and ς , respectively. Provided that the matrix $I - a_1 \pi$ is nonsingular, the equilibrium exists and is unique. Obviously, that is the case if $\alpha \geq 0$. In any case, nonsingularity of the matrix $I - a_1 \pi$ is generic.

Our basic approach to estimation follows Cosslett and Lee (1985) and is consistent with our assumption that agents see $(\gamma_t, \mu_t, \varsigma_t)$ before deciding on their gold holdings. That assumption forces us to estimate not only the processes for $(\gamma_t, \mu_t, \varsigma_t)$, but also the realizations. We specify magnitudes for M_γ , M_μ and M_ς and assume that the data on the stock of gold and its price are measured with multiplicative errors. Our version of the estimation procedure has two stages.⁵

⁵ See Appendix 2 for a description of how we would jointly estimate the model and for a discussion of why we chose not to do that.

First, we use only the data on the stock of gold to estimate the processes and realizations for (γ_t, μ_t) , the measurement error for the stock of gold production, and the initial stock of gold. Second, taking those as given, we use the price data to estimate the rest of the model: the process and realization for ς_t , the measurement error in p_t , and α – the main parameter of interest.

4. Estimating the process for the stock of gold

Our model of the stock of gold⁶ for $t = 1, 2, \dots, T$ is

$$G_t = G \left(\prod_{i=1}^t \mu_i \right) \gamma_t \exp(\varepsilon_t^G), \quad (13)$$

where G , a parameter to be estimated, is the stock of gold at $t = 0$ and $\varepsilon_t^G \sim N(0, \sigma_G^2)$, the measurement error, is independent and identically distributed (i.i.d.) and is independent of γ_t, μ_t and ς_t .

As noted above, we estimate some features of the Markov processes for γ_t and μ_t and impose others. Because scaling the support of γ_t by one factor and scaling the support of μ_t by its inverse leaves the true process for the stock of gold unaffected, we normalize γ_t by assuming that the support of γ_t is symmetric around unity. We do not estimate the number of elements in the supports. We use $M_\gamma = 3$ and $M_\mu = 5$, which yield a reasonably good fit.⁷ And, finally, for both supports, we assume equally spaced elements; that is, $\gamma^{j+1} = \gamma^j + \Delta_\gamma$ and $\mu^{j+1} = \mu^j + \Delta_\mu$. Given the normalization imposed on the support of γ_t , that leaves us to estimate only Δ_γ for the support of γ_t . For the support of the μ_t process, we estimate the lower endpoint of the support, μ^1 and Δ_μ . To estimate π^γ and π^μ , we use a logit representation for each row (distribution) of the transition matrices, i.e.:

$$\pi_{ij} = \begin{cases} \frac{e^{w_{ij}}}{1 + \sum e^{w_{ij}}} & \text{if } i \neq j, \\ \frac{1}{1 + \sum e^{w_{ij}}} & \text{if } i = j. \end{cases} \quad (14)$$

We impose additional bound constraints on w_{ij} , which ensure that the estimated transition matrices are ergodic. The latter helps us to avoid the inconvenience of having to deal with the multiple stationary distributions implied by π^γ and π^μ during the course of estimation. Thus, $\kappa = (G, \Delta_\gamma, \pi^\gamma, \mu^1, \Delta_\mu, \pi^\mu, \sigma_G)$ are the parameters to be estimated for this part of the model.

⁶ An equivalent alternative is to model gold production.

⁷ We arrived at these choices as follows. We began with $M_\gamma = 3$ and $M_\mu = 1$, which were a poor fit. We further found that increasing M_γ while maintaining $M_\mu = 1$ did not substantially improve the fit. Then we tried $M_\gamma = 3$ and $M_\mu = 3$ and $M_\gamma = 3$ and $M_\mu = 5$. The latter fit substantially better than the former; as we show later, the implied measurement error ε_t^G is less than 1% of the gold stock $\tilde{G}_t$.

Letting $\gamma = \{\gamma_t\}_{t=1}^T$ and $\mu = \{\mu_t\}_{t=1}^T$ denote possible realizations for the γ_t and μ_t processes, the likelihood function is:

$$L_1 = \left[\prod_{t=1}^T \Psi \left[\ln G_t - \ln G - \sum_{i=1}^t \ln \mu_i - \ln \gamma_t ; \sigma_G \right] \right] P(\gamma)P(\mu), \quad (15)$$

where $\Psi(x; \sigma_G)$ is the density of the normal distribution,

$$G_t = G + \sum_{i=1}^t Z_i, \quad (16)$$

where Z_i is the date i data for gold production, and where $P(\gamma)$ and $P(\mu)$ are the probabilities of the sequences γ and μ implied by the transition matrices π^γ and π^μ .

Because the number of γ and μ sequences is large and the state space is discrete, we use a genetic algorithm as part of our procedure for maximizing L_1 . In particular, we proceed as follows.

1. Generate a population of pairs of sequences which determine γ and μ , but only by their order in the respective domains. That is, the sequences corresponding to γ are drawn from the set $\{1, 2, 3\}^T$ and those for μ are drawn from the set $\{1, 2, 3, 4, 5\}^T$.

2. For each pair of sequences in the step 1 population, find the parameters κ that maximize the objective L_1 . (Notice that for given parameters that complete the description of the domains for γ_t and μ_t – $\Delta_\gamma, \mu^1, \Delta_\mu$ – a pair of sequences in the step 1 population determine γ and μ . Those realizations and the remaining parameters in κ – $G, \pi^\gamma, \pi^\mu, \sigma_G$ – determine a magnitude for L_1 . Given the logit representation of the transition matrices π^γ and π^μ , the maximization problem is a standard constrained maximization problem with bound constraints. However, the objective is not concave, so there may be many local maxima.)

3. Apply standard genetic operators (such as simple crossover and mutation) to amend the population of step 1 sequences.

4. Repeat step 2 until the best pair γ, μ is found.

5. Compute confidence intervals for the estimated parameters, $\hat{G}, \hat{\Delta}_\gamma, \hat{\mu}^1, \hat{\Delta}_\mu, \hat{\sigma}_G$, using a Monte Carlo procedure.

Because our concern about the post-1913 data does not apply to gold production, we use the entire data set, 1561 to 1972, for gold production. However, prior to 1851, the gold production data reported in Jastram (1977) (and in the original source) are averages: 20-year averages prior to 1811 and 10-year averages during 1811–1850. We therefore use an interpolation procedure which is consistent with the maximum likelihood estimation procedure that we carry out. Given a vector κ of parameters and a γ and μ sequences, the likelihood function (15) is concave with respect to the gold production data Z_i . Consequently, if some annual data are missing, then maximum likelihood

interpolation of the missing data implies that the data should be interpolated in order to keep the terms,

$$\ln G_t - \ln G - \sum_{i=0}^t \ln \mu_i - \ln \gamma_t,$$

in the likelihood function (15) constant over all sample periods for which we must interpolate. This requirement implies an interpolation procedure that makes the interpolated terms dependent on the γ and μ realizations and on the parameters. Each evaluation of the likelihood function therefore has its own interpolation of the missing data. The details appear in Appendix 1.

Figure 3. Gold production: interpolated up to 1850, actual after 1850

Figure 3 contains the interpolated production series (before 1851) and the actual series (after 1851). Given the sparse parametrization (three points in the support of γ_t and five points in the support of μ_t) a change in the state gives rise to large production changes. Moreover, because our interpolation procedure is sensitive to both permanent and transitory components (see Appendix 1, expression (22)), a change in the state creates spikes in the interpolated gold production series.

Some of the spikes, in particular the negative production (loss) of gold, implied by interpolation seem implausible. However, these spikes are consistent with maximization of the likelihood function. We could have eliminated the spikes – for example, by using the averages reported by Jastram (1977) or some other smoothing device. However, given that the data rejects the model even allowing for the spikes, eliminating them would only reinforce our negative result.

The estimated parameters for the gold process are given in Table 1. The estimated initial stock of gold (in 1560) equals 28.5 times gold production in 1930. Measured as percentages, the support of transitory shock is approximately $\{-2.5, 0, 2.5\}$. The support of the permanent component is roughly $\{-0.1, 0.3, 0.7, 1.1, 1.5\}$, so that gold stock growth rates range from -0.1% to 1.5% per year. Given our model for

the gold stock (13), ε_t^G measures the percentage deviation of the gold stock implied by the data from that implied by the model.

Table 1. Estimated parameters for the gold process

Parameter	Estimate	5% confidence interval	
		Lower end	Upper end
G	28.5639	28.5085	29.0073
Δ_r	0.024643	0.024313	0.024835
μ^i	0.999139	0.999137	0.999142
Δ_μ	0.004043	0.004014	0.004048
σ_G	0.002328	0.002151	0.002471
L_1	2048.33		

The estimate $\hat{\sigma}_G = 0.002328$ means that the standard deviation of the discrepancy is about one quarter of 1%. In fact, the difference between gold stock data and the model is less than 1% for almost the entire sample, and less than 1.3% for the entire sample. However, the errors measured in terms of gold production are substantially larger. Relative to gold production, errors of 30–50% are not uncommon and for some years (in the earlier part of sample) the errors exceed 100% of (implied) gold production. We show the plot of $\hat{\varepsilon}_t^G$ in Figure 4. Notice that large spikes in interpolated production data before 1851 are consistent with small errors $\hat{\varepsilon}_t^G$ during that period. In addition, because the likelihood function is concave with respect to the interpolated data, the implied errors are constant over each 20-year (10-year) interpolation subperiod before 1851.

Figure 4. Estimated gold production measurement error $\hat{\varepsilon}_t^G$, 1561–1972

We use the following Monte Carlo procedure to obtain the confidence intervals in Table 1. First, we use $\hat{\sigma}_G$ to generate a sequence of simulated errors $\hat{\varepsilon}_t^G \sim N(0, \hat{\sigma}_G^2)$. After that, we use $\hat{\gamma}$, $\hat{\mu}$ and $\hat{G}$ to compute a simulated gold stock series; namely,

$$\tilde{G}_t = \hat{G} \left(\prod_{i=1}^t \hat{\mu}_i \right) \hat{\gamma}_t \exp(\tilde{\varepsilon}_t^G). \quad (18)$$

Then, we treat $\tilde{G}_t$ as 'data' and obtain new estimates of γ , μ , G and σ_G . In doing this we keep constant the order of states in the estimated sequences of states for γ and μ by order described in step 1 above.

A more complete procedure would reestimate the sequences of states for γ and μ by order. Such a procedure would yield larger confidence intervals. However, our estimates of the ratios $\hat{\Delta}_\gamma / \hat{\sigma}_G \approx 11$ and $\hat{\Delta}_\mu / \hat{\sigma}_G \approx 1.7$ (see Table 1) imply that only a small part of the simulated data would differ from the actual data enough to produce changes in the sequences for γ and μ by order. Given (18), the ratios $\hat{\Delta}_\gamma / \hat{\sigma}_G$ and $\hat{\Delta}_\mu / \hat{\sigma}_G$ are effectively the t-ratios for the null hypothesis that the data distinguish between the two adjoining points in the supports of γ and μ , respectively. The ratio $\hat{\Delta}_\mu / \hat{\sigma}_G$, approximately equal to 1.7, implies that under the more complete procedure only about 10% of the sequence for μ by order would differ from the estimated sequence for μ .

Thus, it did not seem worthwhile to undertake the more complete procedure, which would require a large amount of additional computational time.⁸

The 5% confidence intervals are reported in the two rightmost columns of Table 1. Notice that the estimates $\hat{G}$, $\hat{\Delta}_\gamma$, $\hat{\mu}^1$, $\hat{\Delta}_\mu$, $\hat{\sigma}_G$ are not midpoints of those intervals. Given our procedure, the true gold stock series G_t can be viewed as implied by the sequence of measurement errors, $\tilde{\varepsilon}_t^G$, which is drawn from the normal distribution with zero mean and variance $\hat{\sigma}_G^2$. There is no prior reason to expect that that drawing yields the estimates $\hat{G}$, $\hat{\Delta}_\gamma$, $\hat{\mu}^1$, $\hat{\Delta}_\mu$ and $\hat{\sigma}_G$ which fall exactly in the middle of their 5% confidence intervals. And, as one can see, that is not the case.

The estimated transition matrices are:

$$\hat{\pi}^\gamma = \begin{bmatrix} 0.9749 & 0.0251 & 0 \\ 0.1140 & 0.7170 & 0.1690 \\ 0 & 0.0223 & 0.9777 \end{bmatrix}$$

and

$$\hat{\pi}^\mu = \begin{bmatrix} 0.9264 & 0.0736 & 0 & 0 & 0 \\ 0.1074 & 0.8406 & 0.0520 & 0 & 0 \\ 0.0273 & 0.1131 & 0.7612 & 0.0984 & 0 \\ 0 & 0 & 0.1047 & 0.7398 & 0.1555 \\ 0 & 0 & 0 & 0.0814 & 0.9186 \end{bmatrix}.$$

The plots of the estimated sequences of $\hat{\gamma}$ and $\hat{\mu}$ are given in Figures 5 and 6. The estimates are consistent with high persistence in the gold stock process, so that the most likely state next year is the current state. However, when the

⁸ A limited investigation of the consequences for confidence intervals of varying the sequences for γ and μ is reported in Appendix 3.

transition occurs, the state does not change much; almost all changes are changes to an adjoining state. Furthermore, consistent with a considerable increase in gold production during the course of history (see Figure 1), the estimated sequence $\hat{\mu}$ starts out low and transits to the upper end of the support.

Figure 5. Estimated sequence for the permanent component of the gold stock, $\hat{\mu}$, 1561–1972

Figure 6. Estimated sequence for the temporary component of the gold stock, $\hat{\gamma}$, 1561–1972

5. Estimating the rest of the model

Our model of the gold price data is

$$p_t = p(\zeta_t, \gamma_t) \exp(\varepsilon_t^p), \quad (19)$$

where $p(\cdot, \cdot)$ is the equilibrium price function and where $\varepsilon_t^p \sim N(0, \sigma_p^2)$ is i.i.d. and is independent of γ_t, μ_t, ζ_t , and of ε_t^G . As above, we interpret ε_t^p as measurement error.

We take $M_\zeta = 3$ and assume an equally spaced support for ζ_t , ($\zeta^{j+1} = \zeta^j + \Delta_\zeta$). We also treat as known the estimated stage one parameters and the realizations for the γ_t and μ_t processes (denoted $\hat{\kappa}$, $\hat{\gamma}$ and $\hat{\mu}$). Our choice of equally spaced support for ζ_t and a relatively small M_ζ (as well as adoption of the two-stage estimation procedure in general) is dictated by the size of the computing problem. By using the stage one estimates, we can increase the number of price states from $M_\zeta = 3$ to $M_\gamma M_\zeta = 9$ without increasing the computational burden. The remaining stage two parameters are:

$$\rho = (\alpha, \zeta^1, \Delta_\zeta, \pi^\zeta, \sigma_p). \quad (20)$$

Letting $\varsigma = \{\zeta_t\}_{t=1}^T$ denote a realization for the ζ_t process, the likelihood function is:

$$L_2 = \left[\prod_{t=1}^T \Psi[\ln p_t - \ln(p(\zeta_t, \hat{\gamma}; \rho, \hat{\kappa})); \sigma_p] \right] P(\varsigma), \quad (21)$$

where $\Psi(x; \sigma_p)$ is the density of the normal distribution, $p(\cdot, \cdot; \rho, \hat{\kappa})$ is the solution for the equilibrium price function corresponding to the parameters $(\hat{\kappa}, \rho)$, and $P(\varsigma)$ is the probability of the sequence ς implied by the transition probabilities π^ζ .

Our procedure is similar to the one we used when estimating the production process.

1. Generate a population of sequences which determine ς , but only by order in the respective domain. That is, the sequences corresponding to ς are drawn from the set $\{1, 2, 3\}^T$.

2. For each sequence in the step 1 population, find the parameters ρ that maximize the objective L_2 . This involves the following steps. Pick ρ . This ρ and a stage one sequence imply a corresponding sequence ς . Solve for the equilibrium prices. Together, these imply a magnitude for L_2 . Search over values of ρ . Given a logit representation for the transition probabilities in π^ζ , search over ρ is a standard constrained maximization problem with bound constraints.

3. Apply standard genetic operators (such as simple crossover and mutation) to amend the population of step 1 sequences.

4. Repeat step 2 until the best ς is found.

5. Compute confidence intervals using a Monte Carlo procedure.

As noted above, here we use the sample period from 1561 to 1913. The estimates $\hat{\alpha}$, $\hat{\zeta}^1$, $\hat{\Delta}_\zeta$, $\hat{\sigma}_p$ are given in Table 2 and the estimation transition matrix is

$$\hat{\pi}^\zeta = \begin{bmatrix} 0.9922 & 0.0040 & 0.0038 \\ 0.0032 & 0.9930 & 0.0038 \\ 0.0038 & 0.0037 & 0.9925 \end{bmatrix}.$$

The confidence intervals in Table 2 are computed using the same procedure we used above. In particular, we kept the order sequence corresponding to $\hat{\zeta}$ fixed

during Monte Carlo simulation of the confidence intervals. Because the change in the equilibrium price function associated with a change in the discrete order sequence for ζ is about 21%, i.e. roughly $2\hat{\sigma}_p$, only a small part of the simulated price data would differ from the actual series enough to produce changes in the discrete sequence for ζ if we let the sequence by order change in the course of simulations. Thus, as above, it did not seem worthwhile to undertake the more complete procedure.⁹

Table 2. Estimated parameters for the rest of the model

Parameter	Estimate	5% confidence interval	
		Lower end	Upper end
α	-90.791	-90.907	-88.881
ζ^1	$8.9507 \cdot 10^{-3}$	$8.8387 \cdot 10^{-3}$	$9.0606 \cdot 10^{-3}$
Δ_ζ	$1.7431 \cdot 10^{-7}$	$6.3120 \cdot 10^{-8}$	$2.9184 \cdot 10^{-7}$
σ_p	0.11244	0.10323	0.12014
L	557.68		

The estimated α is large, significant, and has the wrong sign. Interestingly, this holds true if we reestimate the price process for various subsamples within the period for which we have data. We report the estimates of the price process parameters for the periods of the Napoleonic Wars, pre- and post-World-War gold standards, and Bretton Woods monetary system in Appendix 4. Because the number of states M_ζ is relatively small, we have not obtained a close fit of the price data; the standard deviation $\hat{\sigma}_p$ of the error is quite large (about 11%). We show the actual and fitted prices in Figure 7.

Figure 7. Purchasing power of gold: actual and fitted, 1561–1913

⁹ Again, see Appendix 3 for further details.

The estimated sequence $\hat{\xi}$ is persistent, so that the implied price p_t is also persistent. Persistence means that the most likely price tomorrow is the current price. In other words, high persistence implies that the expected rate of return on gold is often zero.

We present the plot of the net expected rate of return on gold in Figure 8, along with the actual return from the data, the first difference of the (logarithm of the) series in Figure 2. At no time does the net expected rate of return implied by the model exceed 0.3% in absolute value. Given the erratic pattern of actual returns, it is, perhaps, not surprising that our model gives rise to little variation in the expected return.

Figure 8. Actual and expected rates of return on gold, 1562–1913

Figure 9. Value of gold stock and realized rate of return on gold. Actual data, 1562–1913

Figure 10. Value of gold stock and expected rate of return on gold. Fitted data, 1561–1913

In Figures 9 and 10, we present two scatter plots. Figure 9 is the data on the value of the stock of gold (except that we use the interpolated stock [see Figure 3])

for the gold stock prior to 1850) and the return on gold. Figure 10 is the analogue using the model's predictions. It is evident from Figure 10 that the model implies that the value of the gold stock is highly sensitive to the rate of return, a sensitivity which happens to be of the wrong sign. However, even if it were the right sign, its magnitude would be implausible.

6. Concluding remarks

Cagan, of course, fitted his demand function to periods of hyperinflation, when there were large 'predictable' variations in the return on money – variations which, he argued, would swamp any changes in a scale variable such as total income or consumption, and any changes in the yields on alternative assets such as the real return on capital. We cannot make that argument. Instead, we use his approach mainly because we do not have access to data on a scale variable or on the yields of alternative assets.

Appendices are available at
[http://rjmf.econs.online/en;](http://rjmf.econs.online/en;dx.doi.org/10.31477/rjmf.201903.122)
dx.doi.org/10.31477/rjmf.201903.122

References

- Cagan, P.** (1956). The Monetary Dynamics of Hyperinflation. In: M. Friedman, ed., *Studies in the Quantity Theory of Money*. Chicago: University of Chicago Press.
- Christiano, L.** (1987). Cagan's Model of Hyperinflation Under Rational Expectations. *International Economic Review*, 28(1), pp. 33–49.
- Cosslett, S. and Lee, L.** (1985). Serial Correlation in Latent Discrete Variable Models. *Journal of Econometrics*, 27(1), pp. 79–97.
- Engsted, T.** (1993). Cointegration and Cagan's Model of Hyperinflation Under Rational Expectations. *Journal of Money, Credit and Banking*, 25(3), pp. 350–360.
- Goodfriend, M.** (1982). An Alternative Method of Estimating the Cagan Money Demand Function in Hyperinflation Under Rational Expectations. *Journal of Monetary Economics*, 9(1), pp. 43–57.
- Jastram, R. W.** (1977). *The Golden Constant*. New York: John Wiley and Sons.
- Salemi, M. and Sargent, T. J.** (1979). The Demand for Money During Hyperinflation Under Rational Expectations: II. *International Economic Review*, 20(3), pp. 741–758.
- Sargent, T. J.** (1977). The Demand for Money During Hyperinflation Under Rational Expectations: I. *International Economic Review*, 18(1), pp. 59–82.
- Taylor, M. P.** (1991) The Hyperinflation Model of Money Demand Revisited. *Journal of Money, Credit and Banking*, 23(3), pp. 327–51.