Testing Alternative Models of Non-neutrality with Disaggregate Data

William A. Bomberger
University of Florida

October 1998

Preliminary

Please Do Not Quote

I. Introduction

Changes in the rate of growth of money seem to have temporary effects on real variables but permanent effects only on inflation. Models consistent with long-run neutrality are easily constructed but no consensus model of the short-run non-neutrality of money has emerged. Two modern candidates which do not rely on money illusion or systematic forecast errors are the "misperceptions" model of Lucas (1973) and Barro (1976) and "menu costs" models such as the one of Ball, Mankiw, and Romer (1988), BMR hereafter.

Interest in these models centers on their differing predictions regarding the short-run behavior of output and employment. Specifically, the BMR model predicts that higher trend inflation leads to a steeper short-run Phillips curve while misperceptions models do not. However, each model provides predictions regarding a second type of non-neutrality: the effect of money growth (and other determinants of inflation) on relative prices. This paper exploits these predictions to provide evidence on the value of the alternative models using individual commodity prices from European hyperinflations.[1].

Such tests link models of non-neutrality with the literature on inflation dispersion. The dispersion literature begins with Mills' (1927) finding that the dispersion of individual prices is positively correlated with the overall inflation rate. This regularity has been examined and largely confirmed in numerous studies using a variety of measures for dependent and independent variables[2]. Most studies examine the robustness of this relationship with respect to data sets rather than the implications for choosing among models of price setting.[3]

An exception to this approach is Hercovitz's 1981 paper in this journal, frequently cited in the dispersion literature. Hercovitz showed that in a misperceptions model where elasticities of supply differ across goods, money shocks increase the dispersion of individual inflation rates. Using data on individual prices from the German hyperinflation he regressed dispersion, as measured by the variance of individual inflation rates across goods, γ^2_t, against the squared first difference of money shocks, $(m_t - m_{t-1})^2$, finding a significant positive relationship:

$$\text{(A)} \qquad \gamma^2_t = a_0 + a_1(m_t - m_{t-1})^2 \quad a_1 > 0$$

Several studies have followed Hercovitz's lead by estimating analogous relationships to test the misperceptions model and rival menu cost models. The predictions of menu cost models (such as BMR) regarding relative prices have not been worked out in equivalent detail, but a common prediction is that higher inflation leads to greater dispersion even if fully anticipated. Misperceptions models predict that anticipated increases in demand raise all prices equally. At the risk of oversimplifying this literature, one may say that a positive relationship between the expected or actual rate of inflation and dispersion is taken as support for the menu cost model. If only shocks affect dispersion, the misperceptions model is supported. Results of different studies have often been contradictory, at least in part because dividing price or money increases into expected and unexpected portions is problematical.[4]

A different approach is suggested here. Below I present estimates of the patterns in the individual prices series. Specifically, the change in the relative price of each good j, d_{jt}, is modeled as a heteroskedastic AR(1) process with a systematic aggregate shock effect:

$$\text{(B)} \qquad d_{jt} = b_j(v_t - v_{t-1}) + e_{jt} + \rho_{jt} e_{jt-1}$$

where v_t is a measure of aggregate shocks (such as m_t) and e_{jt} is an error specific to good j. The BMR model can be shown to imply that as the predictable portion of overall inflation increases: (a) the conditional variance of e_{jt} rises without limit and (b) ρ_{jt} approaches -1. In contrast, the Hercovitz model predicts: (a) the coefficients ρ_{jt} and the variances of e_{jt} should be invariant to the predictable portion of inflation, and (b) the estimated coefficients of (A) and (the N versions of) (B) should be satisfy: $\Sigma\, (b_j)^2/N = a_1$.

Single equation estimates, such as (A), discard much of the information residing in the individual series. In fact, the results show that estimates of dispersion relationships may be consistent with a particular theory even when the patterns in the individual series bear no resemblance to those predicted by the theory. Alternative tests based on individual inflation series are potentially more decisive.

To support this methodological point the Hercovitz study is used as a point of departure. I estimate dispersion relationships and individual relative inflation specifications ((A) and (B)) using disaggregate price data from periods of hyperinflation. To anticipate the results: The German data yield less support for the his model than Hercovitz' results imply, even when his money shocks are used. There is substantially more support for the menu costs model both from an expanded German data set and previously unexploited data from contemporaneous European hyperinflations which are used here and made available for others.

The paper is organized as follows: Section II derives predictions of the Hercovitz and BMR models regarding the behavior of relative prices. Predictions from the BMR model are new and rely heavily on simulations. Section III presents estimates of the specifications above applied to individual price data from periods spanning the hyperinflations and subsequent reforms in Austria, Hungary, and Germany. A brief summary concludes the paper.

II. Predictions of the Alternative Models Regarding the Behavior of Relative Prices

In this section models of inflation dispersion based on the Hercovitz model and the BMR model are sketched, rather than derived. In order to arrive at consistent estimating equations, I attempt to present the models in parallel ways. In particular, to accommodate the discrete time framework of the misperceptions model and the continuous time setup used by BMR, the question of time measurement is addressed in a hybrid way. Changes in the demand for goods occur frequently (hourly). The index h measures the number of hours since time 0. However, the econometrician only has data for the end of each month. The index t measures the number of months since time 0. Thus h = nt where n is the number of hours in a month.

A. The Hercovitz Model

The heart of any misperceptions model is a general specification of the demand in an individual market. Using the hourly time index:

$$(1) \qquad y^d_{jh} = -\alpha^d(P_{jh} - E(P_h)_{h-L}) + \beta(M_h - E(P_h)_{h-L}) + u_h + \epsilon_{jh}$$

where y^d_{jh} is the quantity of good j demanded at time h, P_{jh} is its price, P_h, is the price of the average good, M_h is the nominal money supply, u_h is the cumulative (nonmonetary) change in demand common to all goods since time 0, and ϵ_{jh} captures shifts in demand specific to good j (white noise with variance σ^2_ϵ). $E(\ )_{h-L}$ denotes the expectation of the variable in question conditional on information from L hours ago, when information on aggregate variables was last disseminated (thus L is the information lag in hours). All quantities are in logarithms and y is measured as the deviation from a benchmark level.

Hercovitz makes choices within this general specification. First, (following Lucas and Barro) he assumes that the information lag corresponds exactly to the sampling period of the data (L = n). This restrictive assumption facilitates the application of the model to monthly data. Second, he omits nonmonetary aggregate shocks, u, and sets $\beta = 1$, streamlining the model. Third, he allows the relative price elasticity of demand to vary across markets ($\alpha^d_j \neq \alpha^d_i$). The resulting equation (now using the monthly time index, t) is:[5]

$$(1') \qquad y^d_{jt} = -\alpha^d_j(P_{jt} - E(P_t)_{t-1}) + (M_t - E(P_t)_{t-1}) + \epsilon_{jt}$$

Money growth is given by $M_t - M_{t-1} = g_t + m_t$ where g_t is the rate of money growth which is predictable on the basis of information available at time t-1 and m_t is the unpredictable deviation from g_t distributed with zero mean and variance σ^2_m. The structural equations of the model are completed with a supply curve:[6]

$$(2) \qquad y^s_{jt} = \alpha^s_j(P_{jt} - E(P_t)_{t-1})$$

Solution of the model requires equating demand and supply and imposing rational expectations. The method of undetermined coefficients yields the results[7]:

$$(3) \qquad P_t = M_{t-1} + g_t + (\theta+(1-\theta)/\alpha)m_t$$

$$(4) \qquad P_{jt} = P_t + (\theta+(1-\theta)/\alpha_j)\epsilon_{jt} + \theta((1/\alpha_j)-(1/\alpha))m_t$$

where $\alpha_j \equiv \alpha^d_j + \alpha^s_j$, $\theta \equiv \sigma^2_m/(\sigma^2_m+\alpha\sigma^2_\epsilon)$, and α is the geometric average of α_j across goods. Agents cannot disentangle aggregate and relative demand shifts occurring between t-1 and t. Thus, unexpected money growth, m_t, has a nonunitary effect on average prices (last term in (3)) and a nonzero effect on real output (not shown). Expected money growth, g_t, is not mistaken for a relative shift. All prices are affected proportionally leaving output and relative prices unchanged. Thus, g_t appears with a coefficient of 1 in (3) and does not appear in (4).

I label the last term in (4) the "Hercovitz effect" because it demonstrates his central result. Demand shocks systematically change the prices of some goods more than others. Aggregate shocks are mistaken, in part, for relative shocks in all markets but they have a greater effect on prices where demand and supply curves are less elastic ($\alpha_j < \alpha$). If d_{jt} is relative inflation (defined as $\pi_{jt} - \pi_t = (P_{jt} - P_{jt-1}) - (P_t - P_{t-1})$) and γ^2_t is the dispersion of inflation ($= \Sigma(d_{jt})^2/N$):[8]

$$(5) \qquad d_{jt} = (\theta+(1-\theta)/\alpha_j)(\epsilon_{jt}-\epsilon_{jt-1}) + ((1/\alpha_j)-(1/\alpha))(1-\theta)(m_t-m_{t-1})$$

$$(6) \qquad \gamma^2_t = 2\sigma^2_\epsilon((\theta+(1-\theta)/\alpha)^2+\sigma^2_\alpha(1-\theta)^2) + (1-\theta)^2\sigma^2_\alpha(m_t-m_{t-1})^2$$

where σ^2_α is the variance of $1/\alpha_j$ across goods. (6) corresponds to: (A) $\gamma^2_t = a_0 + a_1(m_t\text{-}m_{t-1})^2$ referred to in the introduction, with $a_1 \equiv (1\text{-}\theta)^2\sigma^2_\alpha$. With identical markets $\sigma^2_\alpha = a_1 = 0$. Note that (5) could also be estimated (for each market) as:

(B) $d_{jt} = b_j(m_t - m_{t-1}) + e_{jt} - e_{jt-1}$

where $e_j \equiv (\theta+(1\text{-}\theta)/\alpha_j)\epsilon_j$ and $b_j \equiv((1/\alpha_j)\text{-}(1/\alpha))(1\text{-}\theta)$. The model predicts a cross-equation restriction: $a_1 = \Sigma\ (b_j)^2/N$. Where a_1 comes from (A) and the b_j from the N estimates of (B).

This last point is worth emphasizing. The Hercovitz model does not simply predict that months with large money shocks should also exhibit a large dispersion of inflation rates. It predicts that goods which tend to show above (below) average inflation rates when $m_t\text{-}m_{t-1}$ is positive should exhibit below (above) average inflation rates when money shocks are negative. Moreover, the systematic effects of shocks in the individual equations, measured by b_j, should be large enough (when squared and summed) to account for the effect of shocks on dispersion, measured by a_1, in the dispersion equation. These predictions can only be addressed by adding estimates of the individual inflation relationships.

B. The BMR Model

One encounters two difficulties in deriving predictions about price dispersion and relative inflation from the BMR model.[9] First, unlike Hercovitz, the authors were not interested in the behavior of relative prices, hence they derived no analogous expressions. Second, their continuous time framework does not give rise to similarly “clean” expressions. My attempt to surmount these difficulties involves three steps. First, the BMR model is used to derive expressions roughly parallel to those in Hercovitz. Second, these expressions are used to provide a basis for conjectures on the behavior of relative inflation rates in this model. Third, simulated data derived from these expressions are used to confirm these conjectures.

BMR do not separate money from other influences of aggregate demand. Thus, demand at time h is given by: $(g/n)h + u_h$ where g is identified with the trend growth (per month) in demand and u is a measure of cumulative demand shocks since time 0. Demand shocks are described by a random walk: $u_h = v_h + v_{h-1} + v_{h-2} \ldots$, where v is the hourly (white noise) innovation in demand. There is a relative demand shock ϵ_{jh} as in Hercovitz. The crucial innovation is that firms change prices only once every λ hours. The fixed price period is identical across firms but the timing of price changes is staggered randomly.

This fixed price period introduces two complications to the determination of relative prices. First, prices depend not on current demand but upon the conditions which prevailed when prices were most recently changed. Second, when firms change prices, they take into account their expectation of conditions over the next λ hours, including a forecast of the future reaction of other firms to recent demand shifts. Specifically, if firm j changes prices at hour h:[10]

$$(7) \qquad P_{jh} = (g/n)h + \epsilon_{jh} + (g/n)\lambda/2 + q_0 v_h + q_1 v_{h-1} + \ldots$$

where the q's are endogenous weights which approach 1 from below as the lag increases.

The terms $(g/n)h$ and ϵ_{jh} show that both predictable growth in aggregate demand and relative shifts affect prices one-for-one. The term $(g/n)\lambda/2$ indicates that price setting at hour h anticipates the average change in demand occurring over the next λ hours during which P_j remains fixed at $P_{jh.}$ The final terms show that demand shocks affect prices less than one-for-one (the q's are less than 1) and have effects on output (not shown). The reason firms do not increase prices in proportion to observed demand shocks is that they know competing firms will not respond until the end of their own fixed price periods.

But (7) applies only to prices that are set anew at hour h. More generally, P_j is the price which seemed optimal at the time it was most recently set, k hours ago, where $0 < k < \lambda$. Thus:

$$(7') \qquad P_{jh} = (g/n)(h-k) + \epsilon_{jh-k} + (g/n)\lambda/2 + q_0 v_{h-k} + q_1 v_{h-k-1} + \ldots$$

Staggered price setting means that k varies randomly from 0 to λ. I emphasize this with the expression k(j,h) or k_{jh} (hours since P_j was last changed as of hour h). P is the mean of the P_j:[11]

$$(8) \qquad P_h = (g/n)h + w_0 v_h + w_1 v_{h-1} + \ldots$$

where the w's are averages of the q's. (7') is parallel to (4) in the Hercovitz model and (8) parallels (3). To construct a parallel to (5), the month-to-month relative inflation rate, one subtracts (8) from (7') and subtracts a similar expression dated n hours (one month) in the past:

$$(9) \quad d_{jt} = g(k_{jh-n} - k_{jh}) + (\epsilon_{jh-k(j,h)} - \epsilon_{hs-n-k(j,h-n)}) + q_0 v_{h-k(j,h)} + q_1 v_{h-k(j,h-1)} + \ldots - w_0 v_h - w_1 v_{h-1} - \ldots$$
$$- q_0 v_{h-n-k(j,h-n)} - q_1 v_{h-n-k(j,h-n)} - \ldots + w_0 v_{h-n} + w_1 v_{h-n-1} + \ldots$$

By defining $q_{-1} = q_{-2} = \ldots = 0$, switching to a monthly time index, and adopting some notational change (9) can be made more compact:

$$(9') \qquad d_{jt} = g(k_{jt-1} - k_{jt}) + \Delta\epsilon_{jt} + a_{0jt} v_h + a_{1jt} v_{h-1} + \ldots - a_{0jt-1} v_{h-n} - a_{1jt-1} v_{h-n-1} - \ldots$$

where k_{jt} is the time between the most recent change in P_j and the end of month t (hour h), k_{jt-1} is the equivalent interval for month t-1, $\Delta\epsilon_{jt}$ is the cumulative relative shock between the most recent change of P_j and the last price change of month t-1, and $a_{ijt} \equiv (q_{i-k(j,t)} - w_i)$.

Comparing this to the Hercovitz model, the second term, $\Delta\epsilon$, corresponds to the first term in (5). Since the menu cost model assumes perfect information, relative demand shocks are not mistaken for aggregate shocks. Thus the coefficient on $\Delta\epsilon$ is 1 in (9') and less than 1 in (5).

However, relative demand shocks are not easily observable, and the difference in results provides no basis for testing the alternative models.

The terms in v from (9') describe the effects of demand shocks (as opposed to trend growth) on relative inflation and correspond to the m_t-m_{t-1} term in (5) Unlike the Hercovitz model, the BMR model does not provide clear predictions regarding the effect of cumulative end-of-month to end-of-month shocks on monthly relative inflation[12]. Predictions based on numerical simulations are derived below. Predictably, the BMR model might or might not generate a positive relationship between demand shocks and inflation dispersion, depending on parameter values. Thus the response of relative prices to aggregate shocks also fails to provide a decisive test of the alternative models.

C. Trend Inflation and Dispersion in the BMR Model

The novel term in (9') is $g(k_{jt-1} - k_{jt})$. It describes the effect of trend inflation, g, on relative inflation. There is no corresponding term in (5) because predictable changes in demand have no effect on relative prices in a misperceptions model. This term is the source of decisive testing of the alternative models. The intuition is clear. Part of the difference between monthly inflation for good j and monthly inflation overall depends on whether the current fixed price period for good j began shortly before or long before the end of month t . If the last hour of month t came toward the end of a fixed price period, k_{jt} is near its maximum, λ, prices overall moved ahead of P_{jt}, and d_{jt} tends to be negative. Conversely, if k_{jt-1} was large, P_{jt-1} began the month behind overall prices and d_{jt} tends to be positive as P_j catches up.

With large trend inflation, the timing of the last price change is important and the effect is large. This is a “timing”effect. The upper panel of Figure 1 depicts paths of the aggregate price level (P) and the price of an individual good (P_j) over several “shock-free” months. In the lower panel a “month” is superimposed to give a visual representation of the relationship between π_{jt} and π_t (as they occur in monthly data) and the values for k_{jt} and k_{jt-1} (unobservable from such data). The dotted line in the lower panel is steeper than the timepath of P in the upper panel. This indicates that $\pi_{jt} > \pi_t$ and $d_{jt} > 0$. d_{jt} is positive because $k_{jt} < k_{jt-1}$ for this month. The difference in slopes would disappear (d = 0) if π or λ =0.

More formally, we can examine the effect of trend inflation and the length of the fixed price period on the dispersion of inflation rates. Squaring (9') and averaging across goods yields:

$$(10)\quad \gamma^2_t = \Sigma(d_{jt})^2/N = g^2\Sigma(\Delta k_{jt})^2/N + g\,\Sigma\Delta k_{jt}(\text{terms in } \epsilon \text{ and } v)/N + \Sigma(\text{terms in } \epsilon^2 \text{ and } v^2)/N$$

Random staggering means k is a random variable with a uniform distribution from 0 to λ. If N is large one may replace the average of random terms with their expected values (as in Barro and Hercovitz). Thus $\Sigma(\Delta k_{jt})^2/N = \lambda^2/6$, and $\Sigma\Delta k_{jt}(...)/N = 0$, yielding:

$$(10')\qquad \gamma^2_t = \Sigma(d_{jt})^2/N = g^2\lambda^2/6 + \text{terms in } \sigma^2_\epsilon + \text{terms in } v^2$$

The term in squared trend inflation is unique to the menu cost model and is an artifact of the timing effect. Hercovitz adds a term in squared actual (as opposed to unexpected) money growth to his own specification as a test of this type of effect (p.355). He refers to Sheshinski and Weiss as a source for this specification. However, the BMR model implies a somewhat different specification. Since λ is endogenous, higher trend inflation leads to more frequent changes in prices (lower λ). No simple functional form arises for this relationship, but Mussa (1981) and Rotemberg (1983) show that a fixed menu cost combined with a quadratic disequilibrium cost yields a fixed price period which decreases with the 2/3 power of trend inflation. That is, $\lambda = Cg^{-2/3}$. In this case (10') becomes:

$$(10'')\qquad \gamma^2_t = \Sigma(d_{jt})^2/N = g^{2/3}C^2/6 + \text{terms in } \sigma^2_\epsilon + \text{terms in } v^2$$

Intuition for this result can also be gleaned from Figure 1. If trend inflation were twice as large, the trajectory of P would be twice as steep. However, the "steps" for P_j would be only 1/3 "higher" because they would also be 2/3 less "wide". The timing of the most recent price change relative to the end-of-month relative would be more important, but not proportionately more important. Therefore, the quadratic form of the misperceptions model is not appropriate for the BMR model. An inclusive functional form is $\lambda = Cg^{-\beta}$ which implies:[13]

(A') $\gamma^2_t = a_0 + a_1(g_t)^{\phi}$ $\quad a_1 = C^2/6 > 0; \quad \phi = 2 - \beta < 2.$

D. Demand Shocks and Dispersion in the BMR Model

Is the relative inflation rate of good j during month t increased or decreased as a result of a positive demand shock i hours before the end of the month? This answer is conveyed by the sign of a_{ijt} in (9') above. It depends crucially on the (random) value of k_{jt} that month. If P_j changed shortly before month t ended (k_{jt} small) then P_j will have adjusted to recent demand shocks more fully than prices overall and the a_{ijt} will tend to be positive. If the month ended just before the end of firm j's fixed price period (k_{jt} close to λ), the average price will have adjusted to recent demand shocks more fully than P_j and the a_{ijt} will tend to be negative.

Figure 2 provides some intuition. Here a single positive demand shock, v_t, has occurred during the month in question. It shows up as a jump in P superimposed on the positive trend. Firm 1 changed prices after the shock and incorporated the shock into the price increase. Firm 2 changed prices earlier in that particular month and has not yet had an opportunity to react ($k_{2t} > k_{1t}$). Thus, the shock has tended to increase the gap between d_{1t} and d_{2t} and increase dispersion overall. Shocks, like trend inflation, contribute to the "timing effect" increasing inflation dispersion. This has two implications: First, the effect of shocks on dispersion is not the unique hallmark of the Hercovitz model and may not be the source of decisive testing of the models. Second, the use of actual inflation as a regressor in a dispersion relationship (hence combining trend and shocks) is supported.

Demand shocks affect dispersion in the BMR model, but there is an important difference. Unlike the Hercovitz result, there should be no consistent relationship (over time) between the demand shocks and relative inflation for a given good. A positive demand shock might increase the relative inflation rate of good j one month and decrease it in a later period (when k_j happens to be larger). That is, the a_{ijt} are themselves random variables and will be positive one month but negative the next for the same good. In the Hercovitz model, goods for which the supply or demand curve is steeper will **consistently** exhibit a larger than average price response to aggregate shocks, That is, the misperceptions model analogues for a_{ijt} (b_j in (B)) should have the same sign each month for a given good.

E. Individual Relative Inflation Rates in the BMR Model

The strategy of Hercovitz and all others of which I am aware is to estimate a version of (A): A single equation with dispersion as the LHS variable. This discards considerable information contained in the individual series. Estimating (5) or its menu cost analogue directly can exploit this information. The proposed specification is: (B) $d_{jt} = b_j(v_t - v_{t-1}) + e_{jt} + \rho_{jt}e_{jt-1}$ where v_t is a measure of aggregate shocks (a money shock or other demand shock), and e_{jt} is an error.

(5) can be viewed as a version of (B) with restrictions: (a) $v_t = m_t$, (b) the b_j are fixed over time but vary across goods ($b_j = ((1/\alpha_j)-(1/\alpha))(1-\theta)$), and (c) $\rho_{jt} = 0$ ($e_{jt} = (\theta+(1-\theta)/\alpha_j)\Delta\epsilon_{jt}$ where $\Delta\epsilon_j$ is white noise). However, Hercovitz's assumption that relative shocks are random walks is a convenient one, rather than an essential. Therefore, in estimates below, violations of (c) ($\rho_j \neq 0$) are not taken as evidence for rejection of the Hercovitz model.

With less precision (B) can be seen as a version of (9') where v_t measures demand shocks. In the BMR model the b_j ($\approx a_{jt}$) vary randomly about 0 over time implying estimated b_j near 0 for the sample. We can examine the structure of the error terms in the BMR version of (B) by putting aside terms in v_t in (9') and again exploiting the uniform distribution of k:

(11a) $E((d_{jt})^2) = g^2E((k_{jt-1}-k_{jt})^2) + E((\Delta\epsilon_{jt})^2) = g^2\lambda^2/6 + \sigma^2_{\epsilon},$

(11b) $E(d_{jt}d_{jt-1}) = -g^2E((k_{jt-1})^2) + E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-1}) = -g^2\lambda^2/12 + \mu_1\sigma^2_{\epsilon},$

where σ^2_{ϵ} is the variance and μ_1 is the first order serial correlation of ϵ (expectation of terms in $k\Delta\epsilon$ disappear). If relative demand is a random walk ($\mu_1 = 0$), and λ depends on g ($\lambda = Cg^{\phi-2}$, if (A') applies). Thus, d_{jt} fits an MA(1) process: $d_{jt} = e_{jt} + \rho e_{jt-1}$ where:

(11a') $E((d_{jt})^2) = (1+(\rho)^2)\sigma^2_e = g^2\lambda^2/6 + \sigma^2_{\epsilon} = g^{\phi}C^2/6 + \sigma^2_{\epsilon}$

(11b') $E(d_{jt}d_{jt-1}) = \rho\sigma^2_e = -g^2(\lambda_j)^2/12 = -g^{\phi}C^2/12$

Solving (11a') and (11b') for ρ and σ^2_e:

(12a) $\rho = -(1 + (6\sigma^2_{\epsilon}/C^2)g^{-\phi}) + ((1 + (6\sigma^2_{\epsilon}/C^2)g^{-\phi}) - 1)^{1/2}$

(12b) $\sigma^2_e = (g^{\phi}C^2/12)(1 + (6\sigma^2_{\epsilon}/C^2)g^{-\phi}) + ((1 + (6\sigma^2_{\epsilon}/C^2)g^{-\phi}) - 1)^{1/2}$

As g becomes large ρ approaches -1 and σ^2_e approaches $g^{\phi}C^2/12$ (assuming $\phi > 0$). This can be seen in (11a'): σ^2_{ϵ} becomes negligible by comparison as the term in g increases without limit. Thus, as g rises, ρ approaches -1 and σ^2_e increases without limit, although not necessarily in proportion to g. In estimates below the assumption that ϵ is a random walk is relaxed as it was with the Hercovitz model. Thus, ρ need not approach 0 as g = 0.

These results are intuitive. Random variation in the timing of the most recent price change contributes to the volatility of individual inflation rates. Higher trend inflation increases volatility unless it greatly reduces the length of the fixed price period as well. Relative inflation for a given good does not rise or fall systematically with trend inflation. Rather the variances of all inflation rates increase during a period of higher trend inflation.

This timing effect also introduces a negative relationship between successive individual inflation rates. If the end of the month occurs soon after the last price change for good j, (k_{jt} near 0), inflation for good j tends to be larger than average inflation (assuming $g > 0$). Good j then starts the next month with a higher price. Since P_j is the base price for the next monthly inflation rate, a lower than average inflation rate will follow most of the time. As g rises this negative correlation comes to dominate other sources of relative inflation. Both of these patterns are found in synthetic data generated by a BMR model in the next section.

F. Simulations Results for the BMR model

The continuous time BMR model does not yield explicit estimating equations applicable to discrete monthly observations . Therefore I have relied heavily on intuition to establish relative price predictions of the BMR model. This section reinforces these predictions using synthetic data generated from a simulated BMR model. I anticipate the results of the empirical section to come by applying the estimation procedures of that section to these synthetic data. The results show that if (A’) and (B) are estimated using data consistent with the BMR model, the patterns described in the previous section do emerge.

The BMR model allows many choices for parameterization. In choosing among them I have attempted to mimic the fluctuations in overall inflation exhibited in the German data described in the next section. Baseline data with 60 observations of 30 individual price series are generated using (7'). Deterministic demand growth fits a trigonometric function which from 0% to 100% per month. The phase and period are selected to produce a single peak roughly at the midpoint of the sample. Aggregate demand shocks come from a pseudo-random number generator mimicking independent errors with a standard deviation of 30% per month. Thus, most of the monthly variation in inflation is due to shocks. Relative shocks have a standard deviation 1/4 as large. This conforms to the notion that aggregate shocks dominate relative shocks in such turbulent times. I deviate from the BMR model in one respect: λ does not depend on g. Rather $\lambda_j = (2/3)n$ for all j (and all t). Thus prices change somewhat more frequently than observations occur. This is roughly consistent with the number of unchanged prices observed in the German data. Starting dates for fixed price periods are staggered randomly within the month.

Table 1 displays the resulting baseline data for γ and π. Estimating (A') or (B) requires separating inflation, π_t, into a predictable part, g_t, and an unpredictable (shock) part, v_t. In the empirical section (and here), I follow the familiar procedure of estimating a timeseries process for inflation over the full sample. Rational expectations are invoked to identify g_t with the conditional forecast for month t based on the estimated relationship. The resulting error, π_t - g_t, is identified with the v_t., The observer is assumed to have no direct knowledge of the process generating the data. Rather he or she fits an ARMA model. The AR(1) passes the usual specification tests as a description of the synthetic π data. Table 1 displays the estimated parameters of the AR(1) and the resulting g_t and v_t series (with π_t and γ_t).

Estimates of the single equation dispersion relationship (A') for the synthetic data should be consistent with the BMR model. Individual inflation rates should be highly dispersed when trend inflation, g, or actual inflation, π, is large, not just when shocks are large. OLS estimates of (A') are displayed in Table 2. Lines (1) - (3) report estimates using baseline data and the standard quadratic form (ϕ = 2) for three alternative RHS variables: actual inflation, trend inflation, and the first difference of inflation shocks. Estimates reported in line (4) use actual inflation and allow $\phi \neq 2$. Three predictable patterns emerge: (a) demand shocks (the Hercovitz term) have no explanatory power, (b) both actual inflation and expected inflation have significant explanatory power but actual inflation dominates, (c) allowing a non-quadratic form ($\phi \neq 2$) makes little difference. The results are robust to broad variation of the parameter values used to generate the data (not shown).

It is not surprising that a quadratic form fits the data well. With λ fixed, the mechanism that generates a non-quadratic relationship has been omitted. The absence of a Hercovitz effect is also to be expected. However, this effect could easily emerge in a menu cost model. To generate his effect Hercovitz simply added cross-market heterogeneity to Barro's model. Heterogeneity could add realism to the BMR model as well. An obvious candidate is heterogeneity in λ. If the price of good j is changed more frequently than the average good ($\lambda_j <$ mean λ) the price of good j will, on average, be more responsive to demand shocks than the average price. Thus d_j will tend, on average, to be positive (negative) when shocks are positive (negative). This pattern will be superimposed on the random variation related to the timing effect and will be reversed for $\lambda_j <$ mean λ.

The results in line (5) use simulated data based on heterogenous λ instead of the baseline data. λ is set at 1/3 of a month for 15 goods while $\lambda = 1$ month for the other 15. Thus, the mean value of λ is identical to the baseline data, but rapidly adjusting firms change prices 3 times as often as the others. Squared changes in shocks are now significant even when combined with squared trend inflation. However, such significant Hercovitz effects are not robust to changes in the model parameters (not shown). Predictably, they emerge for synthetic data with considerable heterogeneity in λ and large aggregate shocks.

The methodological point of this paper is that patterns in individual price series are more informative than dispersion relationships like (A). As a vehicle for verifying such patterns Table 3 presents maximum likelihood estimates of (B) $d_{jt} = b_j(v_t - v_{t-1}) + e_{jt} + \rho_{jt}e_{jt-1}$

based on the synthetic BMR data. As suggested by the discussion of the previous section, ρ_j and σ^2_j (the variance of e_j) are expressed as functions of g. The functional forms are: $\rho_{jt} = (R_j - rg_t)/(1 + rg_t)$ and $\sigma^2_{jt} = S_j(1 + sg_t)$. These are parsimonious versions of (12a) and (12b). As g approaches 0, ρ_{jt} and σ^2_{jt} approach the values R_j and S_j. As g increases without limit σ^2_{jt} does the same (if s >0) and ρ_{jt} approaches -1 (if r > 0). Since b_j, R_j, and S_j are idiosyncratic, there are 30 separate values for each, but r and s are common to all series.[14]

Results in rows (1) - (5) use the baseline data. In (1) all b_j = all R_j = r = s = 0. Thus, Hercovitz effects and serial correlation are not allowed and homoskedasticity (over time) is imposed. Values under S are the mean (standard deviation) of the 30 individual estimated variances ($S_j = \sigma^2_j$). The results in (2) allow heteroskedasticity ($s \neq 0$). The conditional variance of each series rises proportionally with trend inflation (estimated s > 0) . The column labeled σ^2_{max} lists the means (standard deviations) of individual variances which occur if trend inflation is 100% per month (g = 1), the approximate maximum value of g in the sample. Thus, entries under S and σ^2_{max} provide summary statistics for the approximate minimum and maximum estimated conditional variances for the sample.

Row (3) allows AR(1) errors ($R_j \neq 0$) but requires each ρ_j to be constant over time (r = 0). The means (standard deviations) of the individual R_j (= ρ_j) appear under R. They are predominantly negative but heterogeneous (standard deviation = .33). The addition of serial correlation decreases heteroskedasticity (s is smaller) but the effect of trend inflation on variances remains significantly positive. Row (4) displays estimates which allow for the predicted relationship between g and ρ ($r \neq 0$). Since estimates of r are positive, higher trend inflation moves each ρ_j toward -1, reaching the values listed under ρ_{max} when g = 1. Thus, mean ρ_j falls from -.21 to -.94 and the standard deviation falls from .46 to .03 as g varies from 0 to 100% per month. That is, when inflation is low, the correlation patterns of the underlying relative shifts (measured by R_j) are quite different, but during high inflation the timing effect dominates, producing correlation parameters tightly clustered near -1.

Relaxation of the restriction s = 0 (r = 0) causes a change in likelihood from row (1) to (2) (row (3) to (4)). Twice each change in likelihood is distributed as $\chi^2(1)$. The value 570.6 (130.8) is far greater than the 1% threshold of 6.6. Between (2) and (3), 30 restrictions are relaxed ($R_j = 0$). The value 482.0 easily exceeds the 1% threshold (50.9) for a $\chi^2(30)$. Estimates in row (5) allow a Hercovitz effect ($b_j \neq 0$). Predictably, there is little evidence of one. Twice the change in likelihood from (4) to (5) is 42.8 which falls short of 50.9.

This may is consistent with the near zero coefficient in Table 2, line (3), and is forced by the homogeneity of λ_j in the baseline data. The heterogeneous λ version of the synthetic data are used for the results in (6) and (7) . The effects of g on σ^2_j and ρ_j remain (s and r are significantly positive). Removing the restriction that all $b_j = 0$ (moving from (6) to (7)) now yields a significant improvement in likelihood, consistent with a significant coefficient on $(v_t\text{-}v_{t-1})^2$ in line (5) of Table 2. These results verify that if (B) is estimated using data generated by a BMR mechanism: (a) the error terms will be heteroskedastic as conditional variances vary with g. (b) The coefficients on the lagged error terms will approach -1 as g increases. (c) A significant Hercovitz effect (b_j collectively $\neq$ 0) may or may not emerge.

III. Empirical Results

A. Data

In order to see if the patterns generated above occur in real world data, I focus on the commodity data from which Hercovitz derived his dispersion measure. I also extend the sample for these data and present analogous data for two other European hyperinflations.

Hercovitz chose 68 monthly series of wholesale commodity prices with which to calculate the variance of monthly inflation rates across commodities, γ^2. Other series which had a significant number of missing observations over the period Jan 1921 - Jul 1923 were excluded. He used a hypothesized money supply function to divide monthly money growth into anticipated and unanticipated components, g and m, for Nov 1920 - Jul 1923.

For most of these commodities data has been located for the period January 1920 through December 1924[15]. Data on the mean, π, and variance, $\gamma^{2,}$, of individual inflation rates, along with N, the number of non-missing observations for each month, are shown under Germany A in Table 4. About 30 of the series are unavailable before 1921. Therefore, alternative statistics based on the 37 series available over the entire 60 month period (with few missing observations) are presented under Germany B. Using Hercovitz's categories, Germany A contains 27 foodstuffs, 19 textiles and leathers, and 22 metals, oils, and coals. Germany B contains 17, 10, and 10 items in the respective categories. Observations for months which meet Cagan's (1956) definition of hyperinflation appear in bold.

Similar data have been discovered for the nearly contemporaneous Austrian and Hungarian hyperinflations. The Austrian data are monthly observations of consumer prices including 28 fabrics and clothing items, 17 furnishings and hardware, and 6 categories of energy. An additional 5 series were excluded due to missing observations for Jan 1921 - Dec 1923.[16] There are two sets of Hungarian data. The first (Hungary A) contains retail prices for 34 foodstuffs in Budapest for Jan 1923 - Dec 1925. Hungary B contains wholesale prices in Budapest for 43 foodstuffs, 8 metals and oils, and 8 fabrics and building materials for the same period. Table 4 lists means and variances for all data sets and the number of non-missing observations for Hungary B (no missing observations for Austria or Hungary A).[17]

Data of this sort provide extreme variations in the rate of inflation which can potentially be used to discriminate between models. Each set includes a period before the onset of hyperinflation, the period of peak inflation, and some period after the ultimate monetary reform. The German data, in particular, fit into quite heterogenous periods. For Germany A, the hyperinflation period Aug 1922 - Nov 1923 has means (standard deviations) for π and γ^2 of 1.457 (1.784) and .117 (.103). It was preceded by 30 months with values of .064 (.141) and .033 (.028) and followed by a 13 month post-reform period with values of .001 (.047) and .015 (.015). The others show similar, albeit less dramatic, patterns. Because German data display the greatest variation of inflation and have a large number of timeseries observations, results based on Germany B are emphasized below. Results from other data sets are shown as confirmation.

B. Separating Trend Inflation from Inflation Shocks

Anticipated inflation has no effect on relative prices in a misperceptions model. Various effects of anticipated inflation on relative inflation define the menu cost implications for (A') and (B). Hence inflation must be decomposed into trend and shocks to test these models. A similar separation of money growth was Hercovitz's first step. It is difficult to have confidence in any particular ex post decomposition. I deal with this by attempting alternative decompositions and examining the results for robustness with respect to the choice. In each case, a timeseries model of inflation is estimated for the full sample and trend inflation is identified with the one-step-ahead conditional expectation of this process.

An AR(1) process ($\pi_t = \alpha_0 + \alpha_1\pi_{t-1} + v_t$) fits the data fairly well. A look at Table 4 or a formal specification test confirms the pattern of heteroskedasticity in v_t. Hence an ARCH process ($\sigma^2_{vt} = \delta^2_0 + \delta_1(v_{t-1})^2$) is estimated (see Engle (1982)). This is specification (i). Estimates of (i) for all countries appear in Table 5. An appealing, more complicated interpretation of the data is that there are two processes which generate inflation: (a) a low variance AR(1) which describes the pre-hyperinflation and post-reform data, and (b) a much higher variance, more persistent process which generates the hyperinflation proper.

For the second specification I assume process (b) fits data from Cagan's hyperinflation months. For this period, one cannot bound the coefficient on lagged inflation away from 1, hence inflation is modeled as a random walk. An AR(1) process is fitted to the remaining sample. With the divided sample, evidence for heteroskedasticity **within** either process is greatly reduced. Thus, this mixed process is estimated as specification (ii) with separately homoskedastic errors (results in table 5).

A weakness of (ii) is the implicit assumption that contemporaries recognized the start and end of the hyperinflation ahead of time, whereas Cagan identified the dates in retrospect (the beginning date is particularly arbitrary). Specification (iii) which uses Hamilton's model of unobservable switching regimes to address this weakness. Agents (and the econometrician) know that some observations are generated by a high variance random walk and others by a lower variance AR(1) process, but they cannot directly observe the process which generates a given month's inflation. Instead they form ex ante and ex post estimates of the probability that each observation was generated by one or the other. I use estimates of each process (Table 5)[18] to calculate ex ante probabilities for each month and form one-step-ahead inflation forecasts. Resulting measures of forecast inflation (g_t), inflation shocks (v_t), and the conditional variance of inflation about g_t (σ^2_{vt}) for Germany appear in Table 6.[19]

The alternative series for trends and shocks are similar but there are important differences, particularly at the very end of the hyperinflation. In November 1923 Germany initiated a currency reform and revision of fiscal policies that succeeded in ending the hyperinflation. Inflation dropped from over 5 (in logarithmic terms) in November to near 0 in December.

Specification (ii) implicitly assumes that contemporaries knew the reform would succeed (g = .06 for December). Specifications (i) and (iii) assume that the reform into was not even taken into account (g = 3.53 and 4.33, respectively). Therefore (i) and (iii) show large negative demand shocks for December 1923.

While (i) is primitive, (iii) is appealing. Uncertainty about the onset of hyperinflation and its possible return after a reform (inherent in a switching regimes model) add realism. Assuming agents ignored November's reform as they forecast December's inflation seems unrealistic. Rather than devising another, more complicated specification, I take a simpler approach. Below I include aditional estimates for a sample which omits December 1923. Thus, there are six results for each relationship: one for each specification and each sample.

C. Estimates of the Dispersion Relationship.

Estimates of the modified dispersion relationship (A') for Germany appear in Table 7. Austria and Hungary results follow in Tables 8 and 9. Row (1) of each table displays the results using the standard quadratic form with actual inflation. Rows (2) - (4) contain analogous results using the three alternative measures of trend inflation rather than actual inflation. For Germany the fit is poorer for g than for π but the relationship is quite significant in each case. As expected, the removal of the first post-hyperinflation month improves the fit for trend inflation based on (i) and (iii). Rows (5) - (8) contain estimates of the same relationship using the flexible form suggested by the BMR model. The fit improves noticeably and ϕ is significantly below 2 in each case for Germany.[20]

Remaining estimates in Table 7 assess the Hercovitz effect. The change in shocks replaces trend inflation. There is evidence of this effect for (ii). The shock term has no explanatory power in the other specifications.[21] Estimates for Austria and Hungary in Tables 8 and 9 may be summarized: (a) π has more explanatory power than g or shocks in every case. (b) The flexible form (A') provides a superior fit in every case. (c) Omitting the first post-reform observation has no effect if actual inflation is used but improves the performance of some trend inflation specifications (this omission has no noticeable effect on any Hungarian result, hence these results are not shown here or for (B) below). (d) The Hercovitz term sometimes emerges significantly[22], but tends to provide a poorer fitthan trend inflation and often has no explanatory power.

D. Tests Based on Individual Relative Inflation Series

Estimates of (B) for full German sample and the sample with Dec. 1923 omitted appear in Tables 10 and 11. The format is that of Table 3, which was based on synthetic BMR data. Resulting statistics for hypothesis testing appear in Table 12. Row (1) of each table shows benchmark estimates of individual variances assuming each d_{jt} is white noise ($\rho_j = 0$) with fixed variance, σ^2_j (= S_j). Estimates allowing variances to rise with trend inflation ($s \neq 0$) appear in lines (2i, 2ii, 2iii) using g series from specifications (i), (ii), and (iii), respectively. The parameter c_s is significantly positive (see Table 12) causing individual variances to rise more than tenfold as g varies from 0 ($\sigma^2_j = S_j$) to the maximum for the sample ($\sigma^2_j = \sigma^2_{max}$).

Estimates in (3i, 3ii, 3iii) and (4i, 4ii, 4iii) allow autocorrelation ($\rho_j \neq 0$). ρ_{jt} varies across goods in each case but is fixed over time (r = 0) in (3i, 3ii, 3iii) where the R_j (= ρ_j) are near 0 and collectively insignificant.[23] In (4i, 4ii, 4iii) ρ_j depends on g ($\rho_{jt} = (R_j - rg_t)/(1 + rg_t)$). In this specification r is highly significant and the ρ_j become collectively significant.[24] Estimates of ρ_j at g = 0 **(=** R_j) are mostly positive and highly dispersed. As g approaches its maximum the ρ_j (= ρ_{max}) cluster tightly about a value near -1 (means -.87 to -.90 and standard deviations $\leq$ .02). This mimics the pattern found in the synthetic BMR data.

Estimates in (5i, 5ii, 5iii) allow a Hercovitz effect. The term $b_j(v_t-v_{t-1})$ captures systematic differences in the effect of aggregate shocks on individual inflation. The presence of this term cannot provide a decisive test since it can emerge in both models. Nevertheless, the results are illuminating. The b_j are jointly significant for each specification, but the size and significance are sensitive to the scheme for decomposing inflation. A useful calculation of size is the mean of $(b_j)^2$. This measures the effect of a unit change in $(v_t-v_{t-1})^2$ on the dispersion of inflation. Values from Table 10 yield .0025, .0169, and .0016 for (i), (ii), (iii), respectively - somewhat smaller than .0028, .0321, and .0025, the coefficients on $(v_t-v_{t-1})^2$ in Table 7. Thus, the effect of shocks on dispersion found in (B) is smaller than those found in (A) and is sensitive to the measure of shock. In contrast, the effect of trend inflation on relative inflation is insensitive to the decomposition of shocks and trends.

Overall, estimates of (B) for the German data are consistent with a menu-cost mechanism operating in the data. Estimates of (B) were also derived for Austria and Hungary. For brevity, only the statistics useful to hypothesis testing appear in Tables 13 and 14. Inflation induced heteroskedasticity is pervasive (s is always significant). Inflation induced auto-correlation (r >0) is the rule (specification (ii), full Austrian data is the notable exception). The Hercovitz effect (bj ≠ 0) is again sensitive to shock measures and sample. It emerges significantly about half the time (Austrian specification (ii) and all Hungary B results).

<u>E. Estimates Based on Hercovitz' Money Shocks</u>

Sensitivity of the Hercovitz effect to the measurement of shocks suggests that his original money shocks be examined as an alternative. Hercovitz calculated shocks for Nov 1920 - Jul 1923, thus omitting the peak hyperinflation and the post-reform period.[25] Rather than trying to calculate money shocks for the full sample, I estimate (A) and (B) for the shorter sample using both money shocks and the alternative shocks used above. Tables 15 and 16 contain estimates for (A) and (B), respectively. Row (1) of Table 15 reports a reproduction of the Hercovitz base regression with Germany B data.[26] In (2), (3), and (4), my inflation shocks provide no explanatory power. Relationships using trend inflation in (5), (6), and (7) are significant but yield a poorer fit than $(m_t\text{-}m_{t-1})^2$. However, π^2_t provides the best fit as a single regressor (8) and $(m_t\text{-}m_{t-1})^2$ adds no explanatory power when π^2_t is included (9).

The greater explanatory power of $(m_t\text{-}m_{t-1})^2$ compared to the various $(v_t\text{-}v_{t-1})^2$ suggests that Hercovitz devised a better decomposition of shocks and trends. The superiority of π^2_t as a regressor suggests that the data were not generated by a Hercovitz mechanism, rather spurious correlation has indicated one in a short sample. Such fragile results are common to this literature and reinforce the notion that estimates of (A) are unlikely to be conclusive. For a deeper look at the effect of shocks, Table 16 provides estimates of (B), where m_t replaces my v_t . As before, idiosyncratic values for ρ_j, σ^2_j, and b_j appear in each equation. Aggregate shock effects are omitted ($b_j = 0$) in (1). Rows (2) - (5) report the results of using money shocks, m_t, and my three versions of v_t as alternative sources of shock effects.

Money shocks have little effect on relative inflation. Improvement in likelihood from (1) to (2) is insignificant relative to threshold levels for $\chi^2(37)$. The cross-equation restriction on (A) and (B) is also violated. The coefficient on $(m_t\text{-}m_{t-1})^2$ in Table 15 should equal the mean of the $(b_j)^2$ in Table 16. It is 60 times larger: 12.57 vs .21 (= $(-.11)^2 + (.44)^2$). Money shocks correlate with dispersion but do not display the <u>systematic</u> effects on individual inflation which is the source of correlation in misperceptions models. Paradoxically, each of my demand shocks display significant effects in the individual equations. This is further evidence that estimates of the dispersion relationship are uninformative.

IV. Conclusion

I have derived alternative relationships which distinguish between the effects of trend inflation on relative inflation rates predicted by misperceptions and menu cost models of price setting. They have been estimated using individual price data from the European hyperinflations of the 1920's. Strong evidence of menu cost patterns is found even in this turbulent period. The evidence for misperceptions models reported by Hercovitz seems weak when individual inflation rates are examined rather than inflation dispersion.

Of broader importance, traditional methods of testing these models using a dispersion relationship are found to be less useful for two reasons. First, menu cost models as well as misperceptions models can generate shock effects. Second, dispersion contains less information than the individual prices upon which it is based. By exploiting this information, appropriate estimation of individual series proves more useful. Given the often contradictory results of traditional methods, this promises to be a more fruitful tool for decisive testing using modern data as well.

FOOTNOTES

* I thank Steven Webb for 1920-21 German data, and especially, Robert Anderson and Gail Makinen for uncovering the remaining data and for inspiring and aiding this project.

1. This approach is in the spirit of Ball and Mankiw (1995) which, in a very different way, assesses the ability of a menu cost model to account for relative price changes. This paper also examines the ability of a model to "...explain facts it was not designed to explain." (pg. 191)

2. Blejer (1981), Bomberger and Makinen (1993), Danziger (1987), Debelle and Lamont (1997), Domberger (1987), Fischer (1981), Graham (1930), Grier and Perry (1996), Lach and Tsiddon (1992), Parks (1978), Parsley (1996), Reinsdorf (1994), Van Hoomisen (1988), and Vining and Elwertowski (1976) among others.

3. For instance both Domberger and Reinsdorf examine the dispersion of prices within markets while Debelle and Lamont examine the dispersion across cities.

4. Blejer, Grier and Perry, Lach and Tsiddon, and Parks use this approach. Hartman (1991) claims an additional weakness: some functional forms for dispersion equations are guaranteed to yield a positive relationship, and are, thus, uninformative regarding the underlying model.

5. This is equation (2), page 331 in Hercovitz.

6. The supply shocks which appear in equation (1), page 331 of Hercovitz are omitted.

7. These expressions correspond to (12) and (13) p. 334 in Hercovitz.

8. Hercovitz skips a step in his derivation, hence no expression corresponding to (5) occurs. However, (6) corresponds to the expression directly below (15) p.335 in Hercovitz. To derive (6) one replaces of the realized average of $(\epsilon_{jt})^2$ with its expected value σ^2_ϵ

9. I concentrate on this model because its structure parallels the misperceptions model and permits relatively easy simulations. It has no relative price predictions which are qualitatively different from previous models such as Ceccheti (1985), Mussa (1981), Rotemberg (1983), and Sheshinski and Weiss (1977). Accordingly, my analysis breaks no new theoretical ground.

10. This expression can be arrived at by substituting (A6) and (A2) into (A1) on page 61 of BMR. The trend growth in supply is suppressed as before and notation is changed so as to conform to my treatment of the Hercovitz model.

11. This corresponds to (10) in BMR p. 23. BMR do not derive (7').

12. Hercovitz obtains clean results by assuming the information lag exactly equals the sampling period. Otherwise, the distribution of shocks within a month would matter as it does here.

13. Both Blejer and Lach and Tsiddon impose a linear relationship. Danziger has terms in both $\pi^{4/3}$ and π^2. Van Hoomisen combines linear and quadratic terms. It is somewhat surprising that no-one has chosen (A') which allows many of these forms as special cases.

14. Results for the synthetic data and the European data below are insensitive to the choice of alternative functional forms. All parameters are estimated by a grid search procedure.

15. The data for 1920 are from **Wirtshcaft und Statistich**. **Statistiches Reichsamt 1912/22, 1923, 1924** contain the later data. I do not attempt to extend Hercovitz's money shocks over the same period. Instead inflation is divided into expected and unexpected portions below.

16. Data are from **Statistiches Handbuch Fur Die Republik Osterreich** 1921, 1922, 1923.

17. Data are from **Annuaire Statistiques Hungroise**. Both series were selected from a larger group of series (38 and 77, respectively) on the basis of missing observations. A third Hungarian data set runs from Jan 1915 through Dec 1925. Its length makes it interesting but its coverage of only 12 goods reduces its usefulness. All series will be provided on request.

18. Switches between the two regimes are described by a Markov process with fixed transition matrix. Estimation includes values for p_{aa}, the probability of the AR(1) continuing for another period (no switching) and p_{bb}, the probability of continuing in the random walk regime. The probabilities of switching regimes are $p_{ab} = 1 - p_{aa}$, and $p_{ba} = 1 - p_{bb}$. See Hamilton (1989)

19. Although each process is assumed homoskedastic as before, the ex ante probabilities that the coming month's inflation will be drawn from (a) or (b) vary from month to month. Hence, the conditional variance of future inflation varies over time.

20. Since the insignificant constant term is omitted leaving two fitted parameters, this form does not increase the degrees of freedom and improvement of fit is not inevitable. This result ($0 < \phi < 2$) is consistent with the notion that higher g increases the frequency of price changes but less than proportionally. It is also consistent with the results reporteded in the next section.

21. If v is used instead of Δv, the adjusted R^2 drops below .2 for all specifications (not shown). If both g and Δv are included for specification (ii) the result is:

$$\underset{}{\gamma^2_t} = \underset{(.005)}{.038} + \underset{(.0040)}{.0125}(g_t)^2 - \underset{(.0140)}{.0098}(\Delta v_t)^2 \qquad \bar{R}^2 = .58$$

22. In particular this term is significant for Austrian specifications (ii) and (iii) and remains significant if trend inflation is added to the specification (not shown).

23. Twice the change in likelihood from (2) to (3) is distributed as a $\chi^2(37)$ due to relaxation of the restriction that each of 37 $R_j = 0$. This falls short of the 1% threshold of 60.0 (Table 12).

24. Twice the change in likelihood between (3) and (4) is distributed as a $\chi^2(1)$. Row (2) constrains r = all $R_j = 0$. Hence a $\chi^2(38)$ is relevant for hypothesis testing ((2) to (4)). The respective 1% threshold values of 6.6 and 61.3 are easily achieved in each case (Table 12).

25. "The vertiginous monetary expansion initiated in August 1923 differentiates the last phase of hyperinflation, and, thus, is not included in the sample". Hercovitz p.330

26. The LHS variable is the 37 item variance from Germany B. I use the longest sample for which m_t is available, Nov 1920 - July 1923. Despite these slight data differences, the estimates are virtually identical to Hercovitz's (p. 347).