

**Centre for
Economic
and Financial
Research
at
New Economic
School**



August 2006

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

Alexei Deviatov,
Neil Wallace

Working Paper No 80

CEFIR / NES Working Paper series

Estimating a Cagan-type demand function for gold: 1561-1913

Alexei Deviatov*

47 Nahimovskii Ave., # 1721
New Economic School
Moscow, 117418
Russia
E-mail: deviatov@nes.ru

Neil Wallace

612 Kern Graduate bl.,
Pennsylvania State University
University Park, PA 16802
U.S.A.
E-mail: neilw@psu.edu

August 3, 2006

Abstract

Long times series on production of gold 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 real total 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 *negatively*, rather than positively, on the expected return.

Key words: gold, Cagan demand function, estimation

JEL Classification #: E41.

*Corresponding author.

1 Introduction

A Cagan demand function for money posits that its real value depends positively on its expected return. 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 accord 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 [6]: a time series on the price of gold in terms of consumption (Table 3, The Index of Purchasing Power of Gold: England 1560-1976, pages 34-37) and a time series on gold production (Appendix C, The index of World Production of Gold, 1493-1972, pages 221-225).¹ Because the data describe world production of gold, we view our demand function to be one for 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

¹Jastram 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 ([6], page vii).”

variable like total income or consumption. We assume that the processes and the current state are known and that an expected return is formed based on that information and on the true demand function.²

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 *negatively*, rather than positively, on the expected return. 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.

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 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.

²Therefore, our model is part of a large literature on “rational expectations” specifications of the Cagan model. See, for example, [2], [4], [5], [7], [8], and [9].

3 The model

We treat the data on gold production and the price of gold as measured with error. The following is our model for the true values. The demand function 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, α is a constant, and E_t denotes mathematical expectation conditional on information up through date t .

As noted above, there are two exogenous random processes, the date- t realization of which is denoted (G_t, Y_t) . We assume that the process for G_t is subject to temporary and permanent shocks, according to

$$\ln G_t = \ln(G_t^* \gamma_t) \text{ and } \ln G_t^* = \ln(G_{t-1}^* \mu_t), \quad (2)$$

where γ_t (the temporary shock) and μ_t (the permanent shock) are realizations of independent finite-state Markov processes with positive supports and where the supports for μ_t and for γ_t are equally spaced and where that for γ_t is symmetric around unity. We choose a process for Y_t to be consistent with a stationary process for p_t . That, in turn, requires that the permanent

components of Y_t and G_t be related. We assume that

$$\ln Y_t = \ln(Y_t^* \eta_t) \text{ and } \ln Y_t^* = \ln(G_t^* / \theta_t), \quad (3)$$

where η_t and θ_t are finite-state Markov processes. (The second of these is the cointegrating relationship between Y_t^* and G_t^* that permits there to be a stationary equilibrium.) It follows that

$$\ln Y_t = \ln G_t^* + \ln \eta_t - \ln \theta_t \equiv \ln G_t^* - \ln \varsigma_t, \quad (4)$$

where $\varsigma_t \equiv (\theta_t / \eta_t)$ is a realization of an independent finite-state Markov processes with a positive support that is equally spaced.

Using (2)-(4), we can rewrite (1) as

$$\ln p_t = -\frac{1}{1+\alpha} (\ln \varsigma_t + \ln \gamma_t) + \frac{\alpha}{1+\alpha} E_t \ln p_{t+1}. \quad (5)$$

To give a definition of equilibrium, we need additional notation for the Markov processes; namely,

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
ς_t	$(\varsigma^1, \varsigma^2, \dots, \varsigma^{M_\varsigma})$	π_{ij}^ς

Then we have the following definition of a stationary equilibrium.

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

$$\ln p_{hl} = -\frac{1}{1+\alpha}(\ln \varsigma^l + \ln \gamma^h) + \frac{\alpha}{1+\alpha} \sum_{i=1}^{M_\gamma} \sum_{j=1}^{M_\varsigma} \pi_{hi}^\gamma \pi_{lj}^\varsigma \ln p_{ij} \quad (6)$$

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, an 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 S. and Lee [3]. We specify magnitudes for M_γ , M_μ , and M_ς and assume that the data are measured with multiplicative errors. Our version of the estimation procedure has two stages.³ First, we estimate the process and realization for gold production using only the data on gold production. This gives estimates of the parameters for the γ_t and μ_t processes and for the measurement error for gold production, for the initial stock of gold, and for the realizations of the

³See appendix 2 for a description of joint estimation of the entire model.

γ_t and μ_t processes. Second, taking those as given, we estimate the rest of the model: the parameters for the ς_t process and for the measurement error in p_t , the realization for the ς_t process, the initial value of Y_t , and α .

4 Estimating the stock of gold

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

$$G_t = G \left(\prod_{i=1}^t \mu_i \right) \gamma_t \exp(\varepsilon_t^G), \quad (7)$$

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 i.i.d. and is independent of γ_t , μ_t , and ς_t .

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 μ_t by its inverse leaves the true process for 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.⁴ And, finally, for both

supports, we assume equally spaced elements; that is, $\gamma^{j+1} = \gamma^j + \Delta_\gamma$ and

⁴We arrived at these choices as follows. We began with $M_\gamma = 3$ and $M_\mu = 1$, which fit poorly. And we 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 .

$\mu^{j+1} = \mu^j + \Delta_\mu$. 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.:

$$\pi_{ij} = \begin{cases} \frac{e^{w_{ij}}}{1 + \sum e^{w_{ij}}} & \text{if } i \neq j \\ \frac{1}{1 + \sum e^{w_{ij}}} & \text{if } i = j \end{cases}. \quad (8)$$

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 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.

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

$$L_1 = \left[\prod_{t=1}^T \Psi[\ln G_t - \ln G - \sum_{i=1}^t \ln \mu_i - \ln \gamma_t; \sigma_G] \right] P(\gamma)P(\mu), \quad (9)$$

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

$$G_t = G + \sum_{i=1}^t Z_i, \quad (10)$$

where Z_i is the date i datum 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.

Step 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 $\{1, 2, 3, 4, 5\}^T$.

Step 2. For each pair of sequences in the population of step 1, find the parameters κ that maximize the objective L_1 . (Notice that for given parameters that completes 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 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, $\widehat{G}, \widehat{\Delta}_\gamma, \widehat{\mu}^1, \widehat{\Delta}_\mu, \widehat{\sigma}_G$, 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 [6] (and in the original source) are averages: twenty-year averages prior to 1811 and ten-year averages during 1811-1850. Therefore, we use an interpolation procedure, one which is consistent with the maximum likelihood estimation procedure that we carry out. Given a vector κ of parameters and a γ and μ , the likelihood function (9) 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,

$$\ln G_t - \ln G - \sum_{\iota=0}^t \ln \mu_\iota - \ln \gamma_t, \quad (11)$$

in the likelihood function (9) constant over all sample periods for which we must interpolate. This requirement implies an interpolation procedure, one that makes the interpolated terms dependent on the γ and μ realizations and on the parameters. Therefore, each evaluation of the likelihood function has

its own interpolation of the missing data. The details appear in appendix 1.

Figure 3 contains the interpolated production series (before 1851) and the actual series (after 1851). Given the sparse parametrization (3 points in the support of γ_t and 5 points in the support of μ_t), a change in the state gives rise to large production changes. And because our interpolation procedure is sensitive to both permanent and temporary components (see expression 16), a change in the state creates spikes in the interpolated gold production series.

The estimated parameters for the gold process are given in Table 1. The estimated initial stock of gold (in year 1560) equals 28.5 times the year 1930 gold production. Measured as percentages, the support of the temporary 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 (7), ε_t^G measures the percentage deviation of the gold stock implied by the data from that implied by the model. The estimate $\hat{\sigma}_G = 0.002328$ means that the standard deviation of the discrepancy is about one quarter of one percent. In fact, for almost all of the sample, the difference between the data and the model is less than 1% and for all of the sample this difference is less

than 1.3%. We show the plot of $\widehat{\varepsilon}_t^G$ in Figure 4. Notice that large spikes in interpolated production data before 1851 are consistent with small errors $\widehat{\varepsilon}_t^G$ during that period.

We use the following Monte Carlo procedure to obtain the confidence intervals in Table 1. First, we use $\widehat{\sigma}_G$ to generate a sequence of simulated errors $\widetilde{\varepsilon}_t^G$ from the normal distribution with zero mean and variance $\widehat{\sigma}_G^2$. After that, we use $\widehat{\gamma}$, $\widehat{\mu}$, and \widehat{G} to compute a simulated gold stock series; namely,

$$\widetilde{G}_t = \widehat{G} \left(\prod_{i=1}^t \widehat{\mu}_i \right) \widehat{\gamma}_t \exp(\widetilde{\varepsilon}_t^G). \quad (12)$$

Then, we treat \widetilde{G}_t as “data” and obtain new estimates of γ , μ , G , and σ_G . In doing this we keep constant 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 $\widehat{\Delta}_\gamma/\widehat{\sigma}_G \approx 11$ and $\widehat{\Delta}_\mu/\widehat{\sigma}_G \approx 1.7$ (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 (12), the ratios $\widehat{\Delta}_\gamma/\widehat{\sigma}_G$ and $\widehat{\Delta}_\mu/\widehat{\sigma}_G$ 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 $\widehat{\Delta}_\mu/\widehat{\sigma}_G$, 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.⁵

The 5% confidence intervals are reported in the two rightmost columns in 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, $\hat{\varepsilon}_t^G$, 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 $\hat{\gamma}$ and $\hat{\mu}$ are given in Figures 5 and 6.

⁵See appendix 3 for further details.

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 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.

5 Estimating the rest of model

Our model of the gold price data is

$$p_t = p(\varsigma_t, \gamma_t) \exp(\varepsilon_t^p), \quad (13)$$

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 , ς_t , and of ε_t^G . As above, we interpret ε_t^p as measurement error. We take $M_\varsigma = 3$ and assume an equally-spaced support for ς_t , ($\varsigma^{j+1} = \varsigma^j + \Delta_\varsigma$). We also treat as known the estimated stage-1 parameters and the realizations for the γ_t and μ_t processes. These are denoted $\hat{\kappa}$, $\hat{\gamma}$ and $\hat{\mu}$. The remaining parameters are:

$$\rho = (\alpha, \varsigma^1, \Delta_\varsigma, \pi^\varsigma, \sigma_p). \quad (14)$$

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

$$L_2 = \left[\prod_{t=1}^T \Psi[\ln p_t - \ln(p(\varsigma_t, \hat{\gamma}; \rho, \hat{\kappa})); \sigma_p] \right] P(\varsigma), \quad (15)$$

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^ς .

Our procedure is similar to what we did 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 population of step 1, find the parameters ρ that maximize the objective L_2 . This involves the following steps. Pick ρ . This ρ and a step-1 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^ς , search over ρ is a standard constrained maximization problem with bound

constraints.

3. Apply standard genetic operators to amend the population of step 1 sequences.
4. Repeat step 2 until the best ζ is found.
5. Compute confidence intervals using a Monte Carlo procedure.

As noted above, we here 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 estimated 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 the 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.⁶

The estimated α is large, significant, and has the wrong sign. Because

⁶Again, see appendix 3 for further details.

the number of states M_c 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 the fitted price in Figure 7.

The estimated sequence $\hat{\varsigma}$ is persistent, so that the implied price p_t is persistent as well. 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.

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 for the gold stock the interpolated stock [see figure 3] prior to 1850) and the return on gold. Figure 10 is the analogue using the predictions of the model. It is evident from Figure 10 that the model implies huge sensitivity of the value of the gold stock to the rate of return, 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 like total income or consumption and any changes in the yields on alternative assets like the real return on capital. We cannot make that argument. Instead, we use his approach mainly because we do not have data on a scale variable or on the yields of alternative assets.

7 Appendix 1. Interpolation of gold production

Let $\{t_n + 1, t_n + 2, \dots, t_{n+1}\}$ be a list of years in the sample period for which only aggregate gold production datum \widehat{Z}_{n+1} is available. Constancy of all terms (11) implies that

$$\ln G_{t+1} - \ln G_t = \ln \mu_{t+1} + \ln \gamma_{t+1} - \ln \gamma_t,$$

for $t = t_n + 1, \dots, t_{n+1}$. Taking (10) into account, the interpolated sequence of gold production data satisfies

$$\tilde{Z}_{t+1} = (\Upsilon_{t+1} - 1) \left(G_{t_n} + \sum_{\iota=0}^t \tilde{Z}_{\iota} \right),$$

where

$$\Upsilon_{t+1} \equiv \mu_{t+1} \frac{\gamma_{t+1}}{\gamma_t}.$$

Because the sequence of interpolated data \tilde{Z}_t must sum to \hat{Z}_{n+1} , we have

$$G_{t_n} + \hat{Z}_{n+1} = \left(G_{t_n} + \tilde{Z}_{t_{n+1}} \right) \prod_{i=2}^{t_{n+1}-t_n} \Upsilon_{t_n+i}.$$

It follows that

$$\tilde{Z}_{t_{n+1}} = \frac{G_{t_n} + \hat{Z}_{n+1}}{\prod_{i=2}^{t_{n+1}-t_n} \mu_{t_n+i} \frac{\gamma_{t_{n+1}}}{\gamma_{t_n+1}}} - G_{t_n},$$

and

$$\tilde{Z}_{t_n+s} = \left(G_{t_n} + \hat{Z}_{n+1} \right) \frac{\mu_{t_n+s} \frac{\gamma_{t_n+s}}{\gamma_{t_n+s-1}} - 1}{\prod_{i=s+1}^{t_{n+1}-t_n} \mu_{t_n+i} \frac{\gamma_{t_{n+1}}}{\gamma_{t_n+s-1}}}, \quad (16)$$

for all $s = 2, 3, \dots, t_{n+1}$.

Notice that because the states γ_t and μ_t are discrete, a switch in the estimated state can cause a large change in the interpolated data. In particular, if there is a downward switch in the state γ_t , then the corresponding interpolated datum \tilde{Z}_t can be negative. This can formally be interpreted as a loss of gold occurring at the respective dates.

8 Appendix 2. A one-stage estimation procedure

We could estimate the entire model jointly. Let

$$L_3 = \prod_{t=0}^T \Psi[\ln G_t - \ln G - \sum_{\iota=0}^t \mu_\iota - \ln \gamma_t; \sigma_G] \Psi[\ln p_t - \ln(p(\varsigma_t, \gamma_t)); \sigma_p] \quad (17)$$

Then, given a triplet of sequences μ , γ , and ς , we could search for the (κ, ρ) that maximizes

$$L_3 P(\mu) P(\gamma) P(\varsigma).$$

Then, we could apply the genetic algorithm to find the optimal triplet of sequences μ , γ , and ς . Because some parameters—e.g. γ_t —appear in both parts of L_3 , this joint procedure would not give the same estimates and would seem to be more efficient. However, the joint estimation is harder to carry out because of the larger set of parameters that must be searched over.

9 Appendix 3. Confidence intervals

Taking the logarithm of (12) yields:

$$\ln \tilde{G}_t = \ln \hat{G} + \sum_{i=1}^t \ln \hat{\mu}_i + \ln \hat{\gamma}_t + \tilde{\varepsilon}_t^G, \quad (18)$$

which is our “data” generating process. Then, given $\hat{\kappa}$, and the estimated pair of sequences for γ and μ by order, the (logarithm of the) fitted gold

stock series is:

$$\omega_t \equiv \ln \widehat{G} + \sum_{i=1}^t \ln \widehat{\mu}_i + \ln \widehat{\gamma}_t.$$

Then, conditional on $\widehat{\kappa}$ the test between the two alternatives $H_0 : (\gamma, \mu)$ and $H_1 : (\gamma', \mu')$ is the test between the sequence $\omega \equiv \{\omega_t\}_{t=1}^T$ of fitted gold stock implied by (γ, μ) , and the sequence $\omega' \equiv \{\omega'_t\}_{t=1}^T$ implied by (γ', μ') . Given normality of $\widetilde{\varepsilon}_t^G$ the test of the alternative against the null is the test of $H_1 : \omega'$ against $H_0 : \omega$, which is an F-test, while for every given date t the test that ω'_t differs from ω_t is a t-test. Given (18) the t-statistics for the test that one can distinguish between the two adjoining states $\widehat{\mu}^j$ and $\widehat{\mu}^{j+1}$ is

$$\widehat{t}_\mu = \frac{\ln \widehat{\mu}^{j+1} - \ln \widehat{\mu}^j}{\widehat{\sigma}_G} = \frac{\ln \left(\widehat{\mu}^1 + j \widehat{\Delta}_\mu \right) - \ln \left(\widehat{\mu}^1 + (j-1) \widehat{\Delta}_\mu \right)}{\widehat{\sigma}_G} \approx \frac{\widehat{\Delta}_\mu}{\widehat{\sigma}_G}.$$

The value of $\widehat{\Delta}_\mu / \widehat{\sigma}_G \approx 1.7$ implies that there is a roughly 10% chance that simulations alter a single entry in the sequence for μ by order. Thus, one should expect that about 10% of that sequence will differ if one allows for the more complete procedure.

As regards confidence intervals in the second stage, the “data” generating process is:

$$\ln \widetilde{p}_t = \ln \widehat{p}_t + \widetilde{\varepsilon}_t^p,$$

where

$$\widehat{p}_t \equiv p(\widehat{\varsigma}_t, \widehat{\gamma}_t),$$

is the estimated equilibrium price function. Notice that by definition (see 6) the equilibrium price function $p(\varsigma_t, \gamma_t)$ is not necessarily a monotone function of ς . Consequently, the t-statistics for the test that one can distinguish between the two adjoining values (pertaining to distinct points in the support of ς) of the equilibrium price function is:

$$\hat{t}_p = \frac{\ln \hat{p}^{j+1} - \ln \hat{p}^j}{\hat{\sigma}_p} \approx \frac{\frac{\tilde{p}^{j+1}}{\hat{p}^j} - 1}{\hat{\sigma}_p}.$$

Given the estimates in Table 2, the value of \hat{t}_p is approximately equal to 2, which implies that there is a roughly 5% chance that simulations alter a single entry in the estimated equilibrium price function. Thus, one should expect that about 5% of the sequence for ς will differ if one allows for the more complete procedure in the second stage.

References

- [1] Cagan, P. (1956). "The monetary dynamics of hyperinflation." In *Studies in the Quantity Theory of Money*, edited by M. Friedman. University of Chicago Press.
- [2] Christiano, L. (1987). "Cagan's model of hyperinflation under rational expectations." *International Economic Review* 28, 33-49.
- [3] Cosslett, S. and L. Lee. (1985). "Serial correlation in latent discrete variable models." *Journal of Econometrics* 27, 79-97.
- [4] Engsted, T. (1993). "Cointegration and Cagan's model of hyperinflation under rational expectations." *Journal of Money, Credit and Banking* 25, 350-360.
- [5] Goodfriend, M. (1982). "An alternative method of estimating the Cagan money demand function in hyperinflation under rational expectations." *Journal of Monetary Economics* 9, 43-57.
- [6] Jastram, R. W. (1977). *The Golden Constant*. John Wiley and Sons.

- [7] Salemi, M. and T. Sargent. (1979). “The demand for money during hyperinflations under rational expectations, II.” *International Economic Review* 20, 741-58.

- [8] Sargent, T.J. (1977). “The demand for money during hyperinflations under rational expectations, I.” *International Economic Review* 18, 59-82.

- [9] Taylor, M.P. (1991). “The hyperinflation model of money demand revisited.” *Journal of Money, Credit and Banking* 23, 327-51.

Table 1. Estimated parameters of the gold process.

Parameter	Estimate	5% confidence interval	
		Lower end	Upper end
G	28.5639	28.5085	29.0073
Δ_γ	0.024643	0.024313	0.024835
μ^1	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		

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		

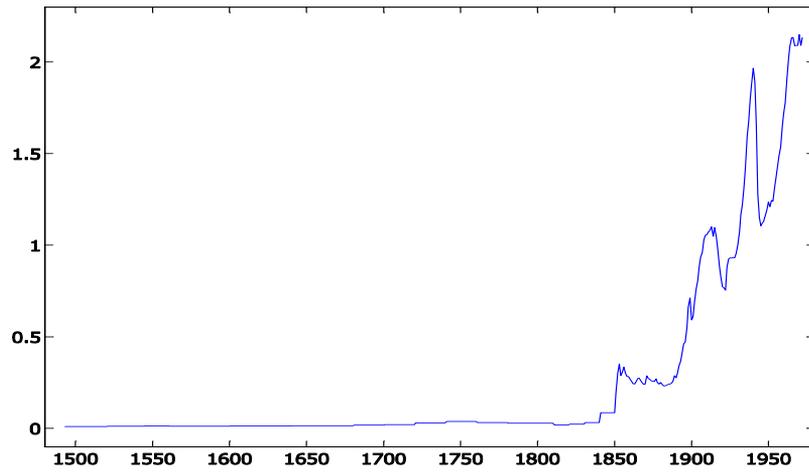


Figure 1: The Index of World Production of Gold, 1492-1972.

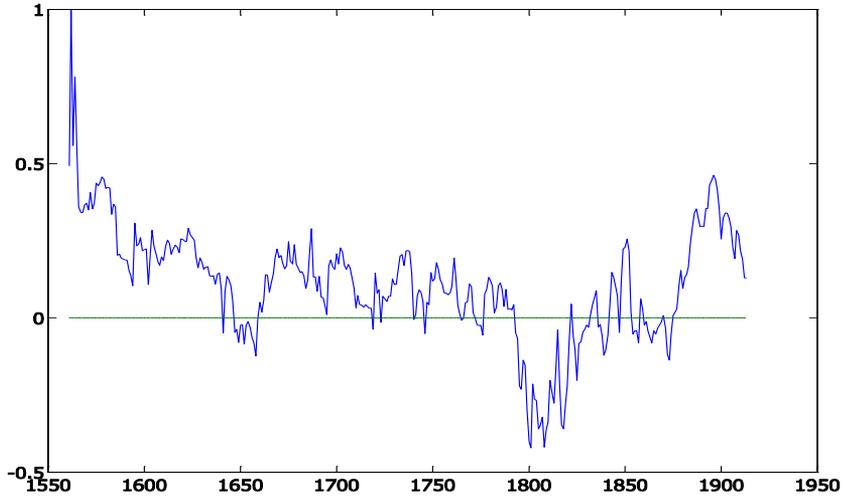


Figure 2: Logarithm of the purchasing power of gold, 1561-1913.

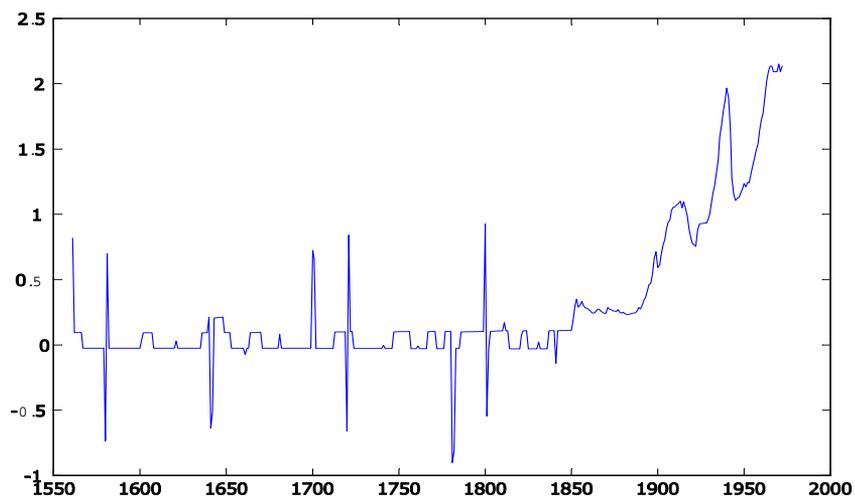


Figure 3: Gold production: interpolated up to 1850, actual after 1850.

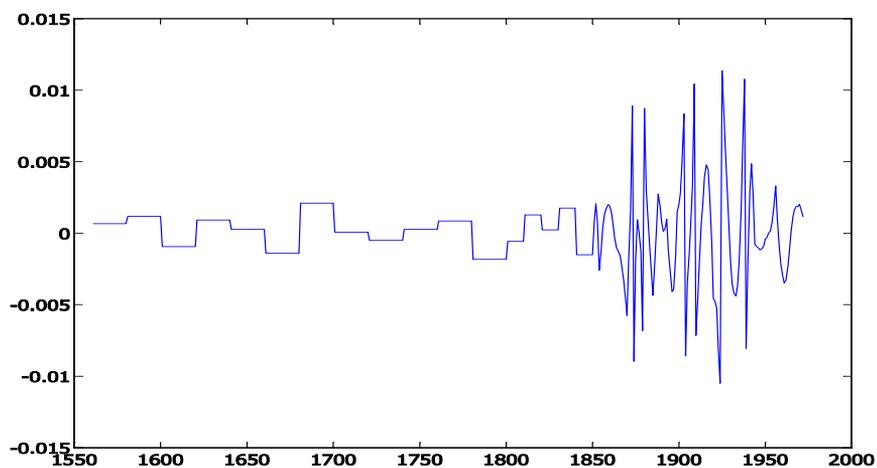


Figure 4: Estimated Gold Measurement Error $\hat{\varepsilon}_t^G$, 1561-1972.

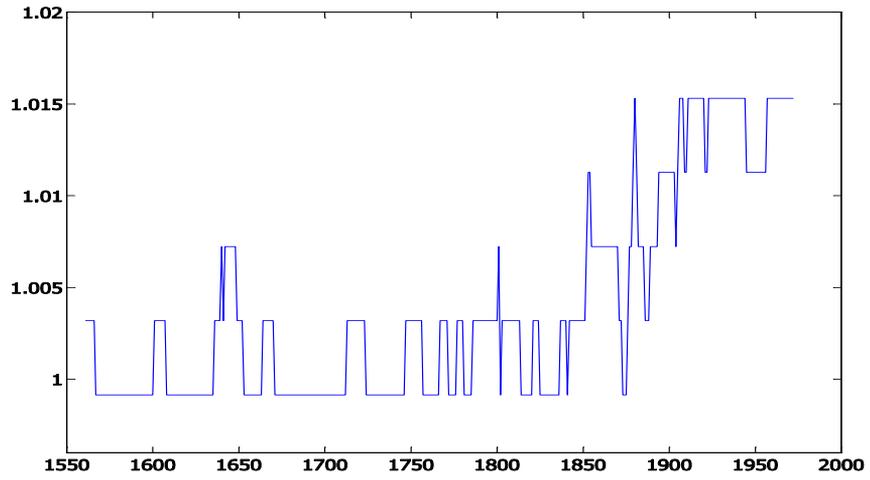


Figure 5: Estimated sequence for the permanent component of the gold stock, $\hat{\mu}$, 1561-1972.

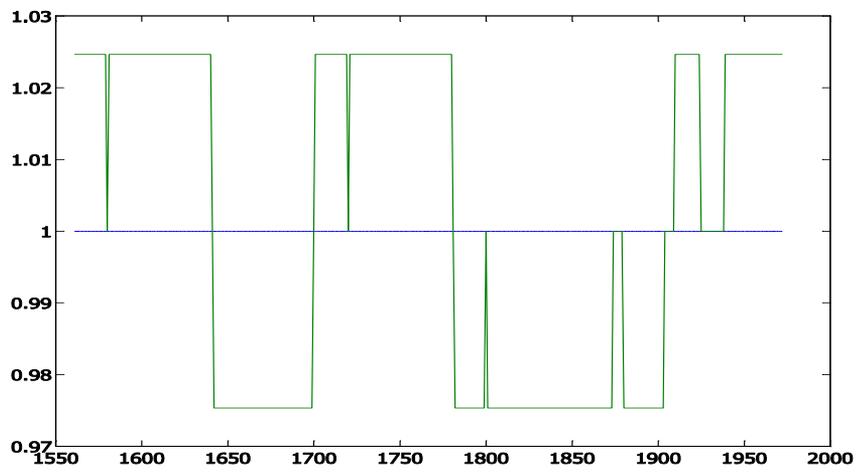


Figure 6: Estimated sequence for the temporary component of the gold stock $\hat{\gamma}$, 1561-1972.

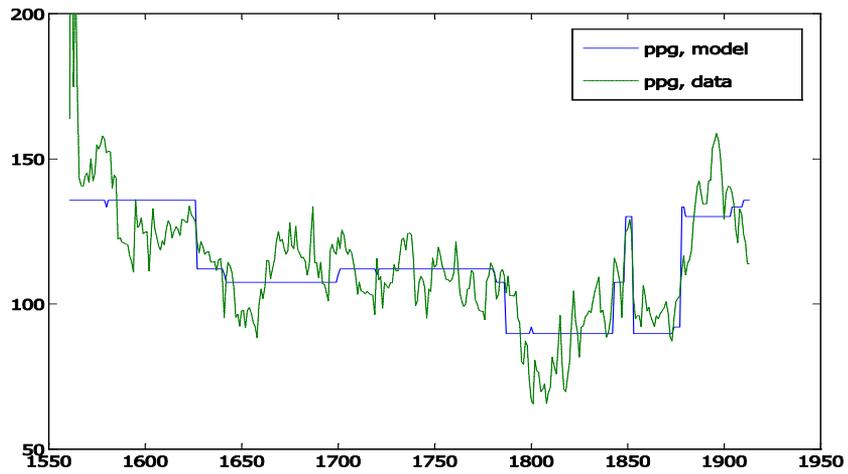


Figure 7: Purchasing power of gold: actual and fitted: 1561-1913.

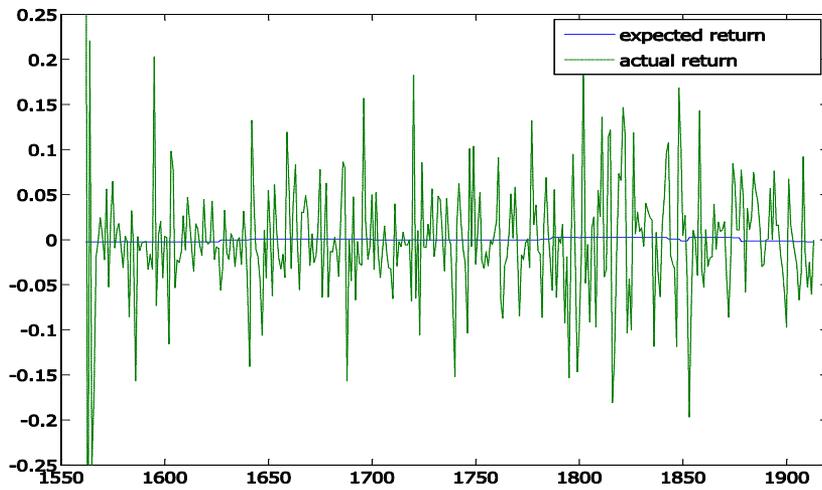


Figure 8: Actual and expected rates of return on gold, 1562-1913.

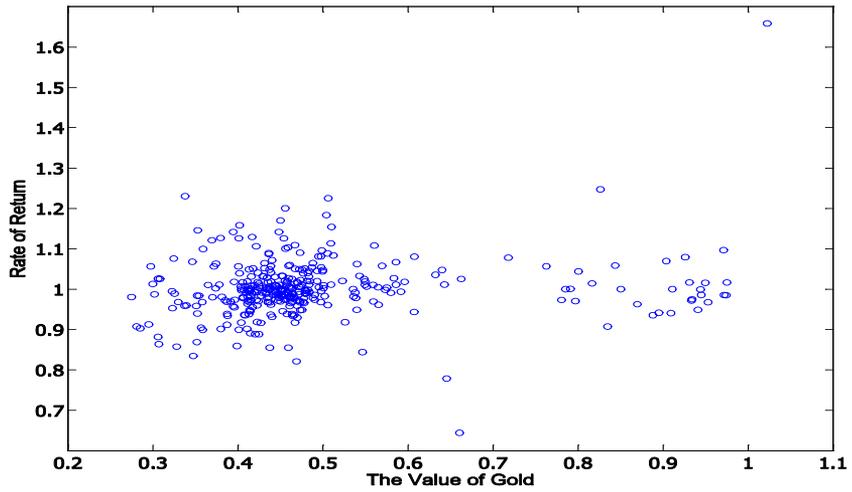


Figure 9: The value of the gold stock and the realized rate of return on gold. Actual data, 1562-1913.

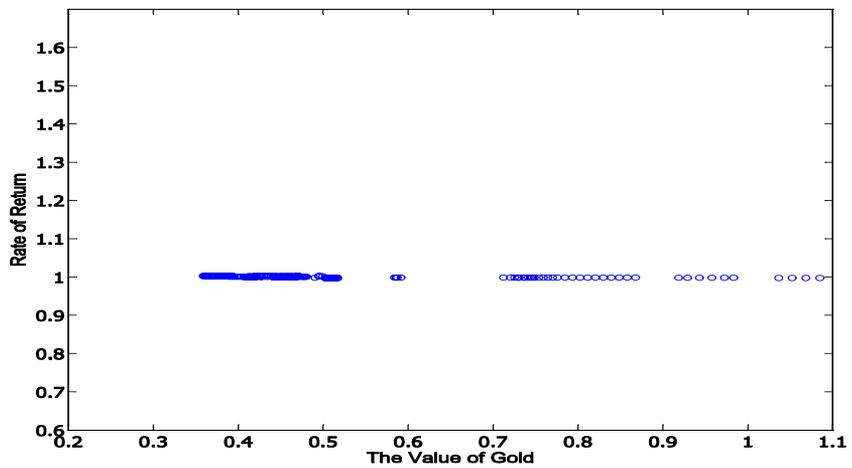


Figure 10: The value of the gold stock and the expected rate of return on gold. Fitted data, 1561-1913.