

Evolving Inflation Dynamics and the New Keynesian Phillips Curve

Andreas Hornstein

In most industrialized economies, periods of above average inflation tend to be associated with above average economic activity, for example, as measured by a relatively low unemployment rate. This statistical relationship, known as the Phillips curve, is sometimes invoked when economic commentators suggest that monetary policy should not try to suppress signs of inflation. But this interpretation of the Phillips curve implicitly assumes that the statistical relationship is structural, that is, the relationship will not break down during periods of persistently high inflation. Starting in the mid-1960s, Friedman and Phelps argued that the Phillips curve is indeed not structural and the experience of the United States and other countries with high inflation and low GDP growth in the late 1960s and 1970s has subsequently borne out their predictions.

Various theories have been proposed to explain the Phillips curve and most of these theories agree that there is no significant long-term tradeoff between inflation and the level of economic activity. One theory that provides a structural interpretation of the short-term inflation-unemployment relationship, and that has become quite popular over the last ten years among central bank economists is based on explicit models of nominal price rigidity. The most well-known example of this theory is the New Keynesian Phillips Curve (NKPC).

In this article, I evaluate how well a structural NKPC can account for the changing nature of inflation in the United States from the 1950s to today. First, I document that changes in average inflation have been associated with

■ I would like to thank Chris Herrington, Thomas Lubik, Yash Mehra, and Alex Wolman for helpful comments, and Kevin Bryan for excellent research assistance. Any opinions expressed in this article are my own and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: andreas.hornstein@rich.frb.org.

changes in the dynamics of inflation as measured by inflation persistence and the co-movement of inflation with measures of real activity that the NKPC predicts are relevant for inflation. Then I argue that the NKPC with fixed structural parameters cannot account for these changes in the inflation process. I conclude that the NKPC does not provide a complete structural interpretation of the Phillips curve. This is troublesome since the changed inflation dynamics are related to changes in average inflation, which are presumably driven by systematic monetary policy. But if the NKPC is not invariant to systematic changes of monetary policy, then its use for monetary policy is rather limited.

In models with nominal rigidities, sticky-price models for short, monopolistically competitive firms set their prices as markups over their marginal cost. Since these firms are limited in their ability to adjust their nominal prices, future inflation tends to induce undesired changes in their relative prices. When firms have the opportunity to adjust their prices they will, therefore, set their prices contingent on averages of expected future marginal cost and inflation. The implied relationship between inflation and economic activity is potentially quite complicated, but for a class of models one can show that to a first-order approximation current inflation is a function of current marginal cost and expected future inflation, the so-called NKPC. The coefficients in this NKPC are interpreted as structural in the sense that they are likely to be independent of monetary policy.

In the U.S. economy, inflation tends to be very persistent, in particular, it tends to be at least as persistent as is marginal cost. At the same time, inflation is not that strongly correlated with marginal cost. This observation appears to be inconsistent with the standard NKPC since here inflation is essentially driven by marginal cost, and inflation is, at most, as persistent as marginal cost. But if inflation is as persistent as is marginal cost then the model also predicts a strong positive correlation between inflation and marginal cost. One can potentially account for this observation through the use of a hybrid NKPC which makes current inflation not only a function of expected future inflation, but also of past inflation as in standard statistical Phillips curves. With a strong enough backward-looking element, inflation persistence then need not depend on the contributions from marginal cost alone.

Another feature of U.S. inflation is that average inflation has always been positive, and it has varied widely: periods of low inflation, such as the 1950s and 1960s, were followed by a period of very high inflation in the 1970s, and then low inflation again since the mid-1980s. Cogley and Sbordone (2005, 2006) point out that the NKPC relates inflation and marginal cost defined in terms of their deviations from their respective trends. In particular, the standard NKPC defines trend inflation to be zero. Given the variations in average U.S. inflation, Cogley and Sbordone (2005, 2006) then argue that accounting for variations in trend inflation will make deviations of inflation from trend less persistent. Furthermore, as Ascari (2004) shows, the first-order

approximation of the NKPC needs to be modified when the approximation is taken at a positive inflation rate.

I build on the insight of Cogley and Sbordone (2005, 2006) and study the implications of a time-varying trend inflation rate for the autocorrelation and cross-correlation structure of inflation and marginal cost. In this I extend the work of Fuhrer (2006) who argues that the hybrid NKPC can account for inflations' autocorrelation structure only through a substantial backward-looking element. In this article, I argue that a hybrid NKPC, modified for changes in trend inflation, cannot account for changes in the autocorrelation and cross-correlation structure of inflation and marginal cost in the United States.

The article is organized as follows. Section 1 describes the dynamic properties of inflation and marginal cost in the baseline NKPC and the U.S. economy. Section 2 describes and calibrates the hybrid NKPC, and it compares the autocorrelation and cross-correlation structure of inflation and marginal cost in the model with that of the 1955–2005 U.S. economy. Section 3 characterizes the inflation dynamics in the NKPC modified to account for nonzero trend inflation. I then study if the changes of inflation dynamics, associated with changes in trend inflation comparable to the transition into and out of the high inflation period of the 1970s, are consistent with the changing nature of inflation dynamics in the U.S. economy for that period.

1. INFLATION AND MARGINAL COST IN THE NKPC

Inflation in the baseline NKPC is determined by expectations about future inflation and a measure of current economic activity. There are two fundamental differences between the NKPC and more traditional specifications of the Phillips curve. First, traditional Phillips curves are backward looking and relate current inflation to lagged inflation rates. Second, the measure of real activity in the NKPC is based on a measure of how costly it is to produce goods, whereas traditional Phillips curves use the unemployment rate as a measure of real activity. More formally, the baseline NKPC is

$$\hat{\pi}_t = \kappa_0 \hat{s}_t + \beta E_t [\hat{\pi}_{t+1}] + u_t, \quad (1)$$

where $\hat{\pi}_t$ denotes the inflation rate, \hat{s}_t denotes real marginal cost, $E_t \hat{\pi}_{t+1}$ denotes the expected value of next period's inflation rate conditional on current information, u_t is a shock to the NKPC, β is a discount factor, $0 < \beta < 1$, and κ_0 is a function of structural parameters described below. The baseline NKPC is derived as the local approximation of equilibrium relationships for a particular model of the economy, the Calvo (1983) model of price adjustment.

For the Calvo model one assumes that all firms are essentially identical, that is, they face the same demand curves and cost functions. The firms are monopolistically competitive price setters, but can adjust their nominal prices

only infrequently. In particular, whether a firm can adjust its price is random, and the probability of price adjustment is constant. Random price adjustment introduces ex post heterogeneity among firms, since with nonzero inflation a firm's relative price will depend on how long ago the firm last adjusted its price. Since firms are monopolistically competitive they set their nominal (and relative) price as a markup over their real marginal cost, and since firms can adjust their price only infrequently they set their price conditional on expected future inflation and marginal cost.

The NKPC is a linear approximation to the optimal price-setting behavior of the firms in the Calvo model. Furthermore, the approximation is local to a state that exhibits a zero-average inflation rate. The inflation rate $\hat{\pi}_t$ should be interpreted as the log-deviation of the gross inflation rate from one, that is, the net-inflation rate, and real marginal cost \hat{s}_t should be interpreted as the log-deviation from its long-run mean. For a derivation of the NKPC, see Woodford (2003).¹ The optimal pricing decisions of firms with Calvo-type nominal price adjustment are reflected in the parameter κ_0 of the NKPC,

$$\kappa_0 = \frac{1 - \alpha}{\alpha} (1 - \alpha\beta), \quad (2)$$

where α is the probability that a firm cannot adjust its nominal price, $0 \leq \alpha < 1$.

The shock to the NKPC is usually not derived as part of the linear approximation to the optimal price-setting behavior of firms. Most of the time the shock is simply "tacked on" to the NKPC, although it can be interpreted as a random disturbance to the firms' static markup. Given the absence of serious microfoundations of the cost shock one would not want the shock to play an independent role in contributing to the persistence of inflation. We, therefore, assume that the shock to the NKPC is i.i.d. with mean zero.²

Persistence of Inflation in the NKPC

The NKPC represents a partial equilibrium relationship within a more comprehensive model of the economy. Thus, inflation and marginal cost will be simultaneously determined as part of a more complete description of the economy. Conditional on the equilibrium process for marginal cost we can, however, solve equation (1) forward by repeatedly substituting for future inflation and obtain the current inflation rate as the discounted expected value

¹ The NKPC approximated at the zero inflation rate is also a special case of the NKPC approximated at a positive inflation rate. For a derivation of the latter, see Ascari (2004), Cogley and Sbordone (2005, 2006), or Hornstein (2007).

² The shock to the NKPC is often called a "cost-push" shock, but this terminology can be confusing since the shock is introduced independently of marginal cost.

of future marginal cost

$$\hat{\pi}_t = \kappa_0 \sum_{j=0}^{\infty} \beta^j E_t [\hat{s}_{t+j}] + u_t. \quad (3)$$

The behavior of the inflation rate, in particular its persistence, is therefore closely related to the behavior of marginal cost. To get an idea of what this means for the joint behavior of inflation and marginal cost, assume that equilibrium marginal cost follows a first-order autoregressive process [AR(1)],

$$\hat{s}_t = \delta \hat{s}_{t-1} + \varepsilon_t, \quad (4)$$

with positive serial correlation, $0 < \delta < 1$, and ε_t is an i.i.d. mean zero shock with variance σ_ε^2 . This AR(1) specification is a useful first approximation of the behavior of marginal cost since, as we will see below, marginal cost is a highly persistent process. For such an AR(1) process the conditional expectation of marginal cost j -periods-ahead is simply

$$E_t [\hat{s}_{t+j}] = E_t [\delta \hat{s}_{t+j-1} + \varepsilon_{t+j}] = \delta E_t [\hat{s}_{t+j-1}] = \dots = \delta^j \hat{s}_t. \quad (5)$$

Substituting for the expected future marginal cost in (3), we get

$$\hat{\pi}_t = \kappa_0 \sum_{j=0}^{\infty} \beta^j \delta^j \hat{s}_t + u_t = \frac{\kappa_0}{1 - \beta\delta} \hat{s}_t + u_t = a_0 \hat{s}_t + u_t. \quad (6)$$

This is a reduced form relationship between current inflation and marginal cost. The relationship is in reduced form since it incorporates the presumed equilibrium law of motion for marginal cost, which is reflected in the fact that the coefficient on marginal cost, a_0 , depends on the law of motion for marginal cost. If the law of motion for marginal cost changes, then the relation between inflation and marginal cost will change.

Given the assumed law of motion for marginal cost, inflation is positively correlated with marginal cost and is, at most, as persistent as is marginal cost. The second moments of the marginal cost process are

$$E [\hat{s}_t \hat{s}_{t-k}] = \delta^k \frac{\sigma_\varepsilon^2}{1 - \delta^2} = \delta^k \sigma_s^2, \quad (7)$$

where σ_s^2 is the variance of marginal cost. The implied second moments of the inflation rate and the cross-products of inflation and marginal cost are

$$E [\hat{\pi}_t \hat{\pi}_{t-k}] = a_0^2 E [\hat{s}_t \hat{s}_{t-k}] + I_{[k=0]} \sigma_u^2 = \delta^k (a_0 \sigma_s)^2 + I_{[k=0]} \sigma_u^2, \quad (8)$$

$$E [\hat{\pi}_t \hat{s}_{t+k}] = a_0 E [\hat{s}_t \hat{s}_{t+k}] = \delta^k a_0 \sigma_s^2, \quad (9)$$

where $I_{[\cdot]}$ denotes the indicator function. The autocorrelation coefficients for inflation and the cross-correlations of inflation with marginal cost are

$$\text{Corr}(\hat{\pi}_t, \hat{\pi}_{t-k}) = \delta^k \frac{a_0^2}{a_0^2 + \sigma_u^2 / \sigma_s^2}, \text{ and} \quad (10)$$

$$\text{Corr}(\hat{\pi}_t, \hat{\delta}_{t+k}) = \delta^k \frac{a_0}{[a_0^2 + \sigma_u^2 / \sigma_s^2]^{1/2}}. \quad (11)$$

As we can see, the autocorrelation coefficients for inflation are simply scaled versions of the autocorrelation coefficients for marginal cost, and the scale parameter depends on the relative volatility of the shocks to the NKPC and marginal cost. If there are no shocks to the NKPC, $\sigma_u = 0$, then inflation is an AR(1) process with persistence parameter δ , and it is perfectly correlated with marginal cost. If, however, there are shocks to the NKPC, $\sigma_u > 0$, then inflation and marginal cost are imperfectly correlated and inflation is less persistent than is marginal cost.

Inflation and Marginal Cost in the U.S. Economy


In order to make the NKPC operational, we need measures of the inflation rate and marginal cost. For the inflation rate we will use the rate of change of the GDP deflator.³ We measure aggregate marginal cost through the wage income share in the private nonfarm business sector. This choice can be motivated as follows. Suppose that all firms use the same production technology with labor as the only input. In particular, assume that the production function is Cobb-Douglas, $y = zn^\omega$, with constant input elasticity ω . Then the nominal marginal cost is the nominal wage divided by the marginal product of labor

$$S_t = \frac{W_t}{MPL_t} = \frac{W_t}{\omega y_t / n_t}, \quad (12)$$

and nominal marginal cost is proportional to nominal average cost. We use the unit labor cost index for the private nonfarm business sector as our measure of average labor cost. Deflating nominal average cost with the price index of the private nonfarm business sector yields real average labor cost, that is, the labor income share. The log deviation of real marginal cost from its mean is

³ This is the most commonly used price index in the implementation of the NKPC. Other price indices used include the price index of the private nonfarm business sector or the price index for Personal Consumption Expenditures (PCE), the consumption component of the GDP deflator. Although the choice of price deflator affects the results described below, the differences are not dramatic, e.g., Galí and Gertler (1999). We should also note that only consumption based indices, such as the PCE index, are commonly mentioned by central banks in their communications on monetary policy.

Figure 1 Inflation and Marginal Cost in the United States, 1955–2005



Notes: Inflation and marginal cost are defined in the Appendix. The solid line in Panel A represents the inflation rate and its sample mean, and the dashed line represents marginal cost and its sample mean. In Panel B, the circles (diamonds) denote the sample autocorrelations for inflation (marginal cost). In Panel C, the squares denote the cross-correlations of inflation and marginal cost. In Panels B and C, the boxes denote the 5-percentile to 95-percentile range of the statistic calculated from 1,000 bootstraps of the data.

then equal to the log-deviation of the labor income share from its mean

$$\hat{s}_t = \frac{\widehat{W_t n_t}}{P_t y_t}. \tag{13}$$

The detailed source information for our data is listed in the Appendix.

In Figure 1.A, we graph the quarterly inflation rate and marginal cost for the time period 1955Q1 to 2005Q4. Inflation varies widely over this time period, from about 1 percent at the low end in the early 1960s, to more than 10 percent in the 1970s, with a 3 1/2 percent average inflation rate, Table 1, column 1. Inflation and marginal cost are both highly persistent, the first-order autocorrelation coefficient is about 0.9 for both variables, Figure 1.B. To the

Table 1 Inflation and Marginal Cost

Sample	$\bar{\pi}$	$\sigma_{\hat{\pi}}$	\bar{s}	$\sigma_{\hat{s}}$	$\bar{\delta}_{\hat{\pi}}$	$\bar{\delta}_{\hat{s}}$	Corr ($\hat{\pi}, \hat{s}$)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1955Q1–2005Q4	3.6	2.4	0.013	0.021	0.94	0.93	0.33
					[0.88,0.99]	[0.89,0.98]	[0.23,0.43]
1955Q1–1969Q4	2.5	1.4	0.023	0.018	0.97	0.89	-0.12
					[0.83,0.98]	[0.79,1.00]	[-0.30,0.05]
1970Q1–1983Q4	6.5	2.2	0.024	0.016	0.80	0.72	0.29
					[0.62,0.98]	[0.56,0.88]	[0.10,0.46]
1984Q1–1991Q4	3.2	0.9	0.011	0.007	0.60	0.73	0.10
					[0.20,1.03]	[0.51,0.95]	[0.09,0.34]
1992Q1–2005Q4	2.1	0.7	-0.009	0.018	0.76	0.92	-0.06
					[0.50,1.02]	[0.81,1.02]	[-0.32,0.22]

Notes: Columns (1) and (2) contain the average annualized inflation rate, $\bar{\pi}$, and its standard deviation, $\sigma_{\hat{\pi}}$. Columns (3) and (4) contain the average values and standard deviation of marginal cost, \bar{s} and $\sigma_{\hat{s}}$. Marginal cost is in log deviations from its normalized 1992 value. Columns (5) and (6) contain the sum of the autocorrelation coefficients of a univariate OLS regression with four lags for inflation respectively marginal cost, $\bar{\delta}_{\hat{\pi}}$ and $\bar{\delta}_{\hat{s}}$. Column (7) contains the contemporaneous correlation coefficient between inflation and marginal cost. For the sum of autocorrelation coefficients and the correlation coefficient, columns (5), (6), and (7), we list the 5th and 95th percentile of the respective bootstrapped statistic with 1,000 replications in brackets.

extent that the autocorrelation coefficients of inflation do not decline as fast as the ones for marginal cost, inflation appears to be somewhat more persistent than marginal cost. Levin and Piger (2002) use an alternative measure of persistence in their analysis of inflation in the United States, namely the sum of lagged coefficients in a univariate regression of a variable on its own lags. This measure also yields estimates of significant and similar persistence for inflation and marginal cost, Table 1, columns 5 and 6. Inflation and marginal cost tend to move together. The cross-correlations between inflation and marginal cost are positive, 0.33 contemporaneously and above 0.2 at all four lags and leads, Table 1, column 7, and Figure 1.C. Although the co-movement between inflation and marginal cost is significant, it is not particularly strong.⁴

As we have shown previously, in the basic NKPC model, persistence of inflation and marginal cost, and co-movement of inflation with marginal cost go together. The observation that inflation is about as persistent as marginal cost, but only weakly correlated with marginal cost then seems to be inconsistent with the basic NKPC. We now study if two modifications of the basic

⁴The positive cross-correlation coefficients are significant for all four lags and leads. Based on 1,000 bootstraps the 5-percentile to 95-percentile ranges of the coefficients do not include zero, Figure 1.C.

NKPC can resolve this apparent inconsistency. The first approach is to make the NKPC more like a standard Phillips curve by directly introducing lagged inflation. The second approach argues that some of the observed inflation persistence is spurious. Extended apparent deviations of the inflation rate from the sample average inflation rate, for example in the 1970s, are interpreted as sub-sample changes in the mean inflation rate. This approach then suggests that the NKPC has to be modified to take into account changes in trend inflation. We will discuss these two approaches in the following sections.

2. A HYBRID NKPC

The importance of marginal cost for inflation persistence will be reduced if there is a source of persistence that is inherent to the inflation process itself. Two popular approaches that introduce such a backward-looking element of price determination into the NKPC are “rule-of-thumb” behavior and indexation. For the first approach, one assumes that a fraction ρ of the price-setting firms do not choose their prices optimally, rather they index their prices to past inflation. For the second approach one assumes that firms who do not have the option to adjust their price optimally simply index their price to a fraction ρ of past inflation.⁵ The two approaches are essentially equivalent and for the second case the NKPC becomes

$$(1 - \rho L) \hat{\pi}_t = \beta E_t [(1 - \rho L) \hat{\pi}_{t+1}] + \kappa_0 \hat{s}_t + u_t, \quad (14)$$

where L is the lag operator, $L^j x_t = x_{t-j}$ for any integer j .

This modification of the NKPC is also called a hybrid NKPC since current inflation not only depends on expected inflation as in the baseline NKPC, but it also depends on past inflation as in a traditional Phillips curve. The dependence on lagged inflation introduced through backward-looking price determination is called “intrinsic” persistence since it is an exogenous part of the model structure. Complementary to intrinsic persistence is “extrinsic” inflation persistence which comes through the marginal cost process that drives inflation. To the extent that monetary policy affects marginal cost, it influences extrinsic inflation persistence.

Note that the hybrid NKPC, equation (14), is of the same form as the basic NKPC, equation (1), except for the linear transformation of inflation, $\tilde{\pi}_t = \hat{\pi}_t - \rho \hat{\pi}_{t-1}$, replacing the actual inflation rate. Forward-solving equation (14), assuming again that marginal cost follows an AR(1) process, as in equation (4), then yields the following expression for $\tilde{\pi}_t$:

$$\hat{\pi}_t - \rho \hat{\pi}_{t-1} = \frac{\kappa_0}{1 - \beta \delta} \hat{s}_t + u_t = a_0 \hat{s}_t + u_t. \quad (15)$$

⁵ “Rule-of-thumb” behavior was introduced by Galí and Gertler (1999); inflation indexation has been used by Christiano, Eichenbaum, and Evans (2005).

For this specification, inflation can be more persistent than marginal cost because current inflation is indexed to past inflation.

The autocorrelation coefficients for the linear transformation of inflation, $\tilde{\pi}_t$, are the same as defined in equation (10), but the autocorrelation coefficients for the inflation rate itself are now more complicated functions of the persistence of marginal cost and the intrinsic inflation persistence. In Hornstein (2007), I derive the autocorrelation and cross-correlation coefficients for inflation and marginal cost,

$$\text{Corr}(\hat{\pi}_t, \hat{\pi}_{t-k}) = \frac{(\sigma_u/\sigma_s)^2 A(k; \rho) + a_0^2 B(k; \rho, \delta)}{(\sigma_u/\sigma_s)^2 A(0; \rho) + a_0^2 B(0; \rho, \delta)} \text{ and} \quad (16)$$

$$\text{Corr}(\hat{\pi}_t, \hat{s}_{t+k}) = \frac{a_0 C(k; \rho, \delta)}{[(\sigma_u/\sigma_s)^2 A(0; \rho) + a_0^2 B(0; \rho, \delta)]^{1/2}}, \quad (17)$$

where

$$\begin{aligned} A(k; \rho) &= \rho^k \frac{1}{1 - \rho^2}, \\ B(k; \rho, \delta) &= \left[\delta^k - \frac{\rho}{\delta} \frac{1 - \delta^2}{1 - \rho^2} \rho^k \right] \frac{1}{(1 - \rho/\delta)(1 - \rho\delta)}, \\ C(k; \rho, \delta) &= \delta^k \frac{1}{1 - \rho\delta} \text{ if } k \geq 0, \text{ and} \\ C(k; \rho, \delta) &= \left[\delta^{-k} - \rho^{-k} \frac{\rho}{\delta} \frac{1 - \delta^2}{1 - \rho\delta} \right] \frac{1}{1 - \rho/\delta} \text{ if } k < 0. \end{aligned}$$

Inflation Persistence in the Hybrid NKPC

Inflation persistence for the hybrid NKPC depends not only on the persistence of marginal cost and intrinsic inflation persistence, δ and ρ , but also on the relative volatility of the shocks to the NKPC and marginal cost, σ_u/σ_s , and the reduced form coefficient on marginal cost, a_0 . In order to evaluate the implications of the hybrid NKPC for inflation dynamics we, therefore, need estimates of the structural parameters of the NKPC and the relative standard deviation of the NKPC shock. In the following, I study the implications of two alternative calibrations. The first calibration is based on generalized method of moments (GMM) estimates of the structural parameters, α , β , and ρ , and an estimate of the relative volatility of the NKPC shocks that is implicit in the GMM estimates. This calibration has only limited success in matching the autocorrelation and cross-correlation properties of inflation and marginal cost. For the second calibration, I then set intrinsic persistence and the relative volatility of the NKPC shock to directly match the autocorrelation and cross-correlation properties of inflation and marginal cost.

Table 2 New Keynesian Phillips Curve Estimates, 1960 Q1–2005 Q4

	α	ρ	β	$\hat{\pi}_{t-1}$	$\hat{\pi}_{t+1}$	\hat{s}_t
(1)	0.901 (0.028)	0.164 (0.124)	0.990 (0.028)	0.141 (0.091)	0.851 (0.087)	0.010 (0.007)
(2)	0.897 (0.021)	0.469 (0.095)	0.944 (0.043)	0.325 (0.046)	0.654 (0.048)	0.012 (0.005)

Notes: This table reports estimates of the NKPC approximated at a zero inflation rate, equation (14). The first three columns contain estimates of the structural parameters: price non-adjustment probability, α , degree of inflation indexation, ρ , and time discount factor β . The next three columns contain the implied reduced form coefficients on marginal cost, and lagged and future inflation when the coefficient on current inflation is one. The first row represents estimates of the moment conditions from equation (14). The second row represents estimates of the moment conditions from equation (14) when the coefficient of contemporaneous inflation is normalized to one. The covariance matrix of errors is estimated with a 12 lag Newey-West procedure. Standard errors of the estimates are shown in parentheses.

Galí, Gertler, and López-Salido (2005) (hereafter referred to as GGLS) estimate the hybrid NKPC for U.S. data using GMM techniques.⁶ I replicate their analysis for the hybrid NKPC (14) using the data on inflation and marginal cost for the time period 1960–2005. The instrument set includes four lags of the inflation rate, and two lags each of marginal cost, nominal wage inflation, and the output gap.⁷ The results reported in Table 2 are not exactly the same as in GGLS, but they are broadly consistent with GGLS. The time discount factor, β , is estimated close to one, and the coefficient on marginal cost, $\kappa_0 = 0.01$, is smaller than for GGLS. The small coefficient on marginal cost translates to a relatively low price adjustment probability: only about 10 percent, $1 - \alpha$, of all prices are optimally adjusted in a quarter. Similar to GGLS the estimated degree of inflation indexation depends on the normalization of the GMM moment conditions. For the first specification, when equation (14) is estimated directly, we find a relatively low degree of indexation to past inflation, $\rho = 0.16$. For the second specification, when the coefficient on current inflation in equation (14) is normalized to one, we find significantly more indexation, $\rho = 0.47$.

We construct an estimate of the volatility of shocks to the NKPC in two steps. First, we regress current inflation $\hat{\pi}_t$ on the set of instrumental variables. The instrumental variables contain only lagged variables, that is, information

⁶ Other work that estimates the NKPC using the same or similar techniques includes Galí and Gertler (1999) and Sbordone (2002). See also the 2005 special issue of the *Journal of Monetary Economics* vol. 52 (6).

⁷ The data are described in detail in the Appendix.

Table 3 Calibration

	Parameter	Calibration	
		(1)	(2)
β	Time Discount Factor	0.99	0.99
α	Probability of No Price Adjustment	0.90	0.80
ρ	Price Indexation	0.45	0.86
σ_u/σ_s	Relative NKPC Shock Volatility	0.10	2.97
δ	Marginal Cost Persistence	0.90	0.90

available in the previous period. We then use this regression to obtain an estimate of the expected inflation rate conditional on available information, $E_t \hat{\pi}_{t+1}$, and substitute it together with the information on current inflation and marginal cost, and the estimated parameter values in equation (14), and solve for the shock to the NKPC, u_t . The calculated standard deviation of the shock is about 1/10 of the standard deviation of marginal cost.⁸

Based on the GMM estimates for the second specification of the moment conditions, I now choose a parameterization of the hybrid NKPC with some intrinsic inflation persistence, Table 3, column 1.⁹ For the persistence of marginal cost, I choose $\delta = 0.9$, which provides a reasonable approximation of the autocorrelation structure of marginal cost for the period 1955 to 2005.

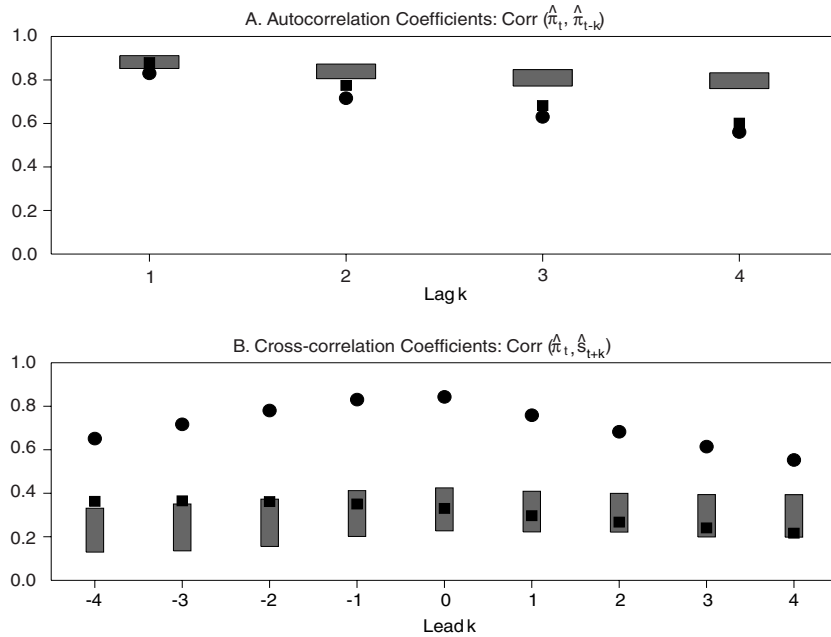
We can now characterize the inflation dynamics implied by the hybrid NKPC. The bullet points in Figure 2 display the first four autocorrelation coefficients of inflation and the cross-correlation coefficients of inflation with marginal cost implied by the calibrated model. Figure 2 also displays the bootstrapped 5th to 95th percentile ranges for the autocorrelation and cross-correlation coefficients of inflation and marginal cost for the U.S. economy from Figure 1.B and 1.C. As we can see, the model does not do too badly for the autocorrelation structure of inflation: the first-order autocorrelation coefficient of inflation is just outside the 5th to 95th percentile range, but then the autocorrelation coefficients are declining too fast relative to the data.¹⁰ The model does generate too much co-movement for inflation and marginal cost

⁸ Depending on the parameter estimates, $\sigma_u = 0.0019$ for specification one and $\sigma_u = 0.0025$ for specification two. For either specification the serial correlation of the shocks is quite low, the highest value is 0.2. Fuhrer (2006) argues for a higher relative volatility of the NKPC shock, about 3/10 of the volatility of marginal cost.

⁹ Choosing a lower value for indexation based on specification, one would generate less inflation persistence.

¹⁰ Fuhrer (2006) assumes a three times larger relative volatility of the NKPC shocks and, therefore, requires substantially more intrinsic persistence, that is, a higher ρ , in order to match inflation persistence.

Figure 2 Inflation Dynamics for the Hybrid NKPC



Notes: The circles (squares) denote autocorrelations and cross-correlations from calibration 1 (2) of the hybrid NKPC. The boxes denote the 5-percentile to 95-percentile range of the statistic calculated from 1,000 bootstraps of data.

relative to the data: the predicted contemporaneous correlation coefficient is about 0.8, well above the observed value of 0.3.

Given the failure of the GMM-based calibration to account for the autocorrelation and cross-correlation structure of inflation and marginal cost, I now consider an alternative calibration that exactly matches the first-order autocorrelation of inflation and the contemporaneous cross-correlation of inflation and marginal cost. As I pointed out above, the estimated price adjustment probability of 10 percent per quarter is quite low. Other work suggests higher price adjustment probabilities, about 20 percent per quarter, e.g., Galí and Gertler (1999), Eichenbaum and Fisher (2007), or Cogley and Sbordone (2006).¹¹ For the alternative calibration I, therefore, assume that $\alpha = 0.8$. Conditional

¹¹ The NKPC specification in equation (14) is based on constant firm-specific marginal cost. Eichenbaum and Fisher (2007) and Cogley and Sbordone (2006) consider the possibility of increasing firm-specific marginal cost. Adjusting their estimates for constant firm-specific marginal cost yields $\alpha = 0.8$.

on an unchanged time discount factor, β , this implies a coefficient on marginal cost, $\kappa_0 = 0.05$, which represents an upper bound of what has been estimated for hybrid NKPCs.

I now choose intrinsic persistence, ρ , and the relative volatility of the NKPC shock, σ_u/σ_s , to match the sample first-order autocorrelation coefficient of inflation, $Corr(\hat{\pi}_t, \hat{\pi}_{t-1}) = 0.88$, and the contemporaneous correlation of inflation and marginal cost, $Corr(\hat{\pi}_t, \hat{s}_t) = 0.33$. This procedure yields a very large value for inflation indexation, $\rho = 0.86$, which makes inflation persistence essentially independent of marginal cost. A very high relative volatility of the NKPC shock, $\sigma_u/\sigma_s = 2.97$, can then reduce the co-movement between inflation and marginal cost without affecting inflation persistence significantly. The implied parameter values of this calibration are summarized in the second column of Table 3.

The autocorrelation and cross-correlation structure of the alternative calibration is represented by the squares in Figure 2. With few exceptions the cross-correlations predicted by the alternative calibration stay in the 5th to 95th percentile ranges of the observed cross-correlations. The autocorrelation coefficients continue to decline at a rate that is faster than observed in the data.

3. THE CHANGING NATURE OF INFLATION

The behavior of inflation has changed markedly over time, Table 1, column (1). Inflation tended to be below the sample mean in the 1950s and 1960s, average inflation was about 2.5 percent, but inflation increased in the second half of the 1960s. In the 1970s, inflation increased even more, averaging 6.5 percent and reaching peaks of up to 12 percent. In the early 1980s, inflation came down fast, averaging 3.2 percent from 1984 to 1991. Finally, in the period since the early 1990s, inflation continued to decline, but otherwise remained relatively stable, averaging about 2 percent.¹²

Most observers attribute the changes in average inflation since the 1960s to changes in monetary policy, as represented by different chairmen of the monetary policy committee of the Federal Reserve System. We have the “Burns inflation” of the 1970s, the “Volker disinflation” of the early 1980s, and the “Greenspan period” with a further reduction and stabilization of inflation from the late 1980s to 2005. Interestingly enough, these substantial changes in the mean inflation rate were not associated with comparable changes in mean marginal cost: average marginal cost differs by at most 3 percent across the sub-samples, Table 1, column 3.

¹² I choose 1970 as the starting point of the high inflation era since mean inflation before 1970 is relatively close to the sample mean. The year 1984 is usually chosen as representing a definite break with the high inflation regime of the 1970s, e.g., Galí and Gertler (1999) or Roberts (2006). Levin and Piger (2003) argue for a break in the mean inflation rate in 1991.

In the following, we will first show that allowing for changes in mean inflation rates affects the inflation dynamics as measured by the autocorrelation and cross-correlation structure. Since it appears that accounting for changes in the mean inflation rate affects the dynamics of inflation, we investigate whether the average inflation rate around which we approximate the optimal price-setting behavior of the firms in the Calvo model affects the dynamics of the NKPC.

Inflation Dynamics and Average Inflation¹³

The persistence and co-movement of inflation and marginal cost have varied across decades. In Figure 3, we display the autocorrelations and cross-correlations of inflation and marginal cost for the four periods we have just mentioned: the 1960s, 1970s, 1980s, and the period beginning in 1992.

In the 1960s, both inflation and marginal cost are highly persistent, with inflation being somewhat more persistent than marginal cost: the autocorrelation coefficients for inflation do not decline as fast as the ones for marginal cost. But in the following periods, it appears as if the persistence of inflation declines, at least relative to marginal cost. This decline of inflation persistence is especially noticeable for the first- and second-order autocorrelation coefficients from 1984 on, Figure 3, A.3 and A.4.¹⁴


The positive correlation between inflation and marginal cost in the full sample hides substantial variation of co-movement across sub-samples. The 1970s is the only period with a strong positive correlation between inflation and marginal cost, Figure 3, B.2. At the other extreme are the 1960s when the correlation between inflation and marginal cost is negative for almost all leads and lags, Figure 3, B.1. In between are the remaining two sub-samples from 1984 on, in which the correlation between inflation and marginal cost tends to be positive, but only weakly so.

The NKPC at Positive Average Inflation

How should we interpret these changes in the time series properties of inflation and marginal cost? In particular, what do these changes tell us about the NKPC as a model of inflation? The decline in persistence is especially intriguing since it coincides with the decline of the average inflation rate. Most observers

¹³ Articles that discuss changes in the inflation process include Cogley and Sargent (2001), Levin and Piger (2003), Nason (2006), and Stock and Watson (2007). Roberts (2006) and Williams (2006) relate the changes in the inflation process to changes in the Phillips curve.

¹⁴ We should note, however, that the sum of autocorrelation coefficients from univariate regressions in the inflation rate and marginal cost do not indicate statistically significant changes in the persistence of inflation or marginal cost across subperiods, Table 1, columns 5 and 6.

Figure 3 Inflation and Marginal Cost Dynamics Over Time

Notes: In Panel A, the circles (squares) denote the sub-sample autocorrelations for inflation (marginal cost). In Panel B, the diamonds denote the cross-correlations of inflation and marginal cost. In Panels A and B, the boxes denote the 5-percentile to 95-percentile range of the statistic calculated from 1,000 bootstraps of the sub-sample data.

attribute the reduction of the average inflation rate to monetary policy, but should one also attribute the reduced inflation persistence to monetary policy?

From the perspective of the reduced form NKPC with no feedback from inflation to marginal cost, equation (15), monetary policy is unlikely to have affected the persistence of inflation. In this framework, monetary policy works through its impact on marginal cost, but if anything, marginal cost has become more persistent rather than less persistent since the 1990s. We now ask if this conclusion may be premature since it relies on an approximation of the inflation dynamics in the Calvo model around a zero-average inflation rate. If one approximates the inflation dynamics around a positive-average inflation rate, then inflation persistence depends on the average inflation rate, even when the other structural parameters of the environment remain fixed.

The modified hybrid NKPC for an approximation at the gross inflation rate $\bar{\pi} \geq 1$ is

$$E_t \left[(1 - \lambda_1 L^{-1}) (1 - \lambda_2 L^{-1}) (1 - \rho L) \hat{\pi}_t \right] = \kappa_1 E_t \left[(1 + \phi L^{-1}) \hat{s}_t \right] + u_t. \quad (18)$$

The derivation of (18) is described in Hornstein (2007).¹⁵ The NKPC is now a third-order difference equation in inflation and involves current and future marginal cost. The coefficients λ_1 , λ_2 , ϕ , and κ_1 are functions of the underlying structural parameters, α , β , ρ , and a new parameter θ , representing the firms' demand elasticity. Furthermore, the coefficients also depend on the average inflation rate, $\bar{\pi}$, around which we approximate the optimal pricing decisions of the firms.


The modified hybrid NKPC (18) simplifies to the hybrid NKPC (14) for zero net-inflation, $\bar{\pi} = 1$. As we increase the average inflation rate, inflation becomes less responsive to marginal cost in the modified NKPC. In Figure 4.A, we plot the coefficient on marginal cost κ_1 in the modified NKPC as a function of the average inflation rate for our two calibrations of the hybrid NKPC. In addition to the parameter values listed in Table 3, we also have to parameterize the demand elasticity of the monopolistically competitive firms, θ . Consistent with the literature on nominal rigidities, we assume that $\theta = 11$, which implies a 10 percent steady-state markup. For both calibrations, the coefficient on marginal cost declines with the average inflation rate, Figure 4.A. This suggests that everything else being equal, inflation will be less persistent and less correlated with marginal cost at higher inflation rates, since marginal cost has a smaller impact on inflation. The first calibration with a low price adjustment probability represents an extreme case, in that respect, since the coefficient on marginal cost converges to zero. On the other hand, for the second calibration with a higher price adjustment probability, the coefficient on marginal cost is relatively inelastic with respect to changes in the inflation rate.

Assuming that marginal cost follows an AR(1) with persistence δ such that the product of δ and the roots of the lead polynomials in equation (18) are less than one, $|\delta \lambda_i| < 1$, we can derive the reduced form of the modified NKPC as

$$(1 - \rho L) \hat{\pi}_t = \kappa_1 \frac{1 + \delta \phi}{(1 - \lambda_1 \delta) (1 - \lambda_2 \delta)} \hat{s}_t + u_t = a_1 \hat{s}_t + u_t. \quad (19)$$

This expression is formally equivalent to the reduced form of the hybrid NKPC, equation (15), but now the coefficient a_1 is a function of the average inflation rate. Since inflation becomes less responsive to marginal cost in the NKPC

¹⁵ Ascari (2004) and Cogley and Sbordone (2005, 2006) also derive the modified NKPC, but choose a different representation. Their representation is based on the hybrid NKPC, equation (14), and adds a term that involves the expected present value of future inflation.

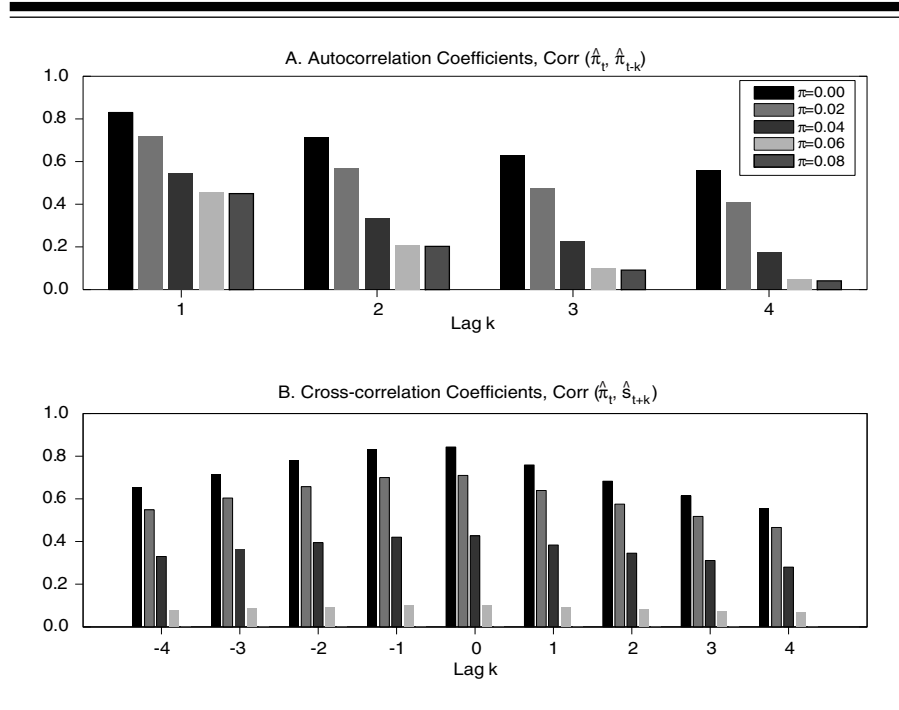
Figure 4 The NKPC and Changes in Average Inflation

when the average inflation rate increases, inflation in the reduced form NKPC also becomes less responsive to marginal cost: a_1 declines with the average inflation rate, Figure 4.B. As with the coefficient on marginal cost in the NKPC, κ_1 , the coefficient on marginal cost in the reduced form NKPC, a_1 , declines much more for the first calibration with the relatively low price adjustment probability. This feature is important since the autocorrelations and cross-correlations of inflation depend on the average inflation rate only through the responsiveness of inflation to marginal cost, a_1 .

We now replicate the analysis of Section 2 and calculate the first four autocorrelation coefficients of inflation and the cross-correlation coefficients of inflation with marginal cost when the average annual inflation rate varies from 0 to 8 percent.¹⁶ In Figures 5 and 6, we display the autocorrelation and cross-correlation coefficients for the two calibrations. With a low price adjustment probability, the first calibration, an increase of the average inflation rate substantially reduces the persistence of inflation and its co-movement with marginal cost, Figure 5. Even moderately high annual inflation rates, about 4

¹⁶ For the parameter values used in the calibration, the “weighted” roots of the lead polynomial are less than one for all of the average annual inflation rates considered.

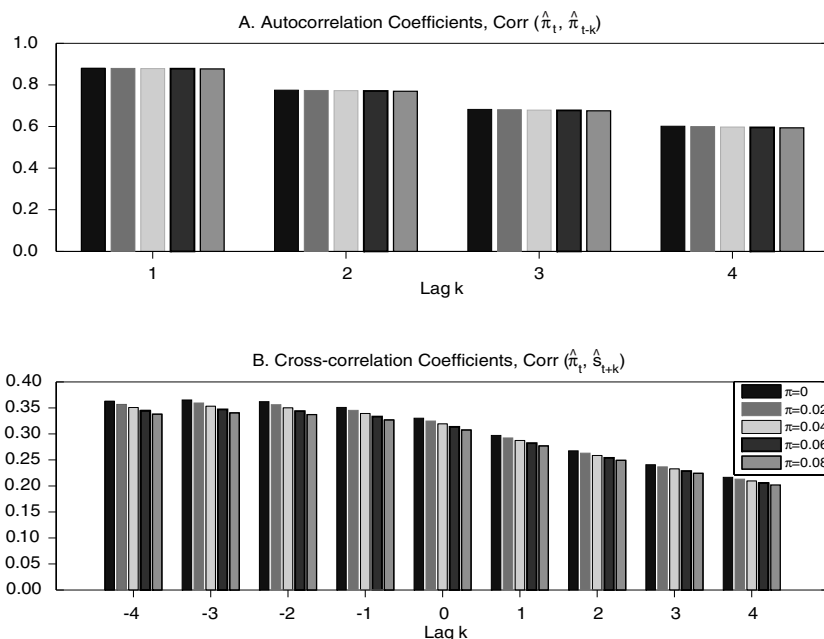
Figure 5 The Effects of Average Inflation, Calibration 1



percent, reduce the first-order autocorrelation and the contemporaneous cross-correlation by half. This pattern follows directly from equations (16) and (17) and the fact that the coefficient a_1 converges to zero for the first calibration. With a higher price adjustment probability, the second calibration, a higher average inflation rate also tends to reduce persistence and co-movement of inflation, but the quantitative impact is negligible, Figure 6. Again, this pattern conforms with the limited impact of changes in average inflation on the reduced form coefficient of marginal cost.

Changing U.S. Inflation Dynamics and the Modified NKPC

Based on the modified NKPC, can changes in average inflation account for the changing U.S. inflation dynamics? Not really. There are two big changes in the average inflation rate between sub-samples of the U.S. economy. First, average inflation increased from 2.5 percent in the 1960s to 6.5 percent in the 1970s, and second, average inflation subsequently declined to 3.2 percent in the 1980s. These changes in average inflation were associated with significant changes in the persistence of inflation and the co-movement of inflation with marginal

Figure 6 The Effects of Average Inflation, Calibration 2

cost. Yet, the predictions of the modified NKPC for inflation persistence and co-movement based on the observed changes in average inflation are inconsistent with the observed changes in persistence and co-movement.

On the one hand, a calibration with relatively low price adjustment probabilities, the first calibration, predicts big changes for persistence and co-movement in response to the changes in average inflation, but the changes either do not take place or are opposite to what the model predicts. In response to the increase of the average inflation rate from the 1960s to the 1970s, inflation persistence and co-movement should have declined substantially, but persistence did not change and co-movement increased. Indeed the correlation between inflation and marginal cost switches from negative, which is inconsistent with the NKPC to begin with, to positive. In response to the reduction of average inflation in the 1980s, the model predicts more inflation persistence and more co-movement of inflation and marginal cost. Yet again, the opposite happens. Inflation persistence declines, at least the first- and second-order autocorrelation coefficients decline, and the correlation coefficients between inflation and marginal cost decline.

On the other hand, a calibration of the modified NKPC with relatively high price adjustment probabilities, the second calibration, cannot account

for any quantitatively important effects on the persistence or co-movement of inflation based on changes in average inflation.

4. CONCLUSION

We have just argued that a hybrid NKPC, modified to account for changes in trend inflation, has problems accounting for the changes of U.S. inflation dynamics over the decades. One way to account for these changes of inflation dynamics within the framework of the NKPC is to allow for changes in the model's structural parameters. For example, inflation indexation, that is, intrinsic persistence, could have increased and decreased to offset the effects of a higher trend inflation in the 1970s. This pattern of inflation indexation in response to the changes in trend inflation looks reasonable. However, attributing changes in the dynamics of inflation to systematic changes in the structural parameters of the NKPC makes this framework less useful for monetary policy analysis. This is troublesome since several central banks have recently begun to develop full-blown Dynamic Stochastic General Equilibrium (DSGE) models with versions of the NKPC as an integral part. Ultimately, these DSGE models are intended for policy analysis, and for this analysis it is presumed that the model elements, such as the NKPC, are invariant to the policy changes considered. Based on the analysis in this article, it then seems appropriate to investigate further the "stability" of the NKPC before one starts using these models for policy analysis.

APPENDIX

We use seasonally adjusted quarterly data for the time period 1955Q1 to 2005Q4. All data are from HAVER with mnemonics in parentheses. From the national income accounts we take real GDP (GDPH@USECON) and for the GDP deflator we take the chained price index (JGDP@USECON). From the nonfarm business sector we take the unit labor cost index (LXNFU@USECON), the implicit price deflator (LXNFI@USECON), and the hourly compensation index (LXNFC@USECON). All of the three nonfarm business sector series are indices that are normalized to 100 in 1992.

We define inflation as the quarterly growth rate of the GDP deflator and marginal cost as the log of the ratio of unit labor cost and the nonfarm business price deflator. We construct the instruments for the GMM estimation other than lagged inflation and marginal cost following Galí, Gertler, and López-Salido (2005). The output gap is the deviation of log real GDP from a quadratic trend, and wage inflation is the growth rate of the hourly compensation index.

REFERENCES

- Ascari, Guido. 2004. "Staggered Prices and Trend Inflation: Some Nuisances." *Review of Economic Dynamics* 7 (3): 642–67.
- Calvo, Guillermo. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.
- Cogley, Timothy, and Thomas Sargent. 2001. "Evolving Post-World War II U.S. Inflation Dynamics." In *NBER Macroeconomics Annual 2001*: 331–72.
- Cogley, Timothy, and Argia M. Sbordone. 2005. "A Search for a Structural Phillips Curve." Federal Reserve Bank of New York *Staff Report* No. 203 (March).
- Cogley, Timothy, and Argia M. Sbordone. 2006. "Trend Inflation and Inflation Persistence in the New Keynesian Phillips Curve." Federal Reserve Bank of New York *Staff Report* No. 270 (December).
- Eichenbaum, Martin, and Jonas D. M. Fisher. 2007. "Estimating the Frequency of Price Re-Optimization in Calvo-Style Models." *Journal of Monetary Economics* 54 (7): 2,032–47.
- Fuhrer, Jeffrey C. 2006. "Intrinsic and Inherited Inflation Persistence." *International Journal of Central Banking* 2 (3): 49–86.
- Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (2): 195–222.
- Galí, Jordi, Mark Gertler, and David López-Salido. 2005. "Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve." *Journal of Monetary Economics* 52 (6): 1,107–18.
- Hornstein, Andreas. 2007. "Notes on the New Keynesian Phillips Curve." Federal Reserve Bank of Richmond Working Paper No. 2007-04.
- Levin, Andrew T., and Jeremy M. Piger. 2003. "Is Inflation Persistence Intrinsic in Industrialized Economies?" Federal Reserve Bank of St. Louis Working Paper No. 2002-023E.
- Nason, James. 2006. "Instability in U.S. Inflation: 1967–2005." Federal Reserve Bank of Atlanta *Economic Review* 91 (2): 39–59.

- Roberts, John M. 2006. "Monetary Policy and Inflation Dynamics." *International Journal of Central Banking* 2 (3): 193–230.
- Sbordone, Argia M. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49 (2): 265–92.
- Stock, James H., and Mark W. Watson. 2007. "Why Has Inflation Become Harder to Forecast?" *Journal of Money, Credit, and Banking* 39 (1): 3–33.
- Williams, John C. 2006. "Inflation Persistence in an Era of Well-Anchored Inflation Expectations." Federal Reserve Bank of San Francisco *Economic Letter* No. 2006-27.
- Woodford, Michael. 2003. *Interest and Prices*. Princeton, NJ: Princeton University Press.

The Evolution of City Population Density in the United States

Kevin A. Bryan, Brian D. Minton, and Pierre-Daniel G. Sarte

The answers to important questions in urban economics depend on the density of population, not the size of population. In particular, positive production or residential externalities, as well as negative externalities such as congestion, are typically modeled as a function of density (Chatterjee and Carlino 2001, Lucas and Rossi-Hansberg 2002). The speed with which new knowledge and production techniques propagate, the gain in property values from the construction of urban public works, and the level of labor productivity are all affected by density (Carlino, Chatterjee, and Hunt 2006, Ciccone and Hall 1996). Nonetheless, properties of the distribution of urban population size have been studied far more than properties of the urban density distribution.

Chatterjee and Carlino (2001) offer an insightful example as to why density can be more important than population size. They note that though Nebraska and San Francisco have the same population, urban interactions occur far less frequently in Nebraska because of its much larger area. Though the differences in the area of various cities are not quite so stark, there are meaningful heterogeneities in city densities. Given the importance of urban density, the stylized facts presented in the article ultimately require explanations such as those given for the evolution of city population.

This article makes two major contributions concerning urban density. First, we construct an electronic database containing land area, population, and urban density for every city with population greater than 25,000 in the

■ We wish to thank Kartik Athreya, Nashat Moin, Roy Webb, and especially Ned Prescott for their comments and suggestions. The views expressed in this article are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System. Data and replication files for this research can be found at http://www.richmondfed.org/research/research_economists/pierre-daniel_sarte.cfm. All errors are our own.

United States. Second, we document a number of stylized facts about the urban density distribution by constructing nonparametric estimates of the distribution of city densities over time and across regions.

We compile data for each decade from 1940 to 2000; by 2000, 1,507 cities meet the 25,000 threshold. In addition, we include those statistics for every “urbanized area” in the United States, decennially from 1950 to 2000. Though we also present data on Metropolitan Statistical Area (MSA) density evolution from 1950 to 1980, this definition of a city can be problematic for work with densities. A discussion of the inherent problems with using MSA data is found in Section 1. To the best of our knowledge, these data have not been previously collected in an electronic format.

Our findings document that the distribution of city densities in the United States has shifted leftward since 1940; that is, cities are becoming less dense. This shift is not confined to any particular decade. It is evident across regions, and it is driven both by new cities incorporating with lower densities, and by old cities adding land faster than they add population. The shift is seen among several different definitions of cities. A particularly surprising result is that “legal cities,” defined in this article as regions controlled by a local government, have greatly decreased in density during the period studied. That is, since 1940, local governments have been annexing territory fast enough to counteract the increase in urban population. Annexation is the only way that cities can simultaneously have increasing population, which is true of the vast majority of cities in our sample, and yet still have decreasing density.

This article is organized as follows. Section 1 describes how our database was constructed, and also discusses which definition of city is most appropriate in different contexts. Section 2 discusses our use of nonparametric techniques to estimate the distribution of urban density. Section 3 presents our results and discusses why cities might be decreasing in density. Section 4 concludes.

1. DATA

What is a city? There are at least three well-defined concepts of a city boundary in the United States that a researcher might use: the legal boundary of the city, the boundary of the built-up, urban region around a central city (an “urbanized area”), and the boundary of a census-defined Metropolitan Statistical Area (MSA). The legal boundary of a city is perhaps most relevant when investigating the area that state and local governments believe can be covered effectively with a single government. Legal boundaries also have the advantage of a consistent definition over the period studied; this is not completely true for urbanized areas, and even less true for MSAs. Urbanized areas parallel nicely with an economist’s mental image of an agglomeration, as they include the built-up suburban areas around a central city. MSAs, though commonly used in the population literature, offer a much vaguer in-

Table 1 Three Definitions of a City

Legal City	The region controlled by a local government or a similar unincorporated region (CDP). <i>Defined by local and state governments.</i>
Urbanized Area	A region incorporating a central city plus surrounding towns and cities meeting a density requirement. <i>Defined by the U.S. Census Bureau.</i>
MSA	A region incorporating a central city, the county containing that city, and surrounding counties meeting a requirement on the percentage of workers commuting to the center. <i>Defined by the U.S. Census Bureau.</i>

terpretation. Figure 1 displays the city, urbanized area, and MSA boundaries for Richmond, Virginia, and Las Vegas, Nevada, in the year 2000.

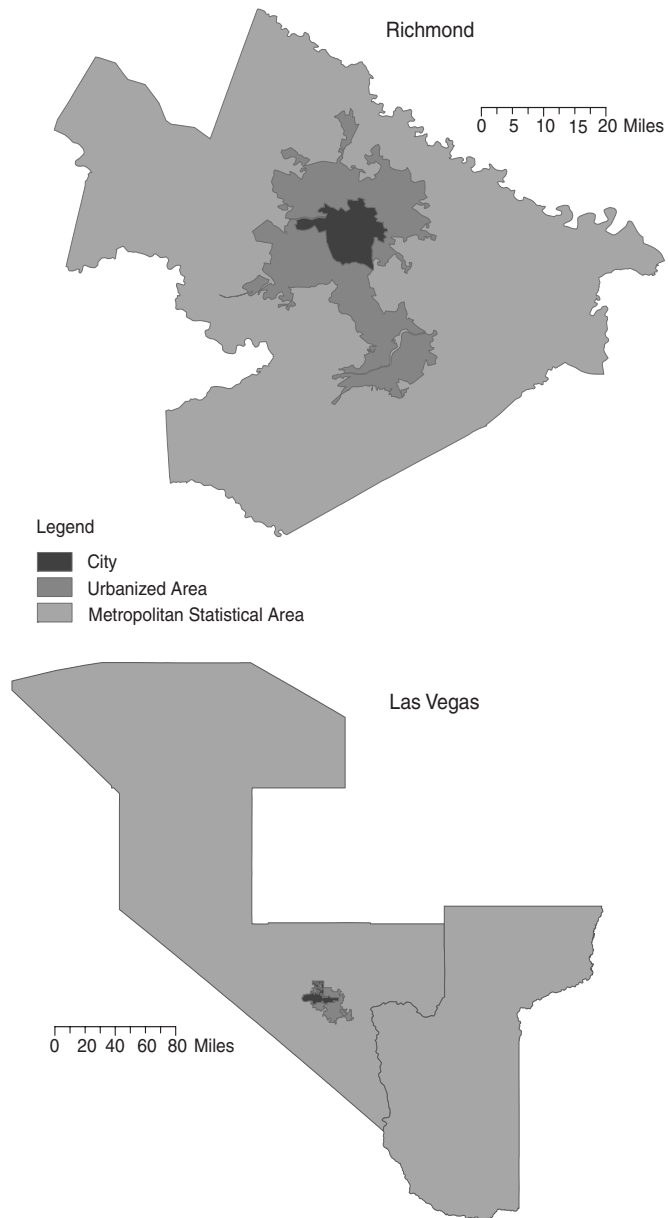
Our database of legal cities is constructed from the decennial U.S. Bureau of the Census *Number of Inhabitants*, which is published two to three years after each census is taken. Population and land area for every U.S. “place” with a population greater than 2,500 are listed. Places include cities, towns, villages, urban townships, and census-designated places (CDPs). Cities, towns, and townships are legally defined places containing some form of local government, while a census-designated place (called an “unincorporated place” before 1980) refers to unincorporated areas with a “settled concentration of population.” Some of these CDPs can be quite large; for instance, unincorporated Metairie, Louisiana has a population of nearly 150,000 in 2000. Though CDPs do not represent any legal entity, they are nonetheless defined in line with settlement patterns determined after census consultation with state and local officials, and are similar in size and density to incorporated cities.¹ Including CDPs in our database, and not simply incorporated cities, is particularly important as some states *only* have CDPs (such as Hawaii), and “towns” in eight states, including all of New England, are only counted as a place when they appear as a CDP.

From this list, we selected every place (including CDPs) with a population greater than 25,000 for each census from 1940 to 2000. There are 412 places in 1940 and 1,507 places in 2000 that meet this restriction. Each place was coded into one of nine geographical regions in line with the standard census region definition.² We also labeled each place as either “new” or “old.” An

¹ 1980 *Census of Population: Number of Inhabitants*. “Appendix A—Area Classification.” U.S. Department of Commerce, 1983. Note that CDPs did not appear in the 1940 Census.

² “Census Regions and Divisions of the United States.” Available online at http://www.census.gov/geo/www/us_regdiv.pdf.

Figure 1 A Graphic Representation of City Definitions



old place is a place that had a population greater than 25,000 in 1940 and still has a population greater than 25,000 in 2000. A new place is one that had a population less than 25,000 or did not exist at all in 1940, yet has a population greater than 25,000 in 2000. There are some places which had a population greater than 25,000 in 1940 but less than 25,000 in 2000 (for instance, a number of Rust Belt cities with declining populations); we considered these places neither new or old. Delineating places in this manner allows us to investigate whether the leftward shift of the distribution of U.S. cities was driven by newly founded cities having a larger area, or by old cities annexing area faster than their population increases.

In addition to legal cities, we also construct a series of urbanized areas from the *Number of Inhabitants* publication. Beginning in 1950, the U.S. Census defined urbanized areas as places with a population of 50,000 or more, meeting a minimum density requirement, plus an “urban fringe” consisting of places roughly contiguous with the central city meeting a small population requirement; as such, urbanized areas are defined in a similar way as agglomerations in many economic models. Aside from 1960, when the density requirement for central cities was lowered from approximately 2,000 people per square mile to 1,000 per square mile, changes in the definition of an urbanized area have been minor.³ Our database includes each urbanized area from 1950 to 2000; there were 157 such areas in 1950 and 452 in 2000.

Much of the literature on city population uses data on Metropolitan Statistical Areas (MSAs). An MSA is defined as a central urban city, the county containing that city, and outlying counties that meet certain requirements concerning population density and the number of residents who commute to the central city for work.⁴ We believe there are a number of reasons that this data can be problematic for investigating city density. First, it is difficult to get consistent data on metro areas. Before 1950, they were not defined at all, though Bogue (1953) constructed a series of MSA populations for 1900–1940 by adding up the population within the area of each MSA as defined in 1950. Because, by definition, Bogue holds MSA area constant for 1900–1950, this data set would not pick up any changes in density caused by the changing *area* of a city over time. Furthermore, there was a significant change in how MSAs are defined in 1983, with the addition of the “Consolidated Metropolitan Statistical Area” (CMSA). Because of this, MSAs between 1980 and 1990 are not comparable. Dobkins and Ioannides (2000) construct MSAs for 1990 using the 1980 definition, but no such series has been constructed for 2000.

Second, the delineation of MSAs is highly dependent on county definitions. Particularly in the West, counties are often much larger than in the

³ See the *Geographic Areas Reference Manual*, U.S. Bureau of the Census, chap. 12. Available online at: <http://www.census.gov/geo/www/garm.html>.

⁴ In New England, the town, rather than the county, is the relevant area.

Midwest and the East. For instance, in 1980, the Riverside-San Bernardino-Ontario, California MSA had an area of 27,279 square miles and a population density of 57 people per square mile.⁵ This MSA has an area three times the size of and a lower population density than Vermont.⁶ When looking solely at population, MSAs can still be useful because the population in outlying rural areas tends to be negligible; this is not the case with area, and therefore density.

Third, the number of MSAs is problematic in that it truncates the number of available cities such that only the far right-hand tail of the population distribution is included. For instance, Dobkins and Ioannides' (2000) MSA database includes only 162 cities in 1950, rising to 334 by 1990. For cities and census-designated places, three to four times as much data can be used. Eeckhout (2004) notes that the distribution of urban population size is completely different when using a full data set versus a truncated selection that includes only MSAs; it seems reasonable to believe that urban density might be similar in this regard. Further, nonparametric density estimation, as used in this article, requires a large data set. For completeness, we show in Section 3 that the distribution of densities in MSAs from 1950 to 1980, when the MSA definition was roughly consistent, follows a similar pattern to that of urbanized areas and legal cities.

Other than the database used in this article, we know of no other complete panel data set of urban density for U.S. cities. For 1990 and 2000, a full listing of places with area and population is available online as part of the U.S. Census Gazetteer.⁷ The County and City Data Books, hosted by the University of Virginia, Geospatial and Statistical Data Center, hold population and area data for 1930, 1940, 1950, 1960, and 1975; these data were entered by hand during the 1970s from the same census books we used.⁸ However, crosschecking this data with the actual census publications revealed a number of minor errors, and further indicated that unincorporated places and urban towns were not included. For some states (for instance, Connecticut and Maryland), this means that very few places were included in the data set at all. Our data set rectifies these omissions.

⁵ The MSA was made up of two counties: Riverside County with an area of 7,214 square miles, and San Bernardino County with an area of 20,064 square miles.

⁶ In fact, the entire planet has a land area of around 58 million square miles and a population of 6.5 billion, giving a density of 112 people per square mile, or twice the density of the Riverside MSA.

⁷ The 1990 data can be found at <http://www.census.gov/tiger/tms/gazetteer/places.txt>. Data for 2000 are available at: <http://www.census.gov/tiger/tms/gazetteer/places2k.txt>.

⁸ County and City Data Books. University of Virginia, Geospatial and Statistical Data Center. Available online at: <http://fisher.lib.virginia.edu/collections/stats/ccdb/>.

2. NONPARAMETRIC ESTIMATION

With these density data, we estimate changes in the probability density function (*pdf*) over time for each definition of a city in order to examine, for instance, how the distribution of urban densities is changing over time. We use nonparametric techniques, rather than parametric estimation, because nonparametric estimators make no underlying assumption about the distribution of the data (for instance, the presence or lack of normality). Assuming, for instance, an underlying normal distribution might mask evidence of a true bimodal distribution, and given our lack of priors concerning the distribution of urban densities, nonparametric estimates offer more flexibility. Potential pitfalls in nonparametric estimation are the requirement of larger data sets, and the computational difficulty of calculating *pdf* estimates with more than two or three variables;⁹ however, our data sets are large and our estimated *pdfs* are univariate. Nonparametric estimates of a *pdf* are closely related to the histogram; a description of this link, and basic nonparametric concepts, is given in Appendix A.

One frequently used nonparametric *pdf* estimator is the Rosenblatt-Parzen estimator,

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^n K(\psi_i),$$

where n is the number of observations, h is a “smoothing factor” to be chosen below, $\psi_i = \frac{x-x_i}{h}$, and K is a nonparametric kernel. The smoothing factor determines the interval of points around x which are used to compute $\hat{f}(x)$, and the kernel determines the manner in which an estimator weighs those points. For instance, a uniform kernel would weigh all points in the interval equally.

In practice, the choice of kernel is relatively unimportant. In this article, we use one of the more common kernels, namely the Gaussian kernel,

$$K(\psi_i) = (2\pi)^{-.5} e^{-\frac{\psi_i^2}{2}}.$$

This kernel uses a weighted average of all observations, with weights declining in the distance of each observation from x_i .

The choice of bandwidth h , on the other hand, can be important, and is often chosen so as to minimize an error function of bias and variance. Given a set of assumptions about the nature of $f(x)$, the Rosenblatt-Parzen estimator

⁹ Nonparametric estimates converge to their true values at a rate slower than \sqrt{n} .

$\hat{f}(x)$ is such that¹⁰

$$\text{Bias} = \frac{h^2}{2} \left[\int \psi^2 K(\psi) d\psi \right] f''(x) + O(h^2) \quad (1)$$

and

$$\text{Variance} = \frac{1}{nh} f(x) \int K^2(\psi) d\psi + O\left(\frac{1}{nh}\right). \quad (2)$$

A low bandwidth, h , gives low bias but high variance, whereas a high h will give high bias but low variance. That is, choosing too small of a value for h will cause the estimated density to lack smoothness since not enough sample points will be used to calculate each $\hat{f}(x_i)$, whereas too high a value for h will smooth out even relevant bumps such as the trough in a bimodal distribution. A description of the assumptions necessary for our bias and variance formulas can be found in Appendix B.

The integrated mean squared error is defined as

$$\int [\text{Bias}(\hat{f}(x))^2 + \text{V}(\hat{f}(x))] dx. \quad (3)$$

This function simultaneously accounts for bias and variance. It is analogous to the conventional mean squared error in a parametric estimation. When h is chosen to minimize (3) after substituting for the bias and variance using expressions (1) and (2) respectively, we obtain

$$h = cn^{-\frac{1}{5}} \text{ where } c = \left[\frac{\int K^2(\psi) d\psi}{\left[\int \psi^2 K(\psi) d\psi \right]^2 \int (f''(x))^2 dx} \right]^{\frac{1}{5}}.$$

Since $f(x)$ is unknown, and the formula for h involves knowing the true $f''(x)$, no more can be said about h without making some assumptions about the nature of $f(x)$. For example, if $f(x) \sim N(\mu, \sigma^2)$, then $c = 1.06\hat{\sigma}$, and therefore $h = 1.06\hat{\sigma}n^{-\frac{1}{5}}$ exactly.¹¹ This formula is called Silverman's Rule of Thumb, and works very well for data that is approximately normally distributed (Silverman 1986). Silverman notes that this rule does not necessarily work well for bimodal or heavily skewed data, and some of the series in this article (for instance, city populations) are heavily skewed. In particular, outliers lead to large increases in the estimated standard deviation, $\hat{\sigma}$, and therefore a very large value for h . Consequently, this article instead uses Silverman's more general specification

$$h = .9Bn^{-\frac{1}{5}}$$

¹⁰ If $\frac{X_n}{n^k} \rightarrow$ some real number c as $n \rightarrow \infty$, then X_n is $O(n^k)$. $O(A)$ is the largest order of magnitude of a sequence of real numbers X_n .

¹¹ Note that this rule does not imply that the nonparametric estimate will look like a parametric normal distribution; it merely says that, given data that are roughly normal, $1.06\hat{\sigma}n^{-\frac{1}{5}}$ is the smoothing factor that minimizes both bias and variance.

given

$$B = \min(\hat{\sigma}, \frac{IQR}{1.34}),$$

where IQR is the interquartile range of sample data. This formula is much less sensitive to outliers than the Rule of Thumb. In practice, this has shown to be nearly optimal for somewhat skewed data.

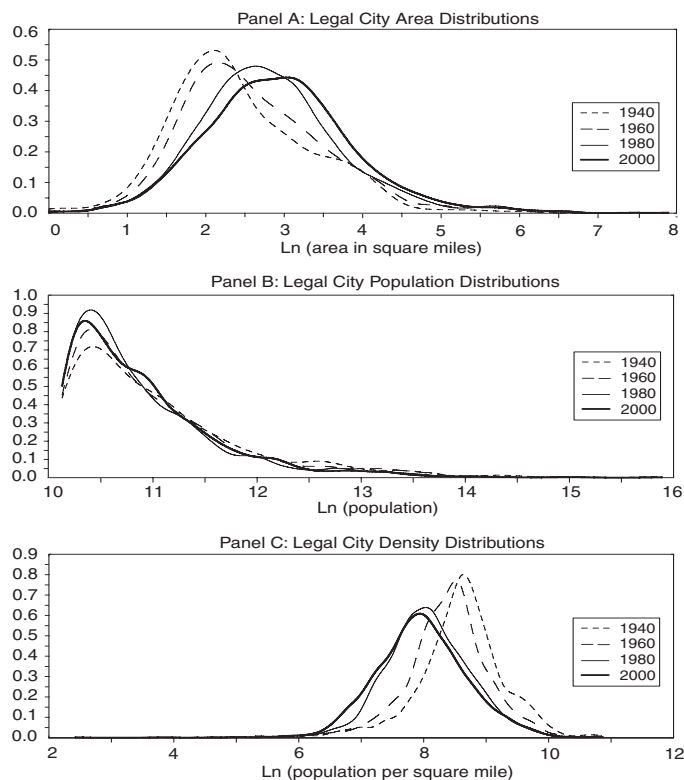
3. RESULTS

Using the kernel and smoothing parameter from the previous section, we can construct estimates of the *pdf* of the distribution of population, area, and urban density in each decade.

Figure 2 shows nonparametric estimations of the distributions of population size, area, and density for legal cities as defined in Section 1. Panel C shows a leftward shift of the distribution of city densities; that is, cities in 2000 are significantly less dense than in 1940. The mean population per square mile during that period fell from 6,742 to 3,802. This is being driven principally by an increase in the area of each city; mean area has increased from 19.2 square miles to 35.1 square miles between 1940 and 2000. The distribution of populations has remained relatively constant during this period.

One might imagine that this shift is being driven only by a subset of cities, such as rapidly-growing suburban and exurban cities, or cities in the West where land is less scarce. Hence, we divide cities into “new” and “old,” as defined in Section 1, as well as categorize each city into one of four regions: East, South, Midwest, and West. Figure 3 shows that the leftward shift in distribution is similar among both old and new cities; that is, city density is decreasing both because existing cities are annexing additional area, and because new cities have lower initial densities than in the past. The number of cities that change their legal boundaries in a given decade is surprising; for instance, between 1990 and 2000, nearly 36 percent of the cities in our data set added or lost at least one square mile. These changes vary enormously by state, however, in a state such as Massachusetts, where all of the land has been divided into towns for decades, there is very little opportunity for a city to add territory. Alternatively, in a state such as Oregon where the majority of land is unincorporated, annexation is much more common. Might it then be the case that the shift in city density is specific to the Midwest and West, where annexation is frequent?

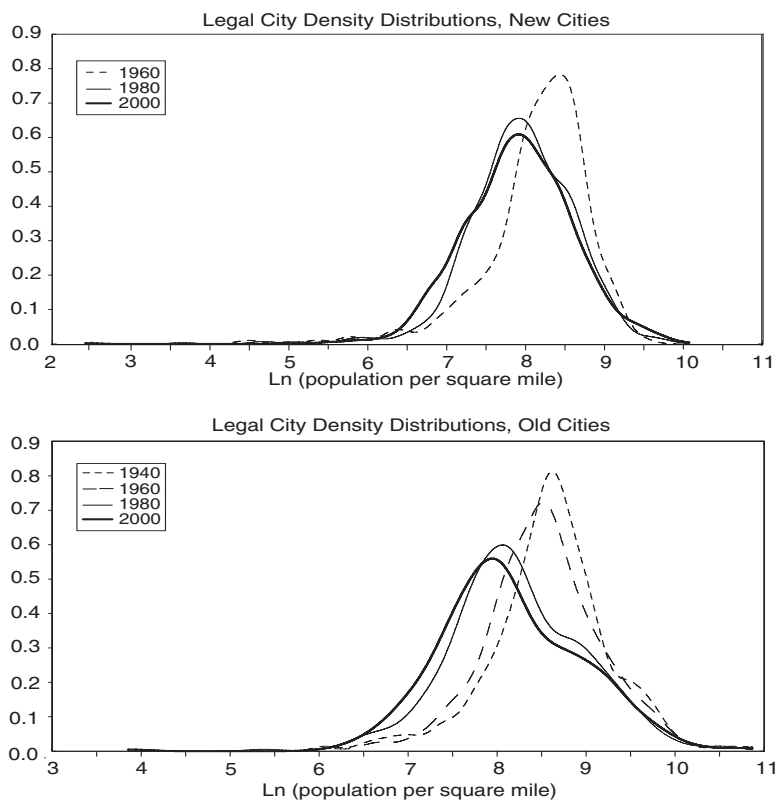
In fact, the leftward shift in city density does not appear to be a regional phenomenon. Figure 4 shows the distribution of densities in the East, South, Midwest, and West during the period 1940–2000. Each region showed a similar decline in density. The full distribution of log density from the Rosenblatt-Parzen estimator is particularly useful when examining the relatively small number of cities in each region when compared to a simple table of moments,

Figure 2 Legal City Area, Population, and Density

as extreme outliers in the data can result in high skewness. For instance, Juneau, Alaska, had an area of 2,716 square miles and a population of 30,711 in 2000, giving a density of approximately 11 people per square mile.

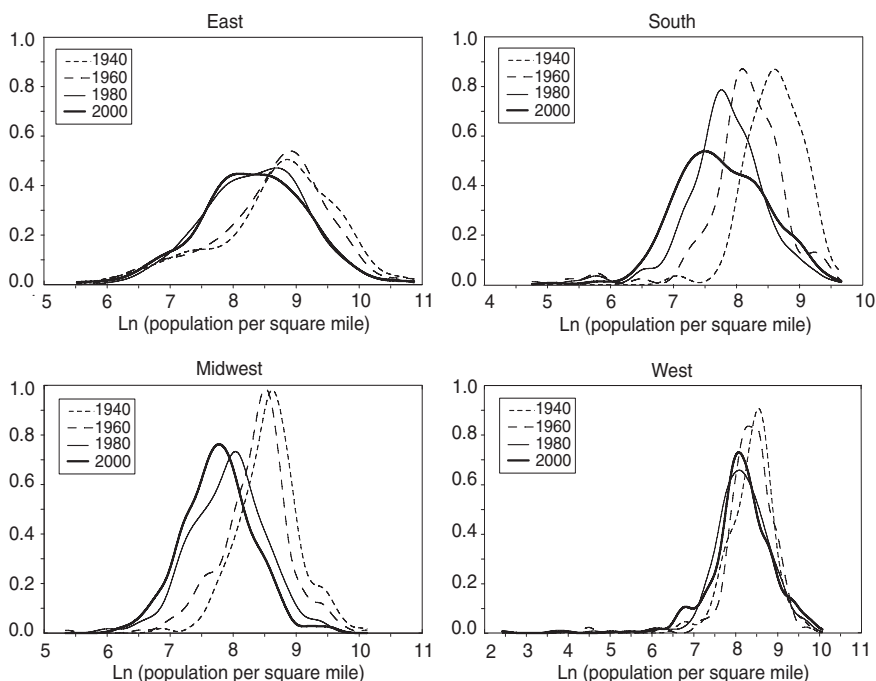
The trend in density is even clearer if we look at urbanized areas. Urbanized areas can be reasonably thought of as urban agglomerations; they represent the built-up area surrounding a central city. Figure 5 shows the estimated distribution of urbanized areas in 1960, 1980, and 2000. As in the case of legal cities, there has been a clear decrease in the density of urbanized areas during this period. Because the boundaries of urbanized areas and legal cities are quite different, it is rather striking that, under both definitions, the decrease in density has been so evident. That is, cities have not simply expanded into a mass of lower-density suburbs, but the individual cities and suburbs themselves have decreased in density, primarily by annexing land.

Finally, we consider the density of Metropolitan Statistical Areas. As noted in Section 1, there are only consistently defined MSA data available for the

Figure 3 Distributions of New and Old Cities

period 1950–1980. Furthermore, a decrease in the distribution of MSA density might simply reflect the increase in the number of MSAs in states with large counties, since each MSA by definition includes its own county. The urban economics literature concerning population size, however, often uses MSAs. Figure 6 shows that the distribution of MSA population density also appears to be shifting leftward in the same manner as legal cities and urbanized areas, but again, it is hazardous to give any interpretation to this shift. The definitional advantages and large data sample size for urbanized areas and legal cities potentially makes them preferable to MSAs for future work concerning urban density.


The importance of these shifts in urban density is underscored by the long-understood link between density and economic prosperity. Lucas (1988) cites approvingly Jane Jacobs' contention that dense cities, not simply cities, are the economic "nucleus of an atom," the central building block of development through their role in spurring human capital transfers. Ciccone and Hall

Figure 4 Distributions of Urban Density by Region

(1996), using county-level data, find that a doubling of employment density in a county increases labor productivity by 6 percent. In addition to knowledge transfer, agglomerations arise in order to facilitate effective matches between employer and employee and to take advantage of external economies of scale such as a common deepwater port.

Measuring the nature of local knowledge transfer, and in particular whether the relevant area has expanded as transportation and communication technologies have fallen, is difficult. Jaffe, Trajtenberg, and Henderson (1993) find evidence that, given the existing distribution of industries and research activity, new patents tend to cite existing patents from the same state and MSA at an unexpectedly high level. Using data on the urbanized portion of a metropolitan area, Carlino, Chatterjee, and Hunt (2006) find that patents per capita rise 20 percent as the employment density of a city doubles. They also find that the benefits of density are diminishing over density, so that cities with employment densities similar to Philadelphia and Baltimore, around 2,100 jobs per square mile, are optimal.


Given the economic benefits of density, the changes in the urban density distribution presented in this article suggest two questions. First, why

Figure 5 Distribution of Urbanized Areas

have agglomeration densities decreased? Second, why have the areas of legal jurisdictions increased?

Decreased densities in urban areas have been explained by a number of processes in the literature, including federal mortgage insurance, the Interstate Highway System, racial tension, and schooling considerations. Mieszkowski and Mills (1993) counter that these explanations tend to be both unique to the United States and are phenomena of the postwar period, whereas a decrease in urban density began as early as 1900 and has occurred across the developed world. Two theories remain.

First, the decreased transportation costs brought about by the automobile and the streetcar has allowed congestion in central cities to be avoided by firms and consumers. Glaeser and Kahn (2003) point out that the automobile also has a supply-side effect in that it allows factories and other places of work to decentralize by eliminating the economies of scale seen with barges and railroads; the rail industry was three times larger than trucking in 1947, but trucks now carry 86 percent of all commodities in the United States. Whereas the wealthy in the nineteenth century might have preferred to live in the center of a city while the poor were forced to walk from the outskirts, the modern

Figure 6 Distribution of MSAs

well-to-do are less constrained by transport times and, therefore, occupy land in less-dense suburban and exurban cities.

Rossi-Hansberg, Sarte, and Owens (2005) present a model in which firms set up non-integrated operations such that managers work in cities in order to take advantage of knowledge transfer externalities but production workers tend to work at the periphery of a city where land costs are lower. They then show that, as city population grows, the internal structure of cities changes along a number of dimensions that are consistent with the data.

A second theory, not entirely independent from the first, posits that cities have become less dense because of a desire for homogenization. When a large group with relatively homogenous preferences for tax rates and school quality is able to occupy its own jurisdiction, it can use land-use controls to segregate itself from potential residents with a different set of preferences. Mieszkowski and Mills (1993) argue that land-use restrictions have become more stringent in the postwar era, and that segregation into income-homogenous areas may be contributing to decreased densities.

There are fewer existent theories about why legal jurisdictions, at a given population level, have increased in area. Glaeser and Kahn (2003) note that effective land use requires larger jurisdictions as transportation costs fall. That

is, if a city wished to limit sprawl in an era with high transportation costs, it could enact effective land-use regulations within small city boundaries. In an era with low transportation costs, however, such a regulation would simply push residents into another bedroom community and have no effect on sprawl or traffic. The growing number of regional land-use planning commissions, such as Portland's Metropolitan Service District and Atlanta's Regional Commission, speak to this trend (Song and Knaap 2004).

Austin (1999) discusses reasons why cities may want to annex territory, including controlling development on the urban fringe, increasing the tax base, lowering the cost of municipal services, lowering municipal service costs by exploiting returns to scale, or altering the characteristics of the city, such as decreasing the minority proportion of population. External areas may wish to be annexed because of urban economies of scale, and because urban areas offer benefits such as cheaper bond issuance than suburban and unincorporated areas. Austin finds evidence that cities annex for both political and economic reasons, but that increasing the tax base does not appear to be a relevant factor, perhaps because of the growing ability of high-wealth areas to avoid annexation by poorer cities.

4. CONCLUDING REMARKS

This article provides two novel contributions. First, it constructs an electronic data set of urban densities in the United States during the previous seven decades for three different definitions of a city. Second, it applies non-parametric techniques to estimate the distribution of those densities, and finds that there has been a stark decrease in density during the period studied. This deconcentration has been occurring continuously since at least 1940, in every area of the United States, and among both new and old cities. This result is striking; increasing population and increasing area across cities do not, by themselves, tell us what will happen to density.

Falling urban densities suggest that, over the past seven decades, the productivity benefits of dense cities have been weakening. Decreasing costs of transportation and communication have allowed firms to move production workers out of high-rent areas, and have allowed residents to move away from downtowns. It is unclear what effect these changes in the urban landscape will have on knowledge accumulation and growth in the future. For instance, it is conceivable that the productivity loss from ever-decreasing spatial density might be counteracted by decreased long-range communication costs. Understanding the broad properties of urban density in modern economies is merely a necessary first step in understanding how these changing properties of cities will affect the broader economy.

APPENDIX A: NONPARAMETRIC ESTIMATORS

Classical density estimation assumes a parametric form for a data set and uses sample data to estimate those parameters. For instance, if an underlying process is assumed to generate normal data, the estimated density is

$$\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\mu)^2}{2\sigma^2}},$$

where σ and μ are the sample standard deviation and mean.

Nonparametric density estimation, on the other hand, allows a researcher to estimate a complete density function from sample data, and therefore estimate each moment of that data, without assuming any underlying functional form. For instance, if a given distribution is bimodal, estimating moments under the assumption of normally distributed data will be misleading. Knowing the full distribution of data also makes clear what stylized facts need to be explained in theory; if the data were skewed heavily to the right and suffered from leptokurtosis, a theory explaining that data should be able to replicate these properties. Nonparametric estimation generally requires a larger data set than parametric estimation to achieve consistency, but is becoming more common in the literature. Given that our city data set is large, we use nonparametric techniques in this article. A brief introduction to these techniques can be found in Greene (2003), while a more complete treatment is found in Pagan and Ullah (1999).

At its core, a nonparametric density estimate is simply a smoothed histogram. Therefore, the nonparametric estimator can be motivated by beginning with a histogram. In a histogram, the full range of n sample values is partitioned into non-overlapping bins of equal width h . Each bin has a height equal to the number of sample observations within the range of that bin divided by the total number of observations. Given an indicator function $\mathbf{I}(A)$, defined as equal to 1 if the statement A is true, and 0 if the statement A is false, the height of a bin centered at some point x_0 , with width h , is

$$H(x_0) = \frac{1}{n} \sum_{i=1}^n \mathbf{I}(x_0 - \frac{h}{2} < x_i \leq x_0 + \frac{h}{2}).$$

That is, we are simply counting the number of sample observations in each bin of width h , and dividing that frequency by the sample size; the resulting height of each bin is the relative frequency. If there are 40 observations, of which 10 are in the bin $(1,2]$, with $h = 1$, then the histogram has height $H(1.5) = .25$ for all x in $(1,2]$.

This concept can be extended by computing a “local” histogram for each point x in the range $(x_{\min} - \frac{h}{2}, x_{\max} + \frac{h}{2}]$, where x_{\min} and x_{\max} are the minimum

and maximum values in the sample data.¹² In the histogram above, we computed $H(x_0)$ for only h points in the range; x_0 was required to be the midpoint of a bin. The local histogram will instead calculate $\hat{f}(x)$ for every x in $(x_{\min} - \frac{h}{2}, x_{\max} + \frac{h}{2})$, where $\hat{f}(x)$ evaluated at a given point x_0 is equal to the number of sample observations within $(x_0 - \frac{h}{2}, x_0 + \frac{h}{2})$, divided by n to give a frequency.¹³ That is,

$$\begin{aligned} \hat{f}(x) &= \frac{1}{n} \sum_{i=1}^n \mathbf{I}(x - \frac{h}{2} < x_i < x + \frac{h}{2}) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{I}(|\frac{x - x_i}{h}| < \frac{1}{2}) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{I}(|\psi(x_i)| < \frac{1}{2}), \end{aligned}$$

where $\psi(x_i) = \frac{x-x_i}{h}$. $\hat{f}(x)$ is a proper density function if, first, it is greater than or equal to zero for all x , which is guaranteed since the indicator function is always either 0 or 1, and second, if $\int_{-\infty}^{\infty} \hat{f}(x)dx = 1$. Dividing $\hat{f}(x)$ by h ensures that the function integrates to one. To see this, observe first that

$$\int_{-\infty}^{\infty} \mathbf{I}(|\psi(x_i)| < \frac{1}{2})d\psi = \int_{-\frac{1}{2}}^{\frac{1}{2}} \mathbf{I}(|\psi(x_i)| < \frac{1}{2})d\psi = \int_{-\frac{1}{2}}^{\frac{1}{2}} d\psi = 1.$$

In addition, since $\psi(x_i) = \frac{x-x_i}{h}$,

$$\begin{aligned} \frac{1}{h} \int_{-\infty}^{\infty} \hat{f}(x)dx &= \frac{1}{nh} \sum_{i=1}^n \int_{-\infty}^{\infty} \mathbf{I}(|\frac{x - x_i}{h}| < \frac{1}{2})dx \\ &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} \mathbf{I}(|\frac{x - x_i}{h}| < \frac{1}{2})d\psi \\ &= 1. \end{aligned}$$

While local histograms certainly provide a nonparametric estimate of density, and are smoother than proper histograms, they are still discontinuous. It seems sensible, then, to attempt to smooth the histogram. This is done by

¹²The local histogram $\hat{f}(x)$ must be computed for $(x_{\min} - \frac{h}{2}, x_{\max} + \frac{h}{2}]$ and not simply for $(x_{\min}, x_{\max}]$, because $\hat{f}(x) > 0$ for points outside of $(x_{\min}, x_{\max}]$. For instance, if $h = 1$ and $(x_{\min}, x_{\max}] = (0, 10]$, $\hat{f}(10.4)$ will be greater than zero because it will count the sample observation $x_0 = 10$.

¹³In practice, $\hat{f}(x)$ can only be computed for a finite number of points. The distributions we display in Section 5 have been computed at 1,000 points evenly divided on the range (x_{\min}, x_{\max}) .

replacing the indicator function in

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^n \mathbf{I}\left(\left|\frac{x - x_i}{h}\right| < \frac{1}{2}\right)$$

with another function called a kernel, $K(\psi)$, such that $\hat{f}(x) \geq 0$, integrates to one and is smooth. An estimator of the form

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^n K(\psi_i), \text{ where } \psi_i = \frac{x - x_i}{h},$$

is a Rosenblatt-Parzen kernel estimator, and the resulting function $\hat{f}(x)$ depends on the choice of h , called a bandwidth or smoothing parameter, and the choice of kernel. A “good” density estimate will have low bias (that is, $E(\hat{f}(x)) - f(x)$, where $f(x)$ is the true density of the data) and low variance.

APPENDIX B: ROSENBLATT-PARZEN BIAS AND VARIANCE

Bias and variance of a nonparametric estimator can be calculated given the following four assumptions:

- 1) The sample observations are i.i.d.
- 2) The kernel is symmetric around zero and satisfies $\int_{-\infty}^{\infty} K(\psi)d\psi = 1$, $\int_{-\infty}^{\infty} \psi^2 K(\psi)d\psi \neq 0$, and $\int_{-\infty}^{\infty} K^2(\psi)d\psi < \infty$.
- 3) The second-order derivatives of \hat{f} are continuous and bounded around x , and
- 4) $h \rightarrow 0$ and $nh \rightarrow \infty$ as $n \rightarrow \infty$.

It can be shown that the Rosenblatt-Parzen estimator $\hat{f}(x)$ has

$$\text{Bias} = \frac{h^2}{2} \left[\int \psi^2 K(\psi)d\psi \right] f''(x) + O(h^2)$$

and

$$\text{Variance} = \frac{1}{nh} f(x) \int K^2(\psi)d\psi + O\left(\frac{1}{nh}\right).$$

The integrated mean squared error (MISE) is defined as

$$\int [\text{Bias}(\hat{f}(x))^2 + \text{V}(\hat{f}(x))] dx.$$

Substituting the formulas for bias and variance, and ignoring the higher order terms, $O(h^2)$ and $O(\frac{1}{nh})$, respectively, gives the asymptotic integrated

mean squared error (AMISE):

$$\begin{aligned} & \frac{h^4}{4} \left[\int \psi^2 K(\psi) d\psi \right]^2 \int (f''(x))^2 dx + \frac{1}{nh} \int f(x) dx \int K^2(\psi) d\psi \\ &= \frac{h^4}{4} \left[\int \psi^2 K(\psi) d\psi \right]^2 \int (f''(x))^2 dx + \frac{1}{nh} \int K^2(\psi) d\psi. \end{aligned}$$

Differentiating with respect to h and setting the result equal to zero, we have

$$h^3 \left[\int \psi^2 K(\psi) d\psi \right]^2 \int (f''(x))^2 dx - \frac{1}{nh^2} \int K^2(\psi) d\psi = 0$$

or

$$h = cn^{-\frac{1}{5}}, \text{ where } c = \left[\frac{\int K^2(\psi) d\psi}{\left[\int \psi^2 K(\psi) d\psi \right]^2 \int (f''(x))^2 dx} \right]^{\frac{1}{5}}.$$

REFERENCES

- Austin, D. Andrew. 1999. "Politics vs. Economics: Evidence from Municipal Annexation." *Journal of Urban Economics* 45 (3): 501–32.
- Bogue, Donald J. 1953. *Population Growth in Standard Metropolitan Areas 1900–1950*. Oxford, Ohio: Scripps Foundation in Research in Population Problems.
- Carlino, Gerald, Satyajit Chatterjee, and Robert M. Hunt. 2006. "Urban Density and the Rate of Invention." Federal Reserve Bank of Philadelphia Working Paper No. 06-14.
- Chatterjee, Satyajit, and Gerald A. Carlino. 2001. "Aggregate Metropolitan Employment Growth and the Deconcentration of Metropolitan Employment." *Journal of Monetary Economics* 48 (3): 549–83.
- Ciccone, Antonio, and Robert E. Hall. 1996. "Productivity and the Density of Economic Activity." *American Economic Review* 86 (1): 54–70.
- Dobkins, Linda, and Yannis Ioannides. 2000. "Dynamic Evolution of the Size Distribution of U.S. Cities." In *The Economics of Cities*, eds. J. Huriot and J. Thisse. New York, NY: Cambridge University Press.
- Eeckhout, Jan. 2004. "Gibrat's Law for (All) Cities." *American Economic Review* 94 (5): 1,429–51.
- Glaeser, Edward L., and Matthew E. Kahn. 2003. "Sprawl and Urban Growth." In *Handbook of Regional and Urban Economics*, eds. J. V.

- Henderson and J. F. Thisse, 1st ed., vol. 4, chap. 56. North Holland: Elsevier.
- Greene, William. 2003. *Econometric Analysis*. 5th ed. Upper Saddle River, NJ: Prentice Hall.
- Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics* 108 (3): 577–98.
- Lucas, Robert E., Jr. 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22 (1): 3–42.
- Lucas, Robert E., Jr., and Esteban Rossi-Hansberg. 2002. "On the Internal Structure of Cities." *Econometrica* 70 (4): 1,445–76.
- Marshall, Alfred. 1920. *Principles of Economics*. 8th ed. London: Macmillan and Co., Ltd.
- Mieszkowski, Peter, and Edwin S. Mills. 1993. "The Causes of Metropolitan Suburbanization." *The Journal of Economic Perspectives* 7 (3): 135–47.
- Pagan, Adrian, and Aman Ullah. 1999. *Nonparametric Econometrics*. Cambridge, UK: Cambridge University Press.
- Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Raymond Owens III. 2005. "Firm Fragmentation and Urban Patterns." Federal Reserve Bank of Richmond Working Paper No. 05-03.
- Silverman, B. W. 1986. *Density Estimation*. London: Chapman and Hall.
- Song, Yan, and Gerritt-Jan Knaap. 2004. "Measuring Urban Form: Is Portland Winning the War on Sprawl?" *Journal of the American Planning Association* 70 (2): 210–25.
- U.S. Bureau of the Census. "Number of Inhabitants: United States Summary." Washington, DC: U.S. Government Printing Office 1941, 1952, 1961, 1971, and 1981.
- U.S. Bureau of the Census. 1994. *Geographic Areas Reference Manual*. Available online at <http://www.census.gov/geo/www/garm.html> (accessed September 4, 2007).

Currency Quality and Changes in the Behavior of Depository Institutions

Hubert P. Janicki, Nashat F. Moin, Andrea L. Waddle,
and Alexander L. Wolman

The Federal Reserve System distributes currency to and accepts deposits from Depository Institutions (DIs). In addition, the Federal Reserve maintains the quality level of currency in circulation by inspecting all deposited notes. Notes that meet minimum quality requirements (fit notes) are bundled to be reentered into circulation while old and damaged notes are destroyed (shredded) and replaced by newly printed notes.

Between July 2006 and July 2007, the Federal Reserve implemented a Currency Recirculation Policy for \$10 and \$20 notes. Under the new policy, Reserve Banks will generally charge DIs a fee on the value of deposits that are subsequently withdrawn by DIs within the same week. In addition, under certain conditions the policy allows DIs to treat currency in their own vaults as reserves with the Fed. It is reasonable to expect that the policy change will result in DIs depositing a smaller fraction of notes with the Fed. While the policy is aimed at decreasing the costs to society of currency provision, it may also lead to deterioration of the quality of notes in circulation since notes that are deposited less often are inspected less often.

This article analyzes the interaction between deposit behavior of DIs and the shred decision of the Fed in determining the quality distribution of currency. For a given decrease in the rate of DIs' note deposits with the Fed, absent any change in the Fed's shred decision, what effect would there be on the quality

■ The authors are grateful to Barbara Bennett, Shaun Ferrari, Juan Carlos Hatchondo, Chris Herrington, Jaelyn Hodges, Larry Hull, Andy McAllister, David Vairo, John Walter, and John Weinberg for their input. The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. Correspondence should be directed to alexander.wolman@rich.frb.org.

distribution of currency in circulation? What kind of changes in the shred criteria would restore the original quality distribution?

To answer these questions, we use the model developed by Lacker and Wolman (1997).¹ In the model, the evolution of the currency quality distribution over time is governed by (i) a quality transition matrix that describes the probabilistic deterioration of notes from one period to the next, (ii) DIs' deposit probabilities for notes at each quality level, (iii) the Fed's shred decision for notes at each quality level, (iv) the quality distribution of new notes, and (v) the growth rate of currency.

We estimate three versions of the model for both \$5 and \$10 notes. We have not estimated the model for \$20 notes because they were redesigned recently, and the new notes were introduced in October 2003. The transition from old to new notes makes our estimation procedure impractical; we discuss this further in the Conclusion.² Although the policy affects \$10 and \$20 notes only, we also estimate the model for \$5 notes because the policy change initially proposed in 2003 included \$5 notes. (It is possible that at some point the recirculation policy might be expanded to cover that denomination.) Also, it is likely that the reduced deposits of \$10 and \$20 notes may induce DIs to change the frequency of transporting notes to the Fed and, hence, affect the deposit rate of other denominations. The model predicts roughly comparable results for both denominations.

In each version of our model, we choose parameters so that the model approximates the age and quality distributions of U.S. currency deposited at the Fed. For each estimated model, we describe the deterioration of currency quality following decreases in DI deposit rates of 20 and 40 percent, and we provide examples of Fed policy changes that would counteract that deterioration. As described in more detail below, we view a 40 percent decrease in deposit rates as an upper bound on the change induced by the recirculation policy.

According to the model(s), a 20 percent decrease in the DI deposit rate would eventually result in an increase in the number of poor quality (unfit) notes of between 0.8 and 2.5 percentage points. While this range corresponds to different specifications of the model, not to a statistical confidence interval, it should be interpreted as indicating the range of uncertainty about our results. For \$10 notes, very small changes in shred policy succeed in preventing a significant increase in the fraction of unfit notes.³ Slightly larger changes in

¹ The Appendix to Lacker (1993) contains a simpler model of currency quality that shares some basic features with the model here.

² New \$10 notes were introduced in March 2006 and new \$5 notes are expected to be introduced in 2008; our data were collected in 2004 and early 2006.

³ We view "fit notes" as referring to any notes that meet a fixed quality standard determined by the Federal Reserve. Prior to a decrease in deposit rates, a fit note is synonymous with a note that meets the Fed's quality threshold for not shredding. If the Fed adjusts its shred policy in

shred policy are required to keep the fraction of unfit \$5 notes from increasing in response to a 20 percent lower deposit rate. Naturally, a 40 percent decrease in deposit rates would cause a larger increase in the number of unfit notes, although the greatest increase we find is still less than 6 percentage points. And even in that case there are straightforward changes in shred policy that would be effective in restoring the level of currency quality.

1. INSTITUTIONAL BACKGROUND

Federal Reserve Banks issue new and fit used notes to DIs and destroy previously circulated notes of poor quality. In order to maintain the quality level of currency in circulation, the Fed uses machines to inspect currency notes deposited by DIs at Federal Reserve currency processing offices. These machines inspect each note to confirm its denomination and authenticity, and measure its quality level on many dimensions. The dimensions that are measured include soil level, tears, graffiti or marks, and length and width of the currency notes. Fit notes are those that pass the threshold quality level on all dimensions. Once sorted, the fit notes are bundled and then recirculated when DIs request currency from the Reserve Banks. To replace destroyed notes and accommodate growth in currency demand, the Federal Reserve orders new notes from the Bureau of Engraving and Printing (B.E.P.) of the U.S. Department of Treasury. The Fed purchases the notes from B.E.P. at the cost of production.⁴ In 2006, the Federal Reserve ordered 8.5 billion new notes from the B.E.P., at a cost of \$471.2 million (Board of Governors of the Federal Reserve System 2006a)—approximately 5.5 cents per note.

In 2006, the Federal Reserve took in deposits of 38 billion notes, paid out 39 billion notes, and destroyed 7 billion notes (Federal Reserve Bank of San Francisco 2006). Of the 19.9 million pounds of notes destroyed every year, approximately 48 percent are \$1 notes, which have a life expectancy of about 21 months. The \$5, \$10, and \$20 denominations last roughly 16, 18, and 24 months, respectively (Bureau of Engraving and Printing 2007). Each day of 2005, the Federal Reserve's largest cash operation, in East Rutherford, New Jersey, destroyed approximately 5.2 million notes, worth \$95 million (Federal Reserve Bank of New York 2006).

response to a decrease in deposit rates, then it will shred some notes that were fit according to this fixed standard.

⁴ Thus, seigniorage for notes accrues initially to the Federal Reserve. In contrast, the Fed purchases coins from the U.S. Mint (a part of the Department of Treasury) at face value, so that seigniorage for coins accrues directly to the Treasury.

Costs and Benefits of Currency Processing and Currency Quality

The Federal Reserve's operating costs for currency processing in 2006 were \$319 million (Federal Reserve Bank of San Francisco 2006). DIs benefit from the Fed's currency processing services in at least two ways. First, the Federal Reserve ships out only fit currency, whereas DIs accumulate a mixture of fit and unfit currency; to the extent that DIs' customers—and their ATMs—demand fit currency, DIs benefit from the Fed's sorting of currency. Second, while DIs need to hold currency to meet their customers' withdrawals, they also incur costs by holding inventories of currency in their vaults. Currency inventories take up valuable space and require expenditures on security systems; in addition, currency in the vault is "idle," whereas currency deposited with the Fed is eligible to be lent out in the federal funds market at a positive nominal interest rate. Thus, the Fed's currency processing services amount to an inventory management service for DIs. The benefits DIs accrue from currency processing may not coincide exactly with the benefits to society. On one hand, positive nominal interest rates make the inventory-management benefit to DIs of currency processing exceed the social benefit (Friedman 1969). On the other hand, the social benefits of improved currency quality may exceed the quality benefits that accrue to DIs: for example, maintaining high currency quality may deter counterfeiting by making counterfeit notes easier to detect (Klein, Gadbois, and Christie 2004). On net, it seems unlikely that the social benefit of currency processing greatly (if at all) exceeds the private benefit. This implies that it would be optimal for DIs to face some positive price for currency processing. Lacker (1993) discusses in detail the policy question of whether the Federal Reserve should subsidize DIs' use of currency.

Historically, the Federal Reserve did not charge DIs for processing currency deposits and withdrawals.⁵ Policy did prohibit a DI's office from cross-shipping currency; cross-shipping is defined as depositing fit currency with the Fed and withdrawing currency from the Fed within the same five-day period. However, as explained in the Federal Reserve Board's request for comments that introduced the proposed recirculation policy (Board of Governors of the Federal Reserve System 2003a), the restriction on cross-shipping was not practical to enforce. Thus, overall the Federal Reserve cash services policy clearly subsidized DIs' use of currency.

Policy Revision

By 2003, the Federal Reserve had come to view existing policy as leading DIs to overuse the Fed's currency processing services (Board of Governors of the

⁵ Note, however, that DIs do pay for transporting currency between their own offices and Federal Reserve offices.

Federal Reserve System 2003b). Factors contributing to this situation included an increase in the number of ATM machines and a decrease in the magnitude of required reserves. The former likely increased the value of the Fed's sorting services, and the latter meant that for a given flow of currency deposits and withdrawals by the DIs' customers, there would be greater demand by DIs to transform vault cash into reserves with the Fed—which requires utilizing the Fed's processing services. In October 2003, the Federal Reserve proposed and requested comments on changes to its cash services policy, aimed at reducing DIs' overuse of the Fed's processing services (Board of Governors of the Federal Reserve System 2003a). In March of 2006, a modified version of the proposal was adopted as the Currency Recirculation Policy (Board of Governors of the Federal Reserve System 2006b).

The Recirculation Policy has two components, both of which cover only \$10 and \$20 denominations. The first component is a custodial inventory program. This program enables qualified DIs to hold currency at the DI's secured facility while transferring it to the Reserve Bank's ledger—thus making the funds available for lending to other institutions but avoiding both the transportation cost and the Fed's processing cost. DIs must apply to be in the custodial inventory program. One criterion for qualifying is that a DI must demonstrate that it can recirculate a minimum of 200 bundles (of 1,000 notes each) of \$10 and \$20 notes per week in the Reserve Bank zone. The policy's second component is a fee of approximately \$5 per bundle of cross-shipped currency. While this new policy is aimed at reducing the social costs incurred because of cross-shipping currency, absent changes in shred policy it is likely to lower the quality of currency in circulation through reduced deposits and thus reduced shredding of unfit currency.⁶ The primary concerns of our study are the effect on currency quality of the anticipated decrease in deposit rates, and the measures the Fed can take to offset that decrease in quality. To address these issues we construct a model of currency quality. We assume that shredding policy is aimed at restoring or maintaining the original quality distribution. If the cost of maintaining quality at current levels exceeds the social benefits of doing so, it would be optimal to let the quality of currency deteriorate somewhat.

2. THE MODEL

The model applies to one denomination of currency.⁷ Time is discrete, and a time period should be thought of as a month. For the purposes of this

⁶ Federal Reserve Banks have estimated that over 10 years, the recirculation policy could reduce their currency processing costs by a present value of \$250 million. Taking into account increased DI costs, the corresponding societal benefit is estimated at \$140 million (Board of Governors of the Federal Reserve System 2006b).

⁷ By changing the parameters appropriately, it can be applied *separately* to more than one denomination; indeed we will do just that.

study, there are three major dimensions to currency quality: soil level front (we will use the shorthand “soil level” or SLF), ink wear worst front (“ink wear” or IWWF), and graffiti worst front (“graffiti” or GWF). There are also at least 18 minor dimensions to currency: soil level back, graffiti total front, etc. For a given denomination, we have separate models for each major dimension.⁸ Those models describe, for example, how the distribution over soil level evolves over time. For each of those models, however, we use data on the other dimensions to more accurately describe the probability that a note of a particular major-dimension quality level will be shredded.⁹

The basic structure of the model is as follows. At the beginning of each period, banks deposit currency with the Fed; their deposit decision may be a function of quality in the major dimension (that is, banks may sort for fitness). The Fed processes deposited notes, shredding those deemed unfit and recirculating the rest at the end of the period. The shred decision is based on quality level in whatever major dimension the model is specified. However, notes that are fit according to their quality level in the major dimension are nonetheless shredded with positive probability; this is to account for the fact that they may be unfit along one of the other (major or minor) dimensions in which the model is not specified. The stock of currency is assumed to grow at a constant rate. Banks make withdrawals from the Fed at the end of the period but these are not specified explicitly; instead, withdrawals can be thought of as a residual that more than offsets deposits in order to make the quantity of currency grow at the specified rate. In order to accommodate growth in currency and replace shredded notes, the Fed must introduce newly printed notes. Meanwhile, the notes that were not deposited with the Fed deteriorate in quality stochastically. The quality of notes in circulation at the end of a period, and thus at the beginning of the next period, is determined by the quality of notes that have remained in circulation and the quality of notes withdrawn from the Fed.

Formal Specification of the Model

Time is indexed by a subscript $t = 0, 1, 2, \dots$. Soil level can take on values $0, 1, 2, \dots, n_s - 1$; ink wear can take on values $0, 1, 2, \dots, n_i - 1$; and graffiti can take on values $0, 1, 2, \dots, n_g - 1$; in general, larger numbers denote poorer quality.¹⁰ We will use q to denote a particular (arbitrary) quality level.

⁸ The models for the three major dimensions are truly separate, in that they will yield different predictions.

⁹ As mentioned earlier, the model was first developed in Lacker and Wolman (1997). That article studied a different policy question, namely expanding the dimensions of quality measurements to include limpness.

¹⁰ The exception is soil level zero, which is assumed to describe currency that has been laundered (i.e., has been through a washing machine) and is deemed unfit.

For the DIs' deposit decision, the vector ρ contains in its q^{th} element the probability that a DI will deposit a note conditional on that note being of quality level q . The vector ρ has length Q , where $Q = n_s$ or n_i or n_g , depending on the particular model in question. For the Fed's fitness criteria, the $Q \times 1$ vector α contains in its q^{th} element the probability that a deposited note of quality q is put back into circulation. If the model were specified in terms of every quality characteristic—so that Q were a huge number describing every possible combination of “soil level front,” “soil level back,” etc.—then the elements of α would each be zero or one and they would be known parameters, taken from the machine settings. Because the model is specified in terms of only one characteristic, the elements of α that would be one according to q are adjusted downward to account for the fact that some quality- q notes are unfit according to other dimensions of quality. The values of α must then be estimated, and we describe in Section 4 how they are estimated.

The net growth rate of the quantity of currency is γ ; that is, if the quantity of currency is M in period t , then it is $(1 + \gamma)M$ in period $(t + 1)$. The $Q \times 1$ vector g describes the distribution of new notes; its q^{th} element is the probability that a newly printed note is of quality q .¹¹ The deterioration of non-deposited notes is described by the $Q \times Q$ matrix π ; the row- r column- c element of π is the probability that a non-deposited note will become quality r next period, conditional on it being quality- c this period.¹² Note that each column of π sums to one, because any column q contains the probabilities of all possible transitions from quality level q .

The model's endogenous variables are the numbers of notes of different quality levels, i.e., the quality distribution of currency. At the beginning of period t , the $Q \times 1$ vector m_t contains in its q^{th} element the number of notes in circulation of quality q . The total number of notes in circulation is $M_t = \sum_{q=1}^Q m_{q,t}$, where $m_{q,t}$ denotes the q^{th} element of the vector m_t .

¹¹ We allow for new notes to have some variation in quality. However, by choosing g appropriately we can impose the highest quality level for all new notes.

¹² We assume that the number of notes is sufficiently large that the probability that a quality c note makes a transition to quality r is the same as the fraction of type c notes that make the transition to type r . That is, the law of large numbers applies.

Combining these ingredients, the number of notes at each quality level evolves as follows:

$$\begin{aligned}
 m_{t+1} = & \pi \cdot \left((1 - \rho) \odot m_t \right) \\
 & + \alpha \odot \rho \odot m_t \\
 & + \left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q m_{q,t} \right) g \\
 & + (\gamma M_t) g .
 \end{aligned} \tag{1}$$

The symbol \odot denotes element-by-element multiplication of vectors or matrices.¹³

Equation (1) is the model, although we will rewrite it in terms of fractions of notes instead of numbers of notes. On the left-hand side, m_{t+1} contains the number of notes at each quality level at the beginning of period $t + 1$. The right-hand side describes how m_{t+1} is determined from the interaction of m_t (the number of notes at each quality level at the beginning of period t) with the model's parameters. The first term on the right-hand side is

$$\pi \cdot \left((1 - \rho) \odot m_t \right) . \tag{2}$$

This term accounts for the fractions $(1 - \rho)$ of notes at each quality level that are not deposited. These notes deteriorate according to the matrix π , and thus the first term is a $Q \times 1$ vector containing in its q^{th} element the number of circulating notes that were not deposited in period t and that begin period $t + 1$ with quality q . If banks were to sort for fitness, then the notes that remain in circulation and deteriorate during the period would be relatively high quality notes, otherwise they would be a random sample of notes. The matrix π has Q^2 elements; assigning numbers to those elements will be the key difficulty we face in choosing parameters for the model.

The second term is

$$\alpha \odot \rho \odot m_t . \tag{3}$$

This term accounts for the fractions $\alpha \odot \rho$ of notes at each quality level that are deposited and not shredded—that is, $\alpha \odot \rho \odot m_t$ comprises the deposited notes at each quality level that are fit and will be put back into circulation at the end of period t . If banks were to sort for fitness in a manner consistent

¹³ For example, if $a = [1, 2]$ and $b = [3, 4]$, then $a \odot b = [3, 8]$.

with the Fed’s fitness definitions, and if banks possessed enough unfit notes to meet their deposit needs, then this term would disappear—all deposited notes would be shredded.

The third term, $\left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q m_{q,t}\right) g$, represents replacement of shredded notes. The object in parentheses is the number of unfit notes that are processed (and shredded) each period. Multiplying by the distribution of new notes g gives the vector of new notes at each quality level that are added to circulation at the end of period t to replace shredded notes.

The fourth term, $(\gamma M_t) g$, represents growth in the quantity of currency. The number of new notes added to circulation to accommodate growth (as opposed to shredding) is γM_t , and the distribution of new notes is g , so this term is a vector containing the numbers of new notes at each quality level added to circulation at the end of period t to accommodate growth.

We noted above that withdrawals are not treated explicitly in the model. The quantity of withdrawals can, however, be calculated. The number of notes withdrawn in period t must be equal to the sum of deposits and currency growth. That is, withdrawals equal

$$\left(\sum_{q=1}^Q \rho_q m_{q,t}\right) + \gamma M_t. \tag{4}$$

Note that the model does not incorporate currency inventories at the Fed. New notes materialize as needed, and fit notes deposited at the Fed are recirculated at the end of the period.

The evolution of currency quality over time is determined entirely by equation (1). Given a vector m_t describing the distribution of currency quality at the beginning of any period t , equation (1) determines the vector m_{t+1} describing the distribution of currency quality at the beginning of period $t + 1$. The law of motion is determined by the parameters π , ρ , g , γ , and α .¹⁴

The Model in Terms of Fractions of Notes

The model has been expressed in terms of the numbers of notes at each quality level. To express the model in terms of fractions of notes at each quality level, we first define f_t to be the vector of fractions, that is the $Q \times 1$ vector of numbers of notes at each quality level divided by the total number of notes:

$$f_t \equiv \left(\frac{1}{M_t}\right) \cdot m_t. \tag{5}$$

¹⁴ We have written the model as if all parameters are constant over time. We maintain that assumption for the quantitative results described in this report. The model remains valid if the parameters change over time, although estimation becomes more challenging.

Likewise, the fraction of notes at a particular quality level is

$$f_{q,t} \equiv \left(\frac{1}{M_t} \right) \cdot m_{q,t}. \quad (6)$$

Note that the elements of f_t sum to one, because $M_t = \sum_{q=1}^Q m_{q,t}$. Using these definitions, we can rewrite the model (1) by dividing both sides by M_t and recalling that $M_{t+1} = (1 + \gamma) M_t$:

$$\begin{aligned} (1 + \gamma) f_{t+1} &= \pi \cdot ((1 - \rho) \odot f_t) \\ &\quad + \alpha \odot \rho \odot f_t \\ &\quad + \left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q f_{q,t} \right) g \\ &\quad + \gamma g. \end{aligned} \quad (7)$$

With this formulation it will be straightforward to study the model's steady state with currency growth.

The Steady-State Distribution of Notes

Under certain conditions, the distribution of currency quality converges to a steady state with the distribution f_t , which is constant over time (see, for example, Stokely, Lucas, with Prescott, chap.11). Assuming that a unique steady-state distribution exists, we will denote it by f^* . In the steady state, the law of motion (7) becomes

$$\begin{aligned} (1 + \gamma) f^* &= \pi \cdot ((1 - \rho) \odot f^*) + \alpha \odot \rho \odot f^* \\ &\quad + \left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q f_q^* \right) g + \gamma g. \end{aligned} \quad (8)$$

Our method of choosing the model's parameters will require us to compute the steady-state distribution—we will assume that our data are generated in a steady-state situation. One way to compute the steady state is to simply iterate on (7) from some arbitrary initial distribution f_0 and hope that the iterations converge. If they converge, we have found the steady state. Alternatively, we can use matrix algebra to solve directly for the steady state from (8). Ultimately, we want to rewrite (8) in the form

$$\Gamma \cdot f^* = \gamma g, \quad (9)$$

where Γ is a $Q \times Q$ matrix. If we can rewrite (8) in this way, then the steady-state distribution is $f^* = \Gamma^{-1} \cdot (\gamma g)$. The first step is to note that for any $Q \times 1$ vector v , we have $v \odot f^* = \text{diag}(v) \cdot f^*$, where $\text{diag}(v)$ denotes the $Q \times Q$ matrix with the vector v on the diagonal and zeros, elsewhere. Using this fact,

we can rewrite (8) as

$$(1 + \gamma) f^* = \pi \cdot \text{diag}(1 - \rho) \cdot f^* + \text{diag}(\alpha \odot \rho) \cdot f^* + \left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q f_q^* \right) g + \gamma g. \tag{10}$$

Next, note that the scalar $\left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q f_q^* \right)$ can be rewritten as $((1 - \alpha) \odot \rho)' f^*$, where “'” denotes transpose. Using this fact, we have

$$\left(\sum_{q=1}^Q (1 - \alpha_q) \rho_q f_q^* \right) g = \underbrace{g}_{Q \times 1} \underbrace{((1 - \alpha) \odot \rho)' f^*}_{1 \times 1}. \tag{11}$$

Now we can express (8) in the same form as (9), $\Gamma \cdot f^* = \gamma g$, where

$$\Gamma \equiv \left[(1 + \gamma) I - \pi \cdot \text{diag}(1 - \rho) - \text{diag}(\alpha \odot \rho) - g \underbrace{((1 - \alpha) \odot \rho)'}_{Q \times 1} \right]^{-1}. \tag{12}$$

Thus, the steady state can be computed directly as

$$f^* = \Gamma^{-1} \cdot (\gamma g).$$

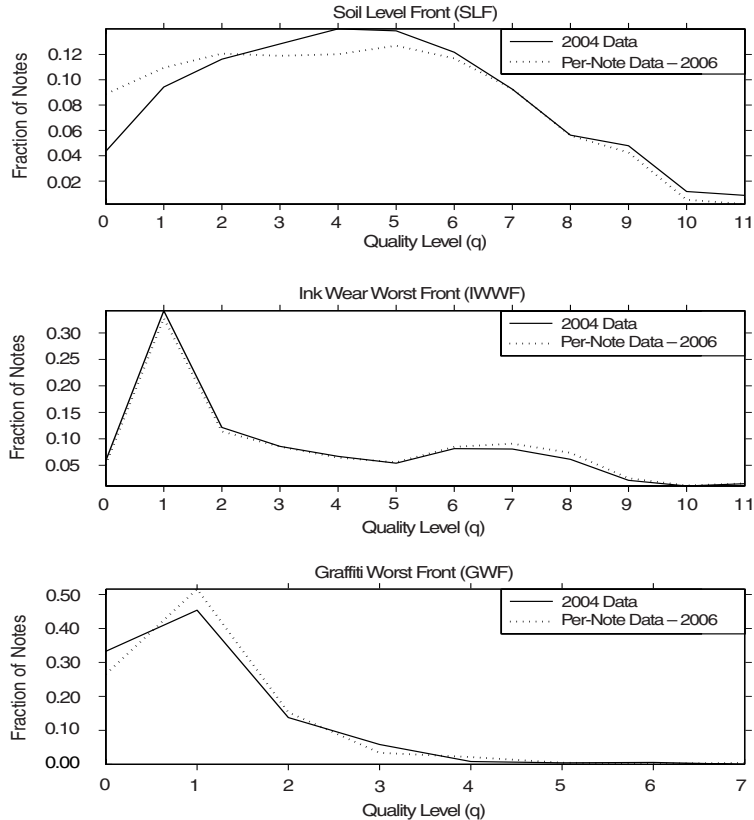
The steady-state distribution f^* contains in its q^{th} element the fraction of notes with quality q , corresponding to a particular measurement of soil level, graffiti or ink wear. Thus, f^* can be thought of as the *marginal distribution* over soil level, graffiti or ink wear. When comparing the model to data, we will use the marginal distributions for each major quality dimension and the distribution of notes by age. We use the age distribution because the quality distribution alone puts few restrictions on the matrix π : we can match a given quality distribution with many π matrices, each implying a different age distribution.

The Appendix contains a detailed description of how to calculate the steady-state age distribution of notes. For now, we simply state the notation: $h_{q,k}$ denotes the fraction of notes that are quality q and age k , and h_k denotes the Q by 1 vector of age k notes, the q^{th} element of which is $h_{q,k}$.

3. THE DATA

The model’s predictions will depend on the numerical values we assign to the matrix π describing deterioration of notes, the vector ρ of deposit probabilities, the vector α of shred probabilities, the quality distribution of new notes g , and the currency growth rate γ . This section describes the basic data whose features we attempt to match in choosing the model’s parameters.

The ideal data set for our purposes would be one with a time series of observations on a large number of currency notes, with observations each month on the quality of every note. Data of this sort would allow for nearly

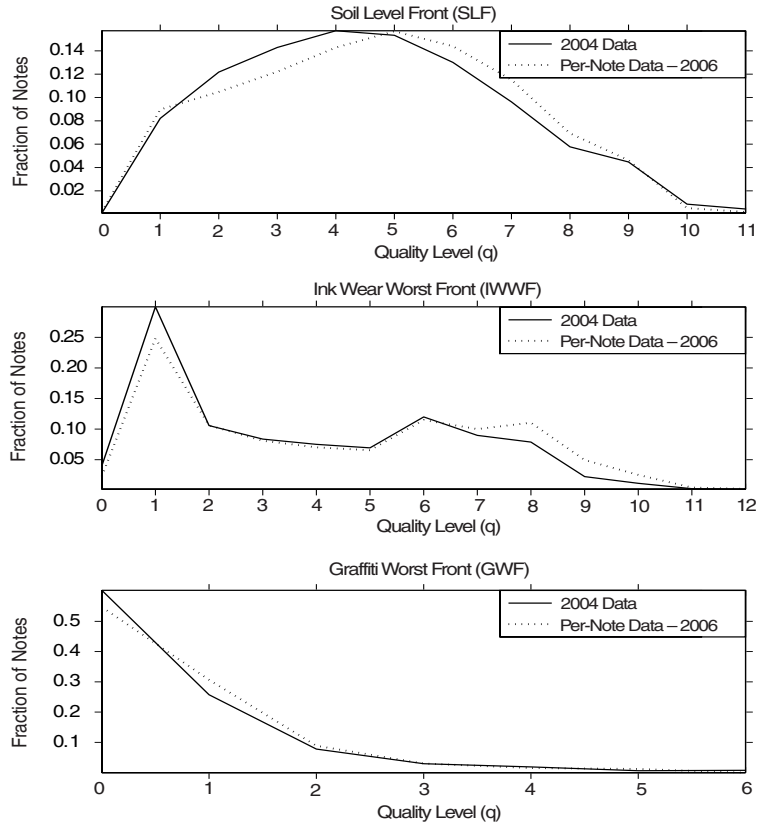
Figure 1 Marginal Quality Distributions, \$5 Notes

direct measurement of the matrix π . Of course such data does not exist, and probably the only way it could exist would be if individual notes had built-in sensors and transmitters. Without such data, we need to estimate the parameters of π . We use two data sets for this purpose. One data set describes the marginal quality distributions only and has extremely broad coverage. The other data set is at the level of individual notes, and contains age of notes as well as quality. It has more limited coverage.

Large Data Set Describing Marginal Distributions

The large data set comprises fitness data for the entire Federal Reserve System for the months of January 2004 and May 2004, provided by the Currency Technology Office (CTO) at the Federal Reserve Bank of Richmond. This data characterizes the marginal quality distributions for more than two-and-

Figure 2 Marginal Quality Distributions, \$10 Notes



a-half billion notes. The data are at the level of office location, date, shift, supervisor, and denomination. For a particular denomination, we assume that summing these data over all dates, supervisors, and shifts generates a precise estimate of the steady-state marginal distribution over each quality level. Figures 1 and 2 plot the marginal distributions over soil level, ink wear and graffiti for the combined January and May 2004 data, for \$5 and \$10 notes (solid lines).¹⁵

The raw data have 26 quality levels for each category. However, for many quality levels there are very few notes, and for speed of computation it is advantageous to decrease the number of quality levels. For each denomination

¹⁵The same data set covers \$20 notes, but as described in the Conclusion, our limited analysis of the 20s has not used this data.

and each category (e.g., SLF) we have, therefore, combined multiple quality levels into one. For example, our new soil level zero for the \$10 notes includes all notes with soil levels zero through 2 in the data. Table 1 contains comprehensive information about how we combine quality levels. Boxes around multiple quality levels indicate that we have combined them, and the columns labeled “ q ” contain the quality level numbers corresponding to our smaller set of quality levels. After combining in this way, we are left with between 7 and 13 quality levels for each denomination and category. For each denomination and each dimension, there are three unfit quality levels. For example, for the \$5 notes SLF, quality levels 9, 10, and 11 are unfit.

Per-Note Data

In addition to the comprehensive data set describing marginal distributions, we use per-note data sets covering approximately 45,000 notes each of \$5 notes and \$10 notes. These data were gathered at nine Federal Reserve offices in February and March 2006. For each note, there is information on the date of issue, as well as quality level in at least 21 categories, including SLF, GWF, and IWWF. The dotted lines in Figures 1 and 2 are the marginal quality distributions for SLF, IWWF, and GWF from the per-note data for the \$5 and \$10 notes. There are minor differences relative to the marginal distributions from the large data set, but the broad patterns are the same. This gives us some confidence that the per-note data are representative samples.

Because the note data contain date of issue for each note, we are able to get an estimate of the age distribution of notes. In Figure 3, the jagged dotted line is a smoothed version of the age distribution of unfit notes from the note data. The smoothing method involves taking a three-month moving average. Without smoothing, the age distributions would be extremely choppy. Note that in Figure 3, we plot the age distribution of *unfit notes*. It is the unfit notes with which we are most concerned for this study, and whose age distribution we care most about matching with the model. Unfit notes are those notes whose quality is worse than the shred threshold in any dimension—major or minor.

4. CHOOSING THE MODEL’S PARAMETERS

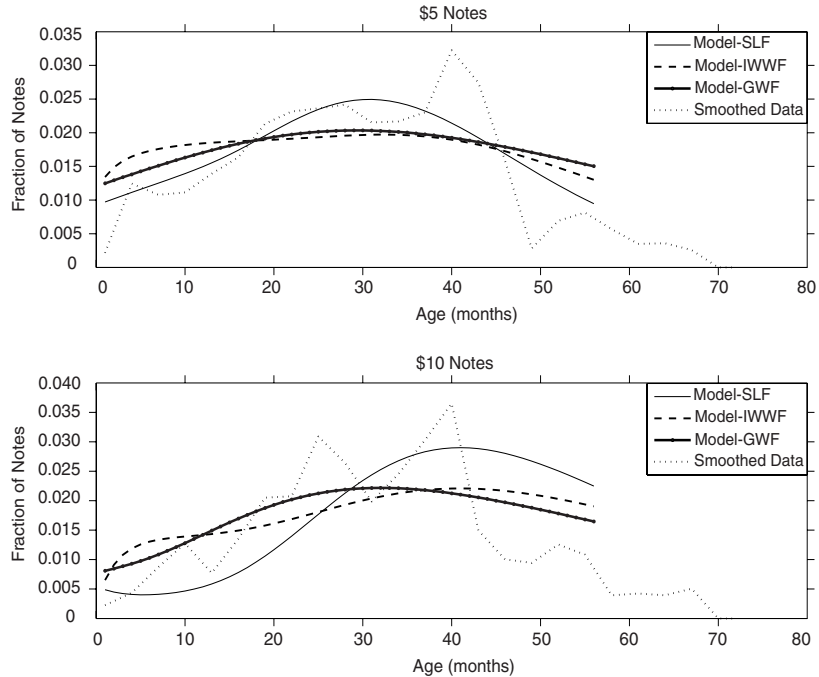
There are $Q^2 + 3Q + 1$ parameters in each model; they comprise the Q^2 elements of π , the $3Q$ elements of α , ρ , and g , and the single parameter γ .¹⁶ Since Q is between 7 and 13, the number of parameters is between 71 and 209. We select the model’s parameters in several stages.¹⁷

¹⁶ Recall that Q is either n_s, n_i , or n_g , depending on the version of the model.

¹⁷ Because our approach to selecting parameters is ad hoc, we hesitate to talk about “estimating the model.” However, in effect that is what we are doing.

Table 1 Marginal Quality Distributions and Combined Quality Levels

Quality Level	\$5 Notes				\$10 Notes							
	SLF	q	IWWF	q	GWF	q	SLF	q	IWWF	q	GWF	q
0	0.000		0.059	0	0.000	0	0.000		0.040	0	0.001	0
1	0.000		0.342	1	0.333	1	0.000	0	0.300	1	0.602	
2	0.000	0	0.121	2	0.454	1	0.001		0.106	2	0.257	1
3	0.004		0.086	3	0.138	2	0.014	1	0.084	3	0.078	2
4	0.039		0.067	4	0.036	3	0.068		0.075	4	0.030	3
5	0.094	1	0.054	5	0.015		0.122	2	0.069	5	0.012	4
6	0.116	2	0.044	6	0.008	4	0.143	3	0.063	6	0.006	
7	0.128	3	0.037		0.005		0.157	4	0.056		0.004	5
8	0.140	4	0.031		0.003	5	0.153	5	0.049	7	0.003	
9	0.139	5	0.027	7	0.002		0.130	6	0.041		0.002	
10	0.122	6	0.023		0.002	6	0.096	7	0.033		0.001	
11	0.093	7	0.019		0.001		0.058	8	0.026	8	0.001	
12	0.056	8	0.017		0.001		0.027		0.019		0.001	
13	0.027		0.014	8	0.001		0.012	9	0.013	9	0.001	
14	0.013	9	0.011		0.001		0.006		0.009		0.000	
15	0.007		0.009		0.000		0.004		0.006		0.000	
16	0.005		0.007	9	0.000		0.003	10	0.003	10	0.000	
17	0.004	10	0.005		0.000		0.002		0.002		0.000	
18	0.003		0.004		0.000	7	0.002		0.001		0.000	6
19	0.002		0.004	10	0.000		0.001		0.001	11	0.000	
20	0.002		0.003		0.000		0.001		0.000		0.000	
21	0.002		0.002		0.000		0.000	11	0.000		0.000	
22	0.001	11	0.002		0.000		0.000		0.000		0.000	
23	0.001		0.002	11	0.000		0.000		0.000	12	0.000	
24	0.001		0.001		0.000		0.000		0.000		0.000	
25	0.000		0.008		0.000		0.000		0.001		0.000	

Figure 3 Age Distribution of Unfit Notes

First, we make some *a priori* assumptions on the transition matrix π that decrease the number of free parameters. Next, we pin down g , α , γ , and ρ based on information from the Federal Reserve System's Currency Technology Office, the Federal Reserve Board, and preliminary analysis of the data. We select the remaining parameters so that the model's steady-state distribution matches the quality and age distributions in Figures 1–3.

At this point, it may be useful to remind the reader where we are: we have specified a model of the evolution of currency quality, and we will now use data from the period before implementation of the currency recirculation policy in order to choose parameters of the model. Once the parameters have been chosen, we will simulate the model under particular assumptions about how DI behavior will change in response to the recirculation policy. The recirculation policy itself is "outside the model"; the model does not address pricing of currency processing by the Fed, and the model does not address (intra-week) cross-shipping because it is specified at a monthly frequency.

A Priori Restrictions on π

We reduce the number of parameters determining π by imposing the restriction that notes never improve in quality, except that soil level may “improve” to zero if a note is laundered (i.e., the note has gone through a washing machine). This restriction means that almost half the elements of π are zeros. For the ink wear and graffiti model, all elements above the main diagonal are zero. For the soil level model, the elements above the main diagonal are zero except in the first row, which may contain nonzero elements in every column to account for the possibility of laundered notes; in the first column, the first row contains a one and all other rows contain zeros, because a laundered note always remains laundered. The numbers of nonzero elements in π are thus $\left(\frac{n_s(n_s+1)}{2} + n_s - 1\right)$, $\frac{n_i(n_i+1)}{2}$ and $\frac{n_g(n_g+1)}{2}$ for the three models. The last restriction we impose on π is an inherent feature of the model: the columns of π must sum to one, and π is a stochastic matrix with each element weakly between zero and one. This adds Q restrictions, subtracting an equal number of parameters.

Choosing α , g , ρ , and γ

The Federal Reserve chooses the definition of fit notes, so there would be no difficulty determining α if the model were specified in terms of all quality dimensions simultaneously; α_q would be one for fit notes and zero for unfit notes. However, since we specify the model in terms of only one dimension, we need to adjust the shred parameter α to reflect the fact that notes may be unfit even though they are fit according to the dimension in which the model is specified. For example, if the model is specified in terms of soil level, a note that is very clean may nonetheless be unfit because of its level of ink wear. We adjust for this possibility as follows, using the soil level example: for each *fit* degree of soil level q , calculate the fraction of notes with soil level q that are unfit according to other dimensions and subtract that fraction from α_q . That calculation is necessarily based on the per-note data, as it requires going beyond marginal distributions. The corrections we make to α are shown in Table 2.

The vector g represents the quality distribution of newly printed notes. Our estimates of g are from the Federal Reserve System’s Currency Technology Office (unpublished data), and these are presented in Table 3. Sorting behavior by DIs is captured by the vector ρ .

Table 2 Corrections to α Vector

q	\$5 Notes				\$10 Notes			
	SLF	q	IWWF	GWF	SLF	q	IWWF	GWF
0	0	0	0.0850	0.0624	0	0	0.0255	0.0374
1	0.0375	1	0.1080	0.1215	0	1	0.0545	0.1142
2	0.0640	2	0.1445	0.2795	2	2	0.0654	0.2294
3	0.0867	3	0.1754	0.5783	3	3	0.0868	0.3392
4	0.1150	4	0.1801	0	4	4	0.0844	0
5	0.1417	5	0.2048	0	5	5	0.0857	0
6	0.1852	6	0.2109	0	6	6	0.0878	0
7	0.2546	7	0.2461	0	7	7	0.1036	
8	0.3658	8	0.2913		8	8	0.1251	
9	0	9	0		9	9	0.1496	
10	0	10	0		10	10	0	
11	0	11	0		11	11	0	
					12	12	0	

Table 3 Quality Distribution of New Notes

\$5 Notes						\$10 Notes					
q	SLF	q	IWWF	q	GWF	q	SLF	q	IWWF	q	GWF
0	0	0	1	0	0.935	0	0	0	1	0	1
1	0.010	1	0	1	0.065	1	0.965	1	0	1	0
2	0.695	2	0	2	0	2	0.035	2	0	2	0
3	0.295	3	0	3	0	3	0	3	0	3	0
4	0	4	0	4	0	4	0	4	0	4	0
5	0	5	0	5	0	5	0	5	0	5	0
6	0	6	0	6	0	6	0	6	0	6	0
7	0	7	0	7	0	7	0	7	0		
8	0	8	0			8	0	8	0		
9	0	9	0			9	0	9	0		
10	0	10	0			10	0	10	0		
11	0	11	0			11	0	11	0		
								12	0		

We assume that DIs do not sort, which implies that all elements of ρ are identical and are equal to the fraction of notes that DIs deposit each period.¹⁸ We set each element of ρ to 0.1165 for the \$5 notes and 0.1322 for the \$10 notes. These numbers are based on data from the Federal Reserve Board (S. Ferrari, pers. comm.). Finally, γ is the growth rate of the stock of currency. We have set the annual growth rate at 1.78 percent for the \$5 notes, and 0.38 percent for the \$10 notes, again based on data from the Federal Reserve Board (S. Ferrari, pers. comm.).

Matching the Quality and Age Data

We select the remaining parameters of the matrix π —for each specification of the model—so that the model’s steady-state distribution matches as closely as possible two features of the data. First, we want to match the marginal quality distribution from the 2004 comprehensive data (Figures 1 and 2, solid line). Second, we want to match the age distribution of unfit notes from the 2006 per-note data (Figure 3). Concretely, we select the parameters of π to minimize a weighted average of (i) the sum of squared deviations between the marginal quality distribution and that predicted by the model, and (ii) the sum of squared deviations between the unfit age distributions and that predicted by

¹⁸ A recent internal Federal Reserve study confirmed that DIs have not been sorting to any appreciable extent, as the quality distribution of currency that the Federal Reserve receives from DIs is close to the quality distribution of currency in circulation (Board of Governors of the Federal Reserve System 2007). However, the recirculation policy—in particular, the fee for cross-shipping *fit* currency—gives DIs an incentive to sort. We address this issue in the Conclusion.

Table 4 π Matrix for \$5 Notes According to GWF

q	0	1	2	3	4	5	6	7
0	0.9469	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
1	0.0531	0.9755	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
2	0.0000	0.0224	0.9647	0.0000	0.0000	0.0000	0.0000	0.0000
3	0.0000	0.0000	0.0353	0.9945	0.0000	0.0000	0.0000	0.0000
4	0.0000	0.0022	0.0000	0.0000	0.8828	0.0000	0.0000	0.0000
5	0.0000	0.0000	0.0000	0.0054	0.1148	0.8294	0.0000	0.0000
6	0.0000	0.0000	0.0000	0.0001	0.0024	0.1706	0.9995	0.0000
7	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0005	1.0000

Notes: The row r , column c element of this matrix is the probability that a note will become quality r next period, conditional on it being of quality c in this period. For example, the probability that a note will be of quality 4 in the next period, given that it is quality 1 in this period is 0.0022, the element in row 4, column 1.

the model.¹⁹ Table 4 contains one example of the π matrix; it is for the GWF model of \$5 notes.

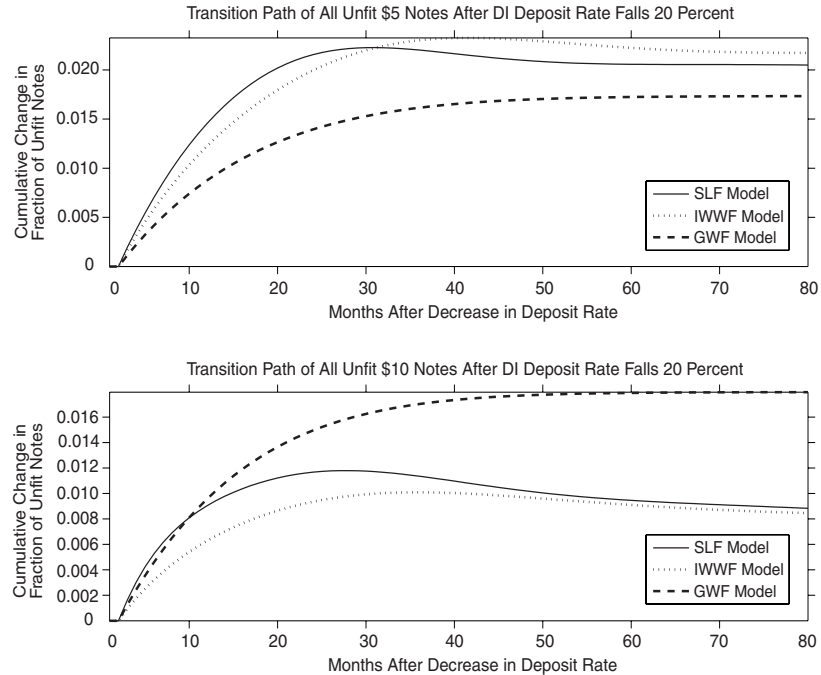
With respect to the marginal quality distributions, we have no trouble matching the data. In all of the model specifications, we match the marginal quality distributions nearly perfectly. The age distributions are a different matter, which perhaps is not surprising given their chopiness in the data—the model wants to make the age distribution of unfit notes smooth. Figure 3 plots the age distributions implied by each specification of the model, along with the age distributions from the data.²⁰ With the exception of the SLF model for \$5 notes, the age distributions implied by the model involve too many unfit notes more than approximately four years old.

5. SIMULATING A CHANGE IN DI BEHAVIOR

Because the response of the quality distribution to a decrease in deposit rates depends on the transition matrix π , the fact that we have multiple models means that we generate a range of responses to a decrease in deposit rates. Figures 4 and 5 plot the time series for the fraction of unfit notes, in response to 20 and 40 percent decreases in DIs' deposit rates, respectively. According to final Currency Recirculation Policy (Board of Governors of the Federal Reserve System 2006b), of the \$10 and \$20 notes processed by the Fed in

¹⁹ We have also experimented with adding to our estimation criterion the fraction of age k notes that are unfit, for $k = 1, 2, \dots$ For moderate weights on this component the results are not materially affected.

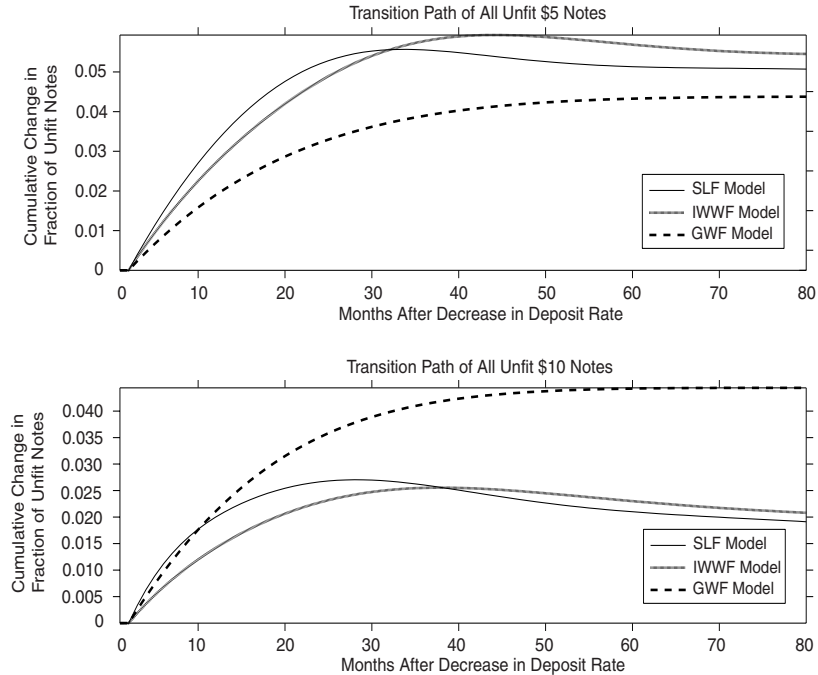
²⁰ In Figure 3, the lines associated with the model stop at 56 months because we did not attempt to match the age distribution beyond 56 months.

Figure 4 Response to 20 Percent Decrease in Deposit Rate

2004, 40.4 percent were cross-shipped. Thus, a 40 percent decrease in deposits corresponds to DIs ceasing entirely to cross-ship. This seems unlikely, so we view the 40 percent number as an upper bound on the effect of the recirculation policy. In addition, cross-shipping is likely more important for \$20 notes than \$10 notes, because of the necessity of having crisp (fit) \$20 notes in ATM machines. Since the DIs always receive fit notes from the Federal Reserve System, a larger volume of \$20 notes are cross-shipped than any other denomination.²¹ Thus, the 40.4 percent upper bound for \$10 notes and \$20 notes combined is higher than the upper bound for the \$10 notes or \$5 notes.

Each line in Figures 4 and 5 represents the transition path for the fraction of unfit notes for a different major dimension model (soil level, ink wear, graffiti). In response to a 20 percent decrease in the deposit rate, the models predict a long-run increase in the *fraction* of unfit notes of between 0.017 and 0.025 for the \$5 notes (i.e., around two percentage points), and between 0.008 and 0.018 for the \$10 notes. In our large data sets, the total fractions


²¹ In 2005, the volume of \$5, \$10, and \$20 notes that were cross-shipped were 12.7 percent, 9.0 percent, and 78.3 percent, respectively.

Figure 5 Response to 40 Percent Decrease in Deposit Rate

of unfit notes are 0.173 for the \$5 notes and 0.150 for the \$10 notes. Note that the model that provides the best fit to the age distribution (\$5 SLF) is also the model that predicts the largest increase in the fraction of unfit notes, 0.025. Not surprisingly, a 40 percent decrease in deposit rates generates a larger increase in the fraction of unfit notes—between 0.044 and 0.055 for the \$5 notes and between 0.019 and 0.044 for the \$10 notes.

Figures 6, 7, 8, and 9 provide a different perspective on the effects of a decrease in deposit rates. These figures plot on the same panel the initial steady-state quality distribution (prior to the drop in deposit rates) and the new steady-state quality distribution corresponding to the lower deposit rate. For the 20 percent experiment (Figures 6 and 7), the long-run effects on quality are generally small, reinforcing the message of Figure 4. There are, however, certain quality levels that are strongly affected. For example, the fraction of \$10 notes at soil level 6 (in Figure 7) eventually rises from 0.13 to 0.1832 in response to the 20 percent drop in deposits. For the 40 percent experiment, things look somewhat more dramatic: for example, the fraction of \$10 notes at soil level 6 increases from 0.13 to 0.27 (in Figure 9). To put this change in perspective though, Table 2 tells us that only 6.2 percent of the level 6 SLF

Figure 6 Effect of 20 Percent Deposit Rate Decrease on Quality Distributions of \$5 Notes



\$10 notes are unfit, so the big increase in notes at that level (which is still fit according to SLF) brings with it an increase of less than one percentage point in unfit notes. Recall that the change in total fraction of unfit notes is shown in Figures 4 and 5.

If the Fed wished to offset the quality deterioration caused by a decrease in deposit rates, a natural policy would be to shred notes of higher quality. Table 5 displays scenarios for fraction of notes to shred at each quality level in order to maintain the fraction of unfit notes at its old steady-state level. For example, if deposit rates fall 20 percent, our SLF model for \$5 notes implies that shredding all notes in the worst-fit category and shredding 35 percent of notes in the second worst-fit category would counteract the deposit decrease, leaving the fraction of notes unchanged. The columns in this table should be read independently, as they each apply to distinct models. In other words, the column labeled \$5 SLF provides a policy change for SLF that is predicted to bring about a stable fraction of unfit notes; no changes are made


Table 5 Policy Response to Offset Effect of Deposit Rate Decrease

20 Percent Decrease in Deposits: Fraction of Notes to Shred							
\$5 Notes				\$10 Notes			
q	SLF	q	IWWF	q	SLF	q	IWWF
0	1	0	0	0	1	0	0
1	0	1	0	1	0	1	0
2	0	2	0	2	0	2	0
3	0	3	0	3	0	3	0
4	0	4	0	4	0	4	0
5	0	5	0	5	0	5	0
6	0	6	0	6	0	6	0
7	0.3512	7	0.0255	7	0	7	0
8	1	8	1	8	0.545	8	0
9	1	9	1	9	1	9	0.6845
10	1	10	1	10	1	10	1
11	1	11	1	11	1	11	1
						12	1

40 Percent Decrease in Deposits: Fraction of Notes to Shred							
\$5 Notes				\$10 Notes			
q	SLF	q	IWWF	q	SLF	q	IWWF
0	1	0	0	0	1	0	0
1	0	1	0	1	0	1	0
2	0	2	0	2	0	2	0
3	0	3	0	3	0	3	0.22
4	0	4	0	4	0	4	1
5	0.198	5	0	5	0	5	1
6	1	6	0.48	6	0	6	1
7	1	7	1	7	0.125	7	1
8	1	8	1	8	1	8	1
9	1	9	1	9	1	9	1
10	1	10	1	10	1	10	1
11	1	11	1	11	1	11	1
						12	1

to shred thresholds for other dimensions. Note that we have omitted GWF from the analysis in Table 1; we were not successful in finding policies that counteracted the quality decline by changing the shred policy for GWF. In order to counteract the effects of a 40 percent decrease in deposits, Reserve Banks would have to shred currency at significantly higher quality levels, depending on the particular model specification. In the most extreme case, which is the IWWF model for \$10 notes, the worst six levels of fit notes would have to be shredded (quality levels four through nine), and 22 percent of notes at quality level 3 would have to be shredded to prevent overall quality from deteriorating. Recall, however, that the 40 percent decrease in deposit rates

Figure 7 Effect of 20 Percent Deposit Rate Decrease on Quality Distributions of \$10 Notes

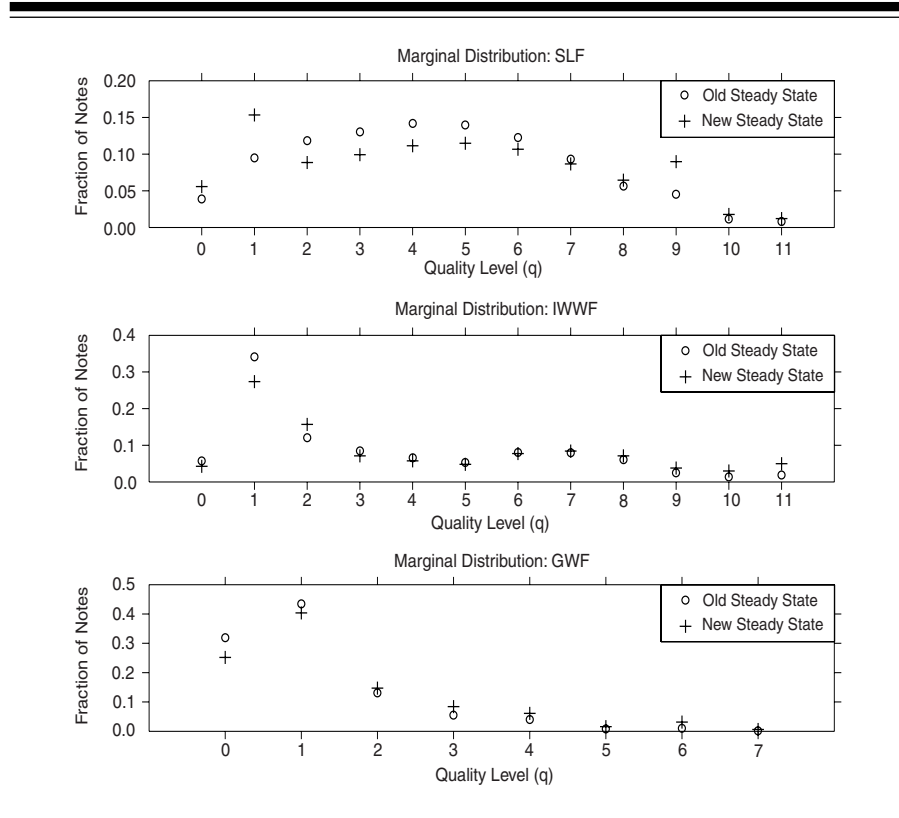


represents an upper bound on how we expect DIs to change their behavior in response to the recirculation policy.

6. CONCLUSION

The quality of currency in circulation is an important policy objective for the Federal Reserve. Changes in the behavior of depository institutions, whether caused by Fed policy or by independent factors, can have implications for the evolution of currency quality. Currently the Fed is implementing a recirculation policy, which is expected to cause changes in the behavior of DIs and, therefore, affect currency quality. The mechanical model of currency quality in this article can be used to study the effects of changes in DI behavior and changes in Fed policy on the quality distribution of currency. In general, the model predicts relatively modest responses of currency quality to decreases in DI deposit rates that are anticipated to occur as a consequence of the recircu-


Figure 8 Effect of 40 Percent Deposit Rate Decrease on Quality Distributions of \$5 Notes



lation policy. For \$5 and \$10 notes, our model is able to match the marginal quality distributions perfectly, and the age distributions of unfit notes reasonably well. Thus, we have some confidence in the range of predictions that the different model specifications make for the effects on currency quality of a decrease in deposit rates. In what follows, we discuss potential extensions to the current analysis.

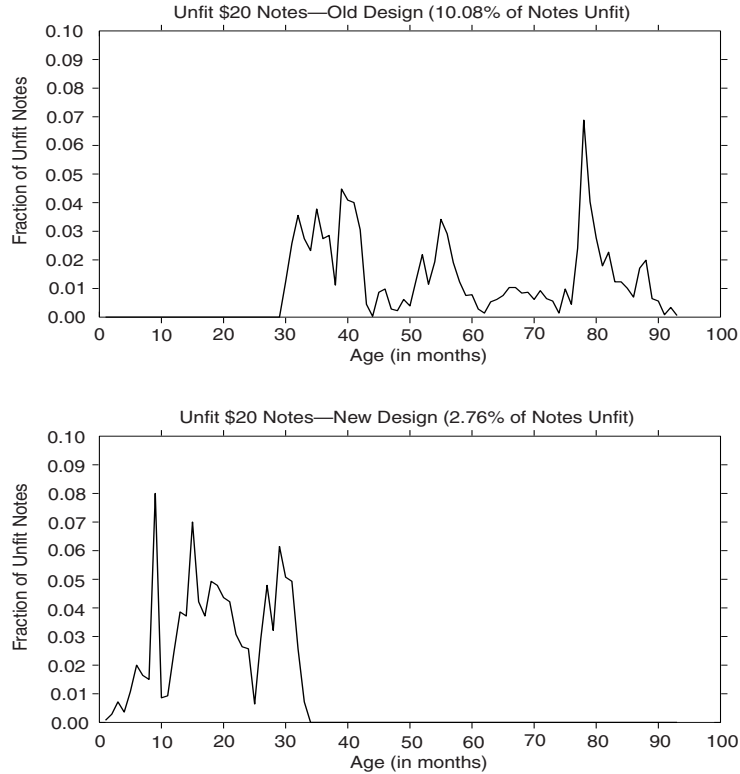
Although our framework allows for sorting by DIs, the quantitative analysis has assumed no sorting occurs. If DIs do sort, then the researcher must take into account that the distribution of currency in circulation is not the same as the distribution of currency that visits the Fed. The derivations in this report do not differentiate between the two distributions, but it is straightforward to do so. If DIs were to sort using the same criteria as the Federal Reserve, then it is likely that the results presented here would overstate the decline in currency quality following implementation of the recirculation policy; by depositing with the Fed only low-quality notes, DIs would offset the deleterious effect of

Figure 9 Effect of 40 Percent Deposit Rate Decrease on Quality Distributions of \$10 Notes



depositing fewer notes. The recirculation policy clearly provides an incentive for at least some DIIs to sort because it imposes fees for cross-shipment of fit currency only.

Our analysis has not addressed \$20 notes. Figure 10 illustrates the difficulty they present: they are not in a steady state but are transiting from the old to the new design. Of the old notes, more than 10 percent are unfit, whereas of the new notes, less than 3 percent are unfit. All the old notes are more than two years old, whereas all the new notes are less than three years old. Our model is not inherently restricted to steady state situations. To apply it to the 20s, one would want to use the form of the model in (7) and also allow for γ (the growth rate of currency) to be time-varying or at least allow γ to vary across designs. The non-steady-state form of the model (7) also could be useful more generally, in providing a check on our estimates. If there is good data on marginal quality distributions available monthly, then that data can be used to generate forecast errors for the model on a real-time basis.

Figure 10 Age Distributions of Unfit \$20 Notes

One reason to question the steady-state assumption is the possibility that the payments system is in the midst of a transition away from the use of currency and toward electronic forms of payment. Although it is difficult to distinguish a change in the trend from a transitory shock, data on the stock of currency does give some credence to this concern: from 2002 to 2007 the growth rate of currency has declined steadily, and at 2 percent for the 12 months ending in June 2007 it is currently growing more slowly than most measures of nominal spending. A decreasing currency growth rate means that there is a decreasing rate of new notes introduced into circulation. This would likely require stronger measures by the Federal Reserve to maintain currency quality in response to a decrease in deposit rates.

The version of the model estimated here is very small and easy to estimate. Expanding the model so that it describes the joint distribution of all three quality dimensions studied here leads to an unmanageably large system. A middle ground that might be worth pursuing would be to specify the model in

terms of two dimensions, say graffiti and soil level, and include information about unfitness in other dimensions, as we have done here.

Finally, it would be useful to embed the currency quality model of this article in an economic model of DIs and households. The DIs' deposit rate and sorting policy (both summarized by ρ) would then be endogenously determined. Such a model could be used to predict the effects of a change in the Federal Reserve's pricing policy on DI behavior. It could also be used to conduct welfare analysis of different pricing and shredding policies. The model in Lacker (1993) is a natural starting point.

APPENDIX: DETAILS OF CALCULATING AGE DISTRIBUTION

It is straightforward to compute the age distribution of notes for any quality level and the quality distribution at any age. Begin by defining the fraction of notes at quality level q and age k to be $h_{q,k}$. These fractions satisfy

$$1 = \sum_{k=0}^{\infty} \sum_{q=1}^Q h_{q,k}. \tag{13}$$

For convenience, define h_k to be the Q -vector containing in element q , the fraction of notes that are k -periods old, and in quality level q :

$$h_k = \begin{bmatrix} h_{1,k} \\ h_{2,k} \\ \vdots \\ h_{Q,k} \end{bmatrix}. \tag{14}$$

We also have that $h_{q,k} = (e_q^Q)' h_k$, where e_q^Q is a $Q \times 1$ selection vector with a 1 in the q^{th} element and zeros elsewhere.

The fraction of brand-new notes is

$$\sum_{q=1}^Q h_{q,0} = \frac{\gamma}{1 + \gamma} + \sum_{j=1}^N (1 - \alpha_j) \rho_j f_j^*, \tag{15}$$

and since the quality distribution of new notes is g , the fractions of notes that are new and in each quality level q are

$$h_0 = \left(\frac{\gamma}{1 + \gamma} + \sum_{j=1}^N (1 - \alpha_j) \rho_j f_j^* \right) \cdot g. \tag{16}$$

For one-period old notes, the fractions are

$$h_1 = \left[\frac{\pi \cdot \text{diag}(1 - \rho) + \text{diag}(\alpha \odot \rho)}{1 + \gamma} \right] \cdot h_0. \quad (17)$$

Likewise, we have

$$h_{k+1} = \left[\frac{\pi \cdot \text{diag}(1 - \rho) + \text{diag}(\alpha \odot \rho)}{1 + \gamma} \right]^{k+1} \cdot h_0, \text{ for } k = 0, 1, \dots, \quad (18)$$

with h_0 determined by (16). Thus, we can calculate the fraction of notes at any age-quality combination as

$$h_{q,k} = (e_q^Q)' \left[\frac{\pi \cdot \text{diag}(1 - \rho) + \text{diag}(\alpha \odot \rho)}{1 + \gamma} \right]^k \cdot h_0. \quad (19)$$

The age distribution of quality- q notes is

$$\frac{1}{\sum_{k=0}^{\infty} h_{q,k}} \begin{bmatrix} h_{q,0} \\ h_{q,1} \\ \vdots \\ h_{q,\infty} \end{bmatrix}. \quad (20)$$

and the quality distribution of age- k notes is

$$\frac{1}{\sum_{q=0}^Q h_{q,k}} \begin{bmatrix} h_{1,k} \\ h_{2,k} \\ \vdots \\ h_{Q,k} \end{bmatrix}. \quad (21)$$

REFERENCES

- Board of Governors of the Federal Reserve System. 2003a. "Federal Reserve Bank Currency Recirculation Policy. Request for Comment, notice." Docket No. OP-1164, October 7. Available at: <http://www.federalreserve.gov/Boarddocs/press/other/2003/20031008/attachment.pdf> (accessed July 13, 2007).
- Board of Governors of the Federal Reserve System. 2003b. Press release, October 8. Available at: <http://www.federalreserve.gov/boarddocs/press/other/2003/20031008/default.htm> (accessed July 13, 2007).
- Board of Governors of the Federal Reserve System. 2006a. "Appendix C. Currency Budget." *Annual Report: Budget Review*: 31.

- Board of Governors of the Federal Reserve System. 2006b. "Federal Reserve Currency Recirculation Policy. Final Policy." Docket No. OP-1164, March 17. Available at: <http://www.federalreserve.gov/newsevents/press/other/other20060317a1.pdf> (accessed July 13, 2007).
- Bureau of Engraving and Printing. 2007. "Money Facts, Fun Facts, Did You Know?" Available at: <http://www.bep.treas.gov/document.cfm/18/106> (accessed October 30, 2007).
- Federal Reserve Bank of New York. 2006. "Currency Processing and Destruction." Available at: <http://www.ny.frb.org/aboutthefed/fedpoint/fed11.html> (accessed May 30, 2007).
- Federal Reserve Bank of San Francisco. 2006. "Cash Counts." *Annual Report*: 6–17.
- Federal Reserve System, Currency Quality Work Group. 2007. "Federal Reserve Bank Currency Quality Monitoring Program." Internal memo, June.
- Ferrari, Shaun. 2005. "Division of Reserve Bank Operations and Payment Systems." Board of Governors of the Federal Reserve System Governors. E-mail message to author, September 9, 2005.
- Friedman, Milton. 1969. *Optimum Quantity of Money: And Other Essays*. Chicago, IL: Aldine Publishing Company.
- Klein, Raymond M., Simon Gadbois, and John J. Christie. 2004. "Perception and Detection of Counterfeit Currency in Canada: Note Quality, Training, and Security Features." In *Optical Security and Counterfeit Deterrence Techniques V*, ed. Rudolf L. van Renesse, SPIE Conference Proceedings, vol. 5310.
- Lacker, Jeffrey M. 1993. "Should We Subsidize the Use of Currency?" Federal Reserve Bank of Richmond *Economic Quarterly* 79 (1): 47–73.
- Lacker, Jeffrey M., and Alexander L. Wolman. 1997. "A Simple Model of Currency Quality." Mimeo, Federal Reserve Bank of Richmond (November).
- Stokely, Nancy L., Robert E. Lucas, Jr., with Edward C. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.

Non-Stationarity and Instability in Small Open-Economy Models Even When They Are “Closed”

Thomas A. Lubik

Open economies are characterized by the ability to trade goods both intra- and intertemporally, that is, their residents can move goods and assets across borders and over time. These transactions are reflected in the current account, which measures the value of a country's export and imports, and its mirror image, the capital account, which captures the accompanying exchange of assets. The current account serves as a shock absorber, which agents use to optimally smooth their consumption. The means for doing so are borrowing and lending in international financial markets. It almost goes without saying that international macroeconomists have had a long-standing interest in analyzing the behavior of the current account.

The standard intertemporal model of the current account conceives a small open economy as populated by a representative agent who is subject to fluctuations in his income. By having access to international financial markets, the agent can lend surplus funds or make up shortfalls for what is necessary to maintain a stable consumption path in the face of uncertainty. The international macroeconomics literature distinguishes between an international asset market that is incomplete and one that is complete. The latter describes a modeling framework in which agents have access to a complete set of state-contingent securities (and, therefore, can share risk perfectly); when markets

■ I am grateful to Andreas Hornstein, Alex Wolman, Juan Carlos Hatchondo, and Nashat Moin for comments that improved the article. I also wish to thank Jinill Kim and Martín Uribe for useful discussions and comments which stimulated this research. The views expressed in this article are those of the author, and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: Thomas.Lubik@rich.frb.org.

are incomplete, on the other hand, agents can only trade in a restricted set of assets, for instance, a bond that pays fixed interest.

The small open-economy model with incomplete international asset markets is the main workhorse in international macroeconomics. However, the baseline model has various implications that may put into question its usefulness in studying international macroeconomic issues. When agents decide on their intertemporal consumption path they trade off the utility-weighted return on future consumption, measured by the riskless rate of interest, against the return on present consumption, captured by the time discount factor. The basic set-up implies that expected consumption growth is stable only if the two returns exactly offset each other, that is, if the product of the discount factor and the interest rate equal one. The entire optimization problem is ill-defined for arbitrary interest rates and discount factors as consumption would either permanently decrease or increase.¹

Given this restriction on two principally exogenous parameters, the model then implies that consumption exhibits random-walk behavior since the effects of shocks to income are buffered by the current account to keep consumption smooth. The random-walk in consumption, which is reminiscent of Hall's (1978) permanent income model with linear-quadratic preferences, is problematic because it implies that all other endogenous variables inherit this non-stationarity so that the economy drifts over time arbitrarily far away from its initial condition. To summarize, the standard small open-economy model with incomplete international asset markets suffers from what may be labelled the unit-root problem. This raises several issues, not the least of which is the overall validity of the solution in the first place, and its usefulness in conducting business cycle analysis.

In order to avoid this unit-root problem, several solutions have been suggested in the literature. Schmitt-Grohé and Uribe (2003) present an overview of various approaches. In this article, I am mainly interested in inducing stationarity by assuming a debt-elastic interest rate. Since this alters the effective interest rate that the economy pays on foreign borrowing, the unit root in the standard linearized system is reduced incrementally below unity. This preserves a high degree of persistence, but avoids the strict unit-root problem. Moreover, a debt-elastic interest rate has an intuitive interpretation as an endogenous risk premium. It implies, however, an additional, essentially ad hoc feedback mechanism between two endogenous variables. Similar to the literature on the determinacy properties of monetary policy rules or models with

¹ Conceptually, the standard current account model has a lot of similarities to a model of intertemporal consumer choice with a single riskless asset. The literature on the latter gets around some of the problems detailed here by, for instance, imposing borrowing constraints. Much of that literature is, however, mired in computational complexities as standard linearization-based solution techniques are no longer applicable.

increasing returns to scale, the equilibrium could be indeterminate or even non-existent.

I show in this article that commonly used specifications of the risk premium do not lead to equilibrium determinacy problems. In all specifications, indeterminacy of the rational expectations equilibrium can be ruled out, although in some cases there can be multiple steady states. It is only under a specific assumption on whether agents internalize the dependence of the interest rate on the net foreign asset position that no equilibrium may exist.

I proceed by deriving, in the next section, an analytical solution for the (linearized) canonical small open-economy model which tries to illuminate the extent of the unit-root problem. Section 2 then studies the determinacy properties of the model when a stationarity-inducing risk-premium is introduced. In Section 3, I investigate the robustness of the results by considering different specifications that have been suggested in the literature. Section 4 presents an alternative solution to the unit-root problem via portfolio adjustment costs, while Section 5 summarizes and concludes.

1. THE CANONICAL SMALL OPEN-ECONOMY MODEL

Consider a small open economy that is populated by a representative agent² whose preferences are described by the following utility function:

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \quad (1)$$

where $0 < \beta < 1$ and E_t is the expectations operator conditional on the information set at time t . The period utility function u obeys the usual Inada conditions which guarantee strictly positive consumption sequences $\{c_t\}_{t=0}^{\infty}$. The economy's budget constraint is

$$c_t + b_t \leq y_t + R_{t-1}b_{t-1}, \quad (2)$$

where y_t is stochastic endowment income; R_t is the gross interest rate at which the agent can borrow and lend b_t on the international asset market. The initial condition is $b_{-1} \geq 0$. In the canonical model, the interest rate is taken parametrically.

The agent chooses consumption and net foreign asset sequences $\{c_t, b_t\}_{t=0}^{\infty}$ to maximize (1) subject to (2). The usual transversality condition applies. First-order necessary conditions are given by

$$u'(c_t) = \beta R_t E_t u'(c_{t+1}), \quad (3)$$

² In what follows, I use the terms "agent," "economy," and "country," interchangeably. This is common practice in the international macro literature and reflects the similarity between small open-economy models and partial equilibrium models of consumer choice.

and the budget constraint (2) at equality. The Euler equation is standard. At the margin, the agent is willing to give up one unit of consumption, valued by its marginal utility, if he is compensated by an additional unit of consumption next period augmented by a certain (properly discounted) interest rate, and evaluated by its uncertain contribution to utility. Access to the international asset market thus allows the economy to smooth consumption in the face of uncertain domestic income. Since the economy can only trade in a single asset such a scenario is often referred to as one of “incomplete markets.” This stands in contrast to a model where agents can trade a complete set of state-contingent assets (“complete markets”).

In what follows, I assume for ease of exposition that y_t is *i.i.d.* with mean \bar{y} , and that the interest rate is constant and equal to the world interest rate $R^* > 1$. The latter assumption will be modified in the next section. Given these assumptions a steady state only exists if $\beta R^* = 1$. Steady-state consumption is, therefore, $\bar{c} = \bar{y} + \frac{1-\beta}{\beta}\bar{b}$. Since consumption is strictly positive, this imposes a restriction on the admissible level of net foreign assets $\bar{b} > -\frac{\beta}{1-\beta}\bar{y}$. The structure of this model is such that it imposes a restriction on the two principally structural parameters β and R^* , which is theoretically and empirically problematic; there is no guarantee or mechanism in the model that enforces this steady-state restriction to hold. Even more so, the steady-state level of a choice variable, namely net foreign assets \bar{b} , is not pinned down by the model’s optimality conditions. Instead, there exists a multiplicity of steady states indexed by the initial condition $b = b_{-1}$.³

Despite these issues, I now proceed by linearizing the first-order conditions around the steady state for some \bar{b} . Denoting $\tilde{x}_t = \log x_t - \log \bar{x}$ and $\hat{x}_t = x_t - \bar{x}$, the linearized system is⁴

$$E_t \tilde{c}_{t+1} = \tilde{c}_t, \quad (4)$$

$$\tilde{c}_t + \hat{b}_t = \bar{y}\tilde{y}_t + \beta^{-1}\hat{b}_{t-1}. \quad (5)$$

It can be easily verified that the eigenvalues of this dynamic system in $[\tilde{c}_t, \hat{b}_t]$ are $\lambda_1 = 1$, $\lambda_2 = \beta^{-1} > 1$. Since \hat{b} is a pre-determined variable this results in a unique rational expectations equilibrium for all admissible parameter values. The dynamics of the solution are given by (a detailed derivation of the solution

³ In the international real business cycle literature, for instance, Baxter and Crucini (1995), \bar{b} is, therefore, often treated as a parameter to be calibrated.

⁴ Since the interest rate is constant, the curvature of the utility function does not affect the time path of consumption and, consequently, does not appear in the linearization. Moreover, net foreign assets are approximated in levels since b_t can take on negative values or zero, for which the logarithm is not defined.

can be found in the Appendix)

$$\tilde{c}_t = \frac{1 - \beta}{\beta} \frac{\widehat{b}_{t-1}}{\bar{c}} + (1 - \beta) \frac{\bar{y}}{\bar{c}} \tilde{y}_t, \quad (6)$$

$$\frac{\widehat{b}_t}{\bar{c}} = \frac{\widehat{b}_{t-1}}{\bar{c}} + \beta \frac{\bar{y}}{\bar{c}} \tilde{y}_t. \quad (7)$$

The contemporaneous effect of a 1 percent innovation to output is to raise foreign lending as a fraction of steady-state consumption by $\beta \frac{\bar{y}}{\bar{c}}$ percent, which is slightly less than unity in the baseline case $\bar{b} = 0$. In line with the permanent income hypothesis only a small percentage of the increase in income is consumed presently, so that future consumption can be raised permanently by $\frac{1-\beta}{\beta}$. The non-stationarity of this solution, the “unit-root problem,” is evident from the unit coefficient on lagged net foreign assets in (7). Temporary innovations have, therefore, permanent effects; the endogenous variables wander arbitrarily far from their starting values. This also means that the unconditional second moments, which are often used in business cycle analysis to evaluate a model, do not exist.

Moreover, the solution is based on an approximation that is technically only valid in a small neighborhood around the steady state. This condition will be violated eventually with probability one, thus ruling out the validity of the linearization approach in the first place. Since an equation system such as (4)–(5) is at the core of much richer open-economy models, the non-stationarity of the incomplete markets solution carries over. The unit-root problem thus raises the question whether (linearized) incomplete market models offer accurate descriptions of open economies. In the next sections, I study the equilibrium properties of various modifications to the canonical model that have been used in the literature to “fix” the unit-root problem.⁵

2. INDUCING STATIONARITY VIA A DEBT-ELASTIC INTEREST RATE

The unit-root problem arises because of the random-walk property of consumption in the linearized Euler equation (4). Following Schmitt-Grohé and Uribe (2003) and Senhadji (2003), a convenient solution is to make the interest rate the economy faces a function of net foreign assets $R_t = F(b_t - \bar{b})$,

⁵ In most of the early international macro literature, the unit-root problem tended to be ignored despite, in principle valid, technical problems. The unit root is transferred to the variables of interest, such as consumption, on the order of the net interest rate, which is quantitatively very small (in the present example, $\frac{1-\beta}{\beta}$). While second moments do not exist in such a non-stationary environment, researchers can still compute sample moments to perform business cycle analysis. Moreover, Schmitt-Grohé and Uribe (2003) demonstrate that the dynamics of the standard model with and without the random walk in endogenous variables are quantitatively indistinguishable over a typical time horizon. Their article, thus, gives support for the notion of using the incomplete market setup for analytical convenience.

where F is decreasing in b , \bar{b} is the steady-state value of b , and $F(0) = R^*$. If a country is a net foreign borrower, it pays an interest rate that is higher than the world interest rate. The reference point for the assessment of the risk premium is the country's steady state. Intuitively, \bar{b} represents the level of net foreign assets that is sustainable in the long run, either by increasing (if positive) or decreasing (if negative) steady-state consumption relative to the endowment.

If a country deviates in its borrowing temporarily from what international financial markets perceive as sustainable in the long run, it is penalized by having to pay a higher interest rate than "safer" borrowers. This has the intuitively appealing implication that the difference between the world interest rate and the domestically relevant rate can be interpreted as a risk premium. The presence of a debt-elastic interest rate can be supported by empirical evidence on the behavior of spreads, that is, the difference between a country's interest rate and a benchmark rate, paid on sovereign bonds in emerging markets (Neumeayer and Perri, 2005). Relative to interest rates on U.S. Treasuries, the distribution of spreads has a positive mean, and they are much more volatile.

A potential added benefit of using a debt-elastic interest rate is that proper specification of F may allow one to derive the steady-state value of net foreign assets endogenously. However, the introduction of a new, somewhat arbitrary link between endogenous variables raises the possibility of equilibrium indeterminacy and non-existence similar to what is found in the literature on monetary policy rules and production externalities. I study two cases. In the first case, the small open economy takes the endogenous interest rate as given. That is, the dependence of the interest rate on the level of outstanding net assets is not internalized. The second case assumes that agents take the feedback from assets to interest rates into account.

No Internalization

The optimization problem for the small open economy is identical to the canonical case discussed above. The agent does not take into account that the interest charged for international borrowing depends on the amount borrowed. Analytically, the agent takes R_t as given. The first-order conditions are consequently (2) and (3). Imposing the interest rate function $R_t = F(b_t - \bar{b})$ yields the Euler equation when the risk premium is not internalized:

$$u'(c_t) = \beta F(b_t - \bar{b}) E_t u'(c_{t+1}). \quad (8)$$

The Euler equation highlights the difference to the canonical model. Expected consumption growth now depends on an endogenous variable, which tilts the consumption path away from random-walk behavior. However, existence of a steady state still requires $R = R^* = \beta^{-1}$. Despite the assumption of an endogenous risk premium, this model suffers from the same deficiency as the

canonical model in that the first-order conditions do not fully pin down all endogenous variables in steady state.⁶

After substituting the interest rate function, the first-order conditions are linearized around some steady state \bar{b} . I impose additional structure by assuming that the period utility function $u(c) = \frac{c^{1-1/\sigma}-1}{1-1/\sigma}$, where $\frac{u''(c)c}{u'(c)} = -1/\sigma$, and $\sigma > 0$ is the intertemporal substitution elasticity. Since I am mainly interested in the determinacy properties of the model, I also abstract from time variation in the endowment process $y_t = y, \forall t$. Furthermore, I assume that $F'(0) = -\psi$.⁷ The linearized equation system is then

$$\begin{aligned} E_t \tilde{c}_{t+1} &= \tilde{c}_t - \beta\sigma\psi\hat{b}_t, \\ \bar{c}\tilde{c}_t + \hat{b}_t &= \left(\frac{1}{\beta} - \psi\bar{b}\right)\hat{b}_{t-1}. \end{aligned} \tag{9}$$

The reduced-form coefficient matrix of this system can be obtained after a few steps:

$$\begin{bmatrix} 1 & -\beta\sigma\psi \\ -\bar{c} & \frac{1}{\beta} + (\beta\sigma\bar{c} - \bar{b})\psi \end{bmatrix}, \tag{10}$$

where $\bar{c} = y + \frac{1-\beta}{\beta}\bar{b}$ as before. I can now establish

Proposition 1 *In the model with additively separable risk premium and no internalization, there is a unique equilibrium for all admissible parameter values.*

Proof. In order to investigate the determinacy properties of this model, I first compute the trace $tr = 1 + \frac{1}{\beta} + (\beta\sigma\bar{c} - \bar{b})\psi$ and the determinant $det = \frac{1}{\beta} - \psi\bar{b}$. Since there is one predetermined variable, a unique equilibrium requires one root inside and one root outside the unit circle. Both (zero) roots inside the unit circle imply indeterminacy (non-existence). The Appendix shows that determinacy requires $|tr| > 1 + det$, while $|det| \leq 1$. The first condition reduces to $\beta\sigma\psi\bar{c} > 0$, which is always true because of strictly positive consumption. Note also that $tr > 1 + det$. Indeterminacy and non-existence require $|tr| < 1 + det$, which cannot hold because of positive consumption. The proposition then follows immediately.

⁶This is an artifact of the assumption of no internalization and the specific assumptions on the interest rate function.

⁷An example of a specific functional form that is consistent with these assumptions and that has been used in the literature (e.g., Schmitt-Grohé and Uribe 2003) is

$$R_t = R^* + \psi \left[e^{-(b_t - \bar{b})} - 1 \right].$$

Internalization

An alternative scenario assumes that the agent explicitly takes into account that the interest rate he pays on foreign borrowing depends on the amount borrowed. Higher borrowing entails higher future debt service which reduces the desire to borrow. The agent internalizes the cost associated with becoming active on the international asset markets in that he discounts future interest outlays not at the world interest rate but at the domestic interest rate, which is inclusive of the risk premium.⁸

The previous assumptions regarding the interest rate function and the exogenous shock remain unchanged. Since the economy internalizes the dependence of the interest rate on net foreign assets, the first-order conditions change. Analytically, I substitute the interest rate function into the budget constraint (2) before taking derivatives, thereby eliminating R from the optimization problem. The modified Euler equation is

$$u'(c_t) = \beta F(b_t - \bar{b}) [1 + \varepsilon_F(b_t)] E_t u'(c_{t+1}), \quad (11)$$

where $\varepsilon_F(b_t) = \frac{F'(b_t - \bar{b})b_t}{F(b_t - \bar{b})}$ is the elasticity of the interest rate function with respect to net foreign assets. Compared to the case of no internalization, the effective interest rate now includes an additional term in the level of net foreign assets. Whether the steady-state level of \bar{b} is determined, therefore, depends on this elasticity. Maintaining the assumption $F'(0) = -\psi$, it follows that $\varepsilon_F(\bar{b}) = -\psi R^* \bar{b}$.

This provides the additional restriction needed to pin down the steady state:

$$\bar{b} = \frac{R^* - 1/\beta}{\psi}. \quad (12)$$

If the country's discount factor is bigger than $1/R^*$, that is, if it is more patient than those in the rest of the world, its steady-state asset position is strictly positive. A more impatient country, however, accumulates foreign debt to finance consumption. Note further that $R = R^*$, but not necessarily equals β^{-1} , while \bar{b} asymptotically reaches zero as ψ grows large. It is worth emphasizing that $\beta R^* = 1$ is no longer a necessary condition for the existence of a steady state, and that \bar{b} is, in fact, uniquely determined. Internalization of the risk premium, therefore, avoids one of the pitfalls of the standard model, but it also nicely captures the idea that some countries appear to have persistent levels of foreign indebtedness.

⁸The difference between internalization and no internalization of the endogenous risk premium is also stressed by Nason and Rogers (2006). Strictly speaking, with internalization the country stops being a price-taker in international asset markets. This is analogous to open-economy models of "semi-small" countries that are monopolistically competitive and price-setting producers of export goods. Schmitt-Grohé (1997) has shown that feedback mechanisms of this kind are important sources of non-determinacy of equilibria.

I now proceed by linearizing the equation system:

$$\begin{aligned} E_t \tilde{c}_{t+1} &= \tilde{c}_t - \beta \sigma \psi (2 - \bar{b}) \hat{b}_t, \\ \tilde{c}_t + \hat{b}_t &= (R^* - \psi \bar{b}) \hat{b}_{t-1}. \end{aligned} \quad (13)$$

The coefficient matrix that determines the dynamics can be derived as:

$$\begin{bmatrix} 1 & -\beta \sigma \psi (2 - \bar{b}) \\ -\bar{c} & \frac{1}{\beta} + \beta \sigma \psi (2 - \bar{b}) \bar{c} \end{bmatrix}, \quad (14)$$

where now $\bar{b} = \frac{R^* - 1/\beta}{\psi}$ and $\bar{c} = \bar{y} + (R^* - 1)\bar{b}$. The determinacy properties of this case are given in

Proposition 2 *In the model with additively separable risk premium and internalization, the equilibrium is unique if and only if*

$$\bar{b} < 2,$$

or

$$\bar{b} > 2 + 2 \frac{1 + \beta}{\beta} \frac{1}{\beta \sigma \psi \bar{c}}.$$

No equilibrium exists otherwise.

Proof. The determinant of the system matrix is $\det = \beta^{-1} > 1$. This implies that there is at least one explosive root, which rules out indeterminacy. Since the system contains one jump and one predetermined variable, a unique equilibrium requires $|tr| > 1 + \det$, where $tr = 1 + \beta^{-1} + \beta \sigma \psi \bar{c} (2 - \bar{b})$. The lower bound of the condition establishes that $\beta \sigma \psi (2 - \bar{b}) \bar{c} > 0$. Since $\bar{c} > 0$, it must be that $\bar{b} < 2$. From $-tr > 1 + \det$, the second part of the determinacy region follows after simply rearranging terms. The proposition then follows immediately.

The proposition shows that a sufficient condition for determinacy is that the country is a net foreign borrower, which implies $\beta^{-1} > R^*$. A relatively impatient country borrows from abroad to sustain current consumption. Since this incurs a premium above the world interest rate, the growth rate of debt is below that of, say, the canonical case, and debt accumulation is, therefore, nonexplosive. Even if the country is a net foreign lender, determinacy can still be obtained for $0 < \bar{b} < 2$ or $R^* < \beta^{-1} + 2\psi$. A slightly more patient country than the rest of the world would imply a determinate equilibrium if the (internalized) interest rate premium is large enough.

From a technical point of view, non-existence arises if both roots in (13) are larger than unity, so that both difference equations are unstable. The budget constraint then implies an explosive time path for assets b which would violate transversality. This is driven by explosive consumption growth financed by interest receipts on foreign asset holdings. In the non-existence region, these are large so as not to be balanced by the decline in the interest rate. Effectively,

the economy both over-consumes and over-accumulates assets, which cannot be an equilibrium. The only possible equilibrium is, therefore, at the (unique) steady state, while dynamics around it are explosive. This highlights the importance of the elasticity term $1 + \varepsilon_F(b_t)$ in equation (11), which has the power to tilt the consumption away from unit-root (and explosive) behavior for the right parameterization.

As the proposition shows, the non-existence region has an upper bound beyond which the equilibrium is determinate again. The following numerical example using baseline parameter values⁹ demonstrates, however, that this boundary is far above empirically reasonable values. Figure 1, Panels A and B depict the determinacy regions for net foreign assets for varying values of σ and ψ . Note that below the lower bound $\bar{b} = 2$ the equilibrium is always determinate, while the size of the non-existence region is decreasing in the two parameters. Recall from equation (12) that the steady-state level \bar{b} depends on the spread between the world interest rate and the inverse of the discount factor. Non-existence, therefore, arises if $\psi < \frac{1}{2}(R^* - \beta^{-1})$. In other words, if there is a large wedge between R^* and β^{-1} , a researcher has to be careful not to choose an elasticity parameter ψ that is too small.

Normalizing output $\bar{y} = 1$, the boundary lies at an asset level that is twice as large as the country's GDP. While this is not implausible, net foreign asset holdings of that size are rarely observed. However, choosing a different normalization, for instance, $\bar{y} = 10$ presents a different picture, in which a plausible calibration for, say, a primary resource exporter, renders the solution of the model non-existent. On the other hand, as \bar{y} becomes large, the upper bound for the non-existence region in Figure 1, Panels A and B moves inward, thereby reducing its size. The conclusion for researchers interested in studying models of this type is to calibrate carefully. Target levels for the net-foreign asset to GDP ratio cannot be chosen independently of the stationarity-inducing parameter ψ if equilibrium existence problems are to be avoided. It is worth pointing out again that indeterminacy, and thus the possibility of sunspot equilibria, can be ruled out in this model.

While it is convenient to represent the boundaries of the determinacy region for net foreign assets b , it is nevertheless an endogenous variable, as is c . The parameter restriction in the above proposition can be rearranged in terms of R^* . That is, the economy has a unique equilibrium if either $R^* < \beta^{-1} + 2\psi$ or $R^* > \beta^{-1} + 2\psi \left[1 + \frac{1+\beta}{\beta} \frac{1}{\beta\sigma\{\psi y + (R^*-1)(R^*-\beta^{-1})\}} \right]$. Again, the equilibrium is non-existent otherwise. Since the second term in brackets is strictly positive, the region of non-existence is nonempty. Although the upper bound is still a function of R^* (and has to be computed numerically), this version presents more intuition.

⁹ Parameter values used are $\beta = 0.98$, $\sigma = 1$, $\psi = 0.001$, and $\bar{y} = 1$.

Figure 1 Determinacy Regions for Net Foreign Assets b and Interest Rates R^*

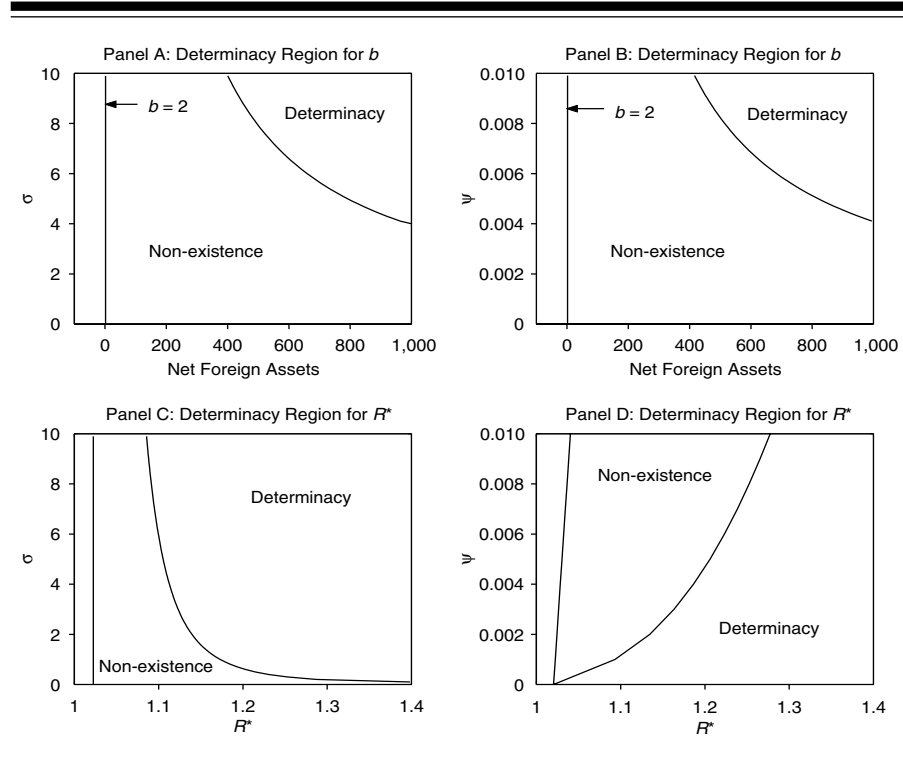


Figure 1, Panels C and D depict the determinacy regions for R^* with varying σ and ψ , respectively. The lower bound of the non-existence region is independent of σ , but increasing in ψ . For a small substitution elasticity, the equilibrium is non-existent unless the economy is more impatient than the rest of the world, inclusive of a factor reflecting the risk premium. This is both consistent with a negative steady-state asset position as well as a small, positive one as long as $b < 2$. Figure 1, Panel D shows that no equilibrium exists even for very small values of ψ . If the economy is a substantial net saver, then the equilibrium is determinate if the world interest rate is (implausibly) high. Analytically, this implies that the asset accumulation equation remains explosive even though there is a large premium to be paid.

To summarize, introducing a debt-elastic interest rate addresses two issues arising in incomplete market models of open economies, viz., the indeterminacy of the steady-state allocation and the induced non-stationarity of the linearized solution. If the derivative of the interest rate function with respect to net asset holdings is nonzero, then the linearized solution is stationary. In the special case when the economy internalizes the dependence of the interest

rate on net foreign assets, the rational expectations equilibrium can be non-existent. However, this situation only arises for arguably extreme parameter values. A nonzero elasticity of the interest rate function is also necessary for the determinacy of the steady state. It is not sufficient, however, as the special case without internalization demonstrated.

3. ALTERNATIVE SPECIFICATIONS

The exposition above used the general functional form $R_t = F(b_t - \bar{b})$, with $F(0) = R^*$ and $F'(0) = -\psi$. A parametric example for this function would be additive in the risk premium term, i.e., $R_t = R^* + \psi [e^{-(b_t - \bar{b})} - 1]$. Alternatively, the risk premium could also be chosen multiplicatively, $R_t = R^* \psi(b_t)$, with $\psi(b) = 1$, $\psi' < 0$. With internalization, the Euler equation can then be written as:

$$u'(c_t) = \beta R^* \psi(b_t) [1 + \varepsilon_F(b_t)] E_t u'(c_{t+1}). \quad (15)$$

$\varepsilon_F(b_t)$ is the elasticity of the risk premium function with respect to foreign assets. Again, the first-order condition shows how a debt-elastic interest rate tilts consumption away from pure random-walk behavior.

A specific example for the multiplicative form of the interest rate function is $R_t = R^* e^{-\psi(b_t - \bar{b})}$, which in log-linear form conveniently reduces to $\tilde{R}_t = -\psi \hat{b}_t$. Assuming no internalization, the steady state is again not pinned down so that $R = R^* = \beta^{-1}$, and the above restrictions on \bar{b} apply. Internalization of the risk premium leads to $\bar{b} = \frac{R^* - 1/\beta}{\psi R^*}$. Again, the economy is a net saver when it is more patient than the rest of the world. As opposed to the case of an additive premium, the equilibrium is determinate for the entire parameter space. This can easily be established in

Proposition 3 *In the model with multiplicative risk premium, with either internalization or no internalization, the equilibrium is unique for all parameter values.*

Proof. See Appendix.

Nason and Rogers (2006) suggest a specification for the risk premium that is additive in net foreign assets relative to aggregate income: $R_t = R^* - \psi \frac{b_t}{y_t}$.¹⁰ The difference to the additive premium considered above is that even without internalization, foreign and domestic rates need not be the same in the steady state. In the latter case, $\bar{b} = \frac{R^* - 1/\beta}{\psi}$, whereas with internalization, $\bar{b} = \frac{1}{2} \frac{R^* - 1/\beta}{\psi}$. This shows that the endogenous risk premium reduces asset

¹⁰ Note that in this case the general form specification of the interest rate function is $R_t = F(b_t)$, and not $R_t = F(b_t - \bar{b})$.

accumulation when agents take into account the feedback effect on the interest rate. The determinacy properties of this specification are established in

Proposition 4 *If the domestic interest rate is given by $R_t = R^* - \psi \frac{b_t}{y_t}$, under either internalization or no internalization, the equilibrium is unique for all parameter values.*

Proof. See Appendix.

It may appear that the determinacy properties are pure artifacts of the linearization procedure. While I *log*-linearized consumption, functions of b_t were approximated in levels as net foreign assets may very well be negative or zero.¹¹ Dotsey and Mao (1992), for instance, have shown that the accuracy of linear approximation procedures depends on the type of linearization chosen. It can be verified,¹² however, that this is not a problem in this simple model as far as the determinacy properties are concerned. The coefficient matrix for all model specifications considered is invariant to the linearization.

4. PORTFOLIO ADJUSTMENT COSTS

Finally, I consider one approach to the unit-root problem that does not rely on feedback from net foreign assets to the interest rate. Several authors, for example, Schmitt-Grohé and Uribe (2003) and Neumeyer and Perri (2005), have introduced quadratic portfolio adjustment costs to guarantee stationarity. It is assumed that agents have to pay a fee in terms of lost output if their transactions on the international asset market lead to deviations from some long-run (steady-state) level \bar{b} . The budget constraint is thus modified as follows:

$$c_t + b_t + \frac{\psi}{2} (b_t - \bar{b})^2 = y_t + R^* b_{t-1}, \quad (16)$$

where $\psi > 0$, and the interest rate on foreign assets is equal to the constant world interest rate R^* . The Euler equation is

$$u'(c_t) [1 + \psi (b_t - \bar{b})] = \beta R^* E_t u'(c_{t+1}). \quad (17)$$

If the economy wants to purchase an additional unit of foreign assets, current consumption declines by one plus the transaction cost $\psi (b_t - \bar{b})$. The payoff for the next period is higher consumption by one unit plus the fixed (net) world interest rate.

Introducing this type of portfolio adjustment costs does not pin down the steady-state value of \bar{b} . The Euler equation implies the same steady-state restriction as the canonical model, namely $\beta R^* = 1$ and $\bar{b} > -\frac{\beta}{1-\beta} \bar{y}$.

¹¹ The interpretation of the linearized system in terms of percentage deviations from the steady state can still be preserved by expressing foreign assets relative to aggregate income or consumption, as in equation (7).

¹² Details are available from the author upon request.

However, the Euler equation (17) demonstrates the near equivalence between the debt-dependent interest rate function and the debt-dependent-borrowing cost formulation. The key to avoiding a unit root in the dynamic model is to generate feedback that tilts expected consumption growth, which can be achieved in various ways.

The coefficient matrix of the two-variable system in $[\tilde{c}_t, \hat{b}_t]$ is given by

$$\begin{bmatrix} 1 & -\sigma\psi \\ -\bar{c} & \beta^{-1} + \sigma\psi\bar{c} \end{bmatrix}.$$

It can be easily verified that both eigenvalues are real and lie on opposite sides of the unit circle over the entire admissible parameter space. The rational expectations solution is, therefore, unique. The same conclusion applies when different linearization schemes, as previously discussed, are used.

It is worthwhile to point out that Schmitt-Grohé and Uribe (2003) have suggested that the model with portfolio adjustment costs and the model with a debt-elastic interest rate imply similar dynamics. Inspection of the two respective Euler equations reveals that the debt-dependent discount factors in the linearized versions are identical for a properly chosen parameterization. However, portfolio costs do not appear in the linearized budget constraint, since they are of second order, whereas the time-varying interest rate changes debt dynamics in a potentially critical way. It follows, that this assertion is true only for that part of the parameter space that results in a unique solution, but a general equivalence result, such as between internalized and external risk premia, cannot be derived.

5. CONCLUSION

Incomplete market models of small open economies imply non-stationary equilibrium dynamics. Researchers who want to work with this type of model are faced with a choice between theoretical rigor and analytical expediency in terms of a model solution. In order to alleviate this tension, several techniques to induce stationarity have been suggested in the literature. This article has investigated the determinacy properties of models with debt-elastic interest rates and portfolio adjustment costs. The message is a mildly cautionary one. Although analytically convenient, endogenizing the interest rate allows for the possibility that the rational expectations equilibrium does not exist. I show that an additively separable risk premium with a specific functional form that is used in the literature can imply non-existence for a plausible parameterization. I suggest alternative specifications that are not subject to this problem. In general, however, this article shows that the determinacy properties depend on specific functional forms, which is not readily apparent a priori.

A question that remains is to what extent the findings in this article are relevant in richer models. Since analytical results may not be easily available, this remains an issue for further research. Moreover, there are other suggested solutions to the unit-root problem. As the article has emphasized, the key is to tilt expected consumption growth away from unity. I have only analyzed approaches that work on endogenizing the interest rate, but just as conceivably the discount factor β could depend on other endogenous variables as in the case of Epstein-Zin preferences. The rate at which agents discount future consumption streams might depend on their utility level, which in turn depends on consumption and net foreign assets. Again, this would provide a feedback mechanism from assets to the consumption tilt factor. Little is known about equilibrium determinacy properties under this approach.

APPENDIX

Solving the Canonical Model

The linearized equation system describing the dynamics of the model is

$$\begin{aligned} E_t \tilde{c}_{t+1} &= \tilde{c}_t, \\ \bar{c} \tilde{c}_t + \hat{b}_t &= \bar{y} \tilde{y}_t + \beta^{-1} \hat{b}_{t-1}. \end{aligned}$$

I solve the model by applying the method described in Sims (2002). In order to map the system into Sims's framework, I define the endogenous forecast error η_t as follows:

$$\tilde{c}_t = \xi_{t-1}^c + \eta_t = E_{t-1} \tilde{c}_t + \eta_t.$$

The system can then be rewritten as:

$$\begin{bmatrix} 1 & 0 \\ \bar{c} & 1 \end{bmatrix} \begin{bmatrix} \xi_t^c \\ \hat{b}_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \beta^{-1} \end{bmatrix} \begin{bmatrix} \xi_{t-1}^c \\ \hat{b}_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ \bar{y} \end{bmatrix} \tilde{y}_t + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \eta_t.$$

Invert the lead matrix $\begin{bmatrix} 1 & 0 \\ \bar{c} & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 \\ -\bar{c} & 1 \end{bmatrix}$, and multiply through:

$$\begin{bmatrix} \xi_t^c \\ \hat{b}_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -\bar{c} & \beta^{-1} \end{bmatrix} \begin{bmatrix} \xi_{t-1}^c \\ \hat{b}_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ \bar{y} \end{bmatrix} \tilde{y}_t + \begin{bmatrix} 1 \\ -\bar{c} \end{bmatrix} \eta_t.$$

Since the autoregressive coefficient matrix is triangular, the eigenvalues of the system can be read off the diagonal: $\lambda_1 = 1$, $\lambda_2 = \beta^{-1} > 1$. This matrix can be diagonalized as follows:

$$\begin{bmatrix} \frac{1-\beta}{\bar{c}\beta} & 0 \\ 1 & \beta^{-1} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \beta^{-1} \end{bmatrix} \begin{bmatrix} \frac{\bar{c}\beta}{1-\beta} & 0 \\ -\frac{\bar{c}\beta^2}{1-\beta} & \beta \end{bmatrix}.$$

Multiply the system by the matrix of right eigenvectors to get:

$$\begin{bmatrix} \frac{\bar{c}\beta}{1-\beta} & 0 \\ -\frac{\bar{c}\beta^2}{1-\beta} & \beta \end{bmatrix} \begin{bmatrix} \xi_t^c \\ \hat{b}_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \beta^{-1} \end{bmatrix} \begin{bmatrix} \frac{\bar{c}\beta}{1-\beta} & 0 \\ -\frac{\bar{c}\beta^2}{1-\beta} & \beta \end{bmatrix} \begin{bmatrix} \xi_{t-1}^c \\ \hat{b}_{t-1} \end{bmatrix} \\ + \begin{bmatrix} 0 \\ \beta\bar{y} \end{bmatrix} \tilde{y}_t + \begin{bmatrix} \frac{\bar{c}\beta}{1-\beta} \\ -\frac{\bar{c}\beta}{1-\beta} \end{bmatrix} \eta_t.$$

Define $w_{1t} = \frac{\bar{c}\beta}{1-\beta}\xi_t^c$ and $w_{2t} = -\frac{\bar{c}\beta^2}{1-\beta}\xi_t^c + \beta\hat{b}_t$, then:

$$\begin{bmatrix} w_{1t} \\ w_{2t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \beta^{-1} \end{bmatrix} \begin{bmatrix} w_{1t-1} \\ w_{2t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ \beta\bar{y} \end{bmatrix} \tilde{y}_t + \begin{bmatrix} \frac{\bar{c}\beta}{1-\beta} \\ -\frac{\bar{c}\beta}{1-\beta} \end{bmatrix} \eta_t.$$

Treat $\lambda_1 = 1$ as a stable eigenvalue. Then the conditions for stability are

$$w_{2t} = 0, \forall t, \\ \beta\bar{y}\tilde{y}_t - \frac{\bar{c}\beta}{1-\beta}\eta_t = 0.$$

This implies a solution for the endogenous forecast error:

$$\eta_t = (1-\beta)\frac{\bar{y}}{c}\tilde{y}_t.$$

The decoupled system can consequently be rewritten as:

$$\begin{bmatrix} w_{1t} \\ w_{2t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} w_{1t-1} \\ w_{2t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ \beta\bar{y} \end{bmatrix} \tilde{y}_t + \begin{bmatrix} \beta\bar{y} \\ -\beta\bar{y} \end{bmatrix} \tilde{y}_t \\ = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} w_{1t-1} \\ w_{2t-1} \end{bmatrix} + \begin{bmatrix} \beta\bar{y} \\ 0 \end{bmatrix} \tilde{y}_t.$$

Now multiply by the matrix of left eigenvectors $\begin{bmatrix} \frac{1-\beta}{\bar{c}\beta} & 0 \\ 1 & \beta^{-1} \end{bmatrix}$ to return to the original set of variables:

$$\begin{bmatrix} \xi_t^c \\ \hat{b}_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \frac{\bar{c}\beta}{1-\beta} & 0 \end{bmatrix} \begin{bmatrix} \xi_{t-1}^c \\ \hat{b}_{t-1} \end{bmatrix} + \begin{bmatrix} (1-\beta)\frac{\bar{y}}{c} \\ \beta\bar{y} \end{bmatrix} \tilde{y}_t.$$

Using the definition of ξ_t^c we find after a few steps:

$$\tilde{c}_t = \tilde{c}_{t-1} + (1-\beta)\frac{\bar{y}}{c}\tilde{y}_t, \\ \hat{b}_t = \hat{b}_{t-1} + \beta\bar{y}\tilde{y}_t.$$

The unit-root component of this model is clearly evident from the solution for consumption. Once the system is disturbed it will not return to its initial level. In fact, it will tend toward $\pm\infty$ with probability one, which raises doubts about the validity of the linearization approach in the first place. Moreover, there is no limiting distribution for the endogenous variables; the variance of

consumption, for instance, is infinite. Strictly speaking, the model cannot be used for business cycle analysis.

Alternatively, one can derive the state-space representation of the solution, that is, expressed in terms of state variables and exogenous shocks. Convenient substitution thus leads to:

$$\begin{aligned}\tilde{c}_t &= \frac{1 - \beta \widehat{b}_{t-1}}{\beta} \frac{\bar{y}}{\bar{c}} + (1 - \beta) \frac{\bar{y}}{\bar{c}} \tilde{y}_t, \\ \frac{\widehat{b}_t}{\bar{c}} &= \frac{\widehat{b}_{t-1}}{\bar{c}} + \beta \frac{\bar{y}}{\bar{c}} \tilde{y}_t.\end{aligned}$$

As in the intertemporal approach to the current account, income innovations only have minor effects on current consumption, but lead to substantial changes in net foreign assets. Purely temporary shocks, therefore, have permanent effects.

Bounding the Eigenvalues

The characteristic equation of a two-by-two matrix A is given by $p(\lambda) = \lambda^2 - tr\lambda + det$, where $tr = trace(A)$ and $det = \det(A)$, are the trace and determinant, respectively. According to the *Schur-Cohn Criterion* (see LaSalle 1986, 27) a necessary and sufficient condition that all roots of this polynomial be inside the unit circle is

$$|det| < 1 \quad \text{and} \quad |tr| < 1 + det.$$

I am also interested in cases in which there is one root inside the unit circle or both roots are outside the unit circle. Conditions for the former can be derived by noting that the eigenvalues of the inverse of a matrix are equal to the inverse eigenvalues of the original matrix. Define $B = A^{-1}$. Then $trace(B) = \frac{trace(A)}{\det(A)}$ and $\det(B) = \frac{1}{\det(A)}$. By Schur-Cohn, B has two eigenvalues inside the unit circle (and therefore both of A 's eigenvalues are outside) if and only if $|\det(B)| < 1$ and $|trace(B)| < 1 + \det(B)$. Substituting the above expressions, I find that $\left| \frac{1}{\det(A)} \right| < 1$, which implies $|\det(A)| > 1$.

The second condition is $-\left(1 + \frac{1}{\det(A)}\right) < \frac{trace(A)}{\det(A)} < 1 + \frac{1}{\det(A)}$. Suppose first that $\det(A) > 0$. It follows immediately that $|trace(A)| < 1 + \det(A)$. Alternatively, if $\det(A) < 0$, I have $|trace(A)| < -(1 + \det(A))$. However, since I have restricted $|\det(A)| > 1$, the latter case collapses into the former for $\det(A) < -1$. Combining these restrictions I can then deduce that a necessary and sufficient condition for both roots lying outside the unit circle is

$$|det| > 1 \quad \text{and} \quad |tr| < 1 + det.$$

Conditions for the case of one root inside and one root outside the unit circle can be found by including all possibilities not covered by the previous

ones. Consequently, I find this requires

$$\begin{aligned} \text{Either } |det| &< 1 \text{ and } |tr| > 1 + det, \\ \text{or } |det| &> 1 \text{ and } |tr| > 1 + det. \end{aligned}$$

As a side note, employing the Schur-Cohn criterion and its corollaries is preferable to using Descartes' Rule of Sign or the Fourier-Budan theorem since I may have to deal with complex eigenvalues (see Barbeau 1989, 170). Moreover, the former can give misleading bounds since it does not treat $det < -1$ as a separate restriction. This is not a problem in the canonical model where $det = \beta^{-1} > 1$, but may be relevant in the other models.

Proof of Proposition 3

With no internalization of the risk premium, the linearized equation system is given by

$$\begin{aligned} \tilde{c}_t &= \tilde{c}_{t-1} - \sigma\psi\hat{b}_{t-1}, \\ \bar{c}\tilde{c}_t + \hat{b}_t &= R^*(1 - \psi\bar{b})\hat{b}_{t-1}. \end{aligned}$$

Its trace and determinant are $tr = 1 + R^*(1 - \psi\bar{b}) + \sigma\psi\bar{c}$ and $det = R^*(1 - \psi\bar{b})$. Since I have $tr = 1 + det + \sigma\psi\bar{c} > 1 + det$, it follows immediately that the system contains one stable and one unstable root, so that the equilibrium is unique for all parameter values.

With internalization of the risk premium, the linearized equation system is given by

$$\begin{aligned} \tilde{c}_t &= \tilde{c}_{t-1} - \sigma\psi(1 + \beta R^*)\hat{b}_{t-1}, \\ \bar{c}\tilde{c}_t + \hat{b}_t &= R^*(1 - \psi\bar{b})\hat{b}_{t-1}. \end{aligned}$$

Its trace and determinant are $tr = 1 + R^*(1 - \psi\bar{b}) + \sigma\psi\bar{c}(1 + \beta R^*)$ and $det = R^*(1 - \psi\bar{b})$. Since I have $tr = 1 + det + \sigma\psi\bar{c}(1 + \beta R^*) > 1 + det$, it follows immediately that the system contains one stable and one unstable root, so that the equilibrium is unique for all parameter values. This concludes the proof of the proposition.

Proof of Proposition 4

With no internalization of the risk premium, the linearized equation system is given by

$$\begin{aligned} \tilde{c}_t &= \tilde{c}_{t-1} - \frac{\sigma\beta\psi}{y}\hat{b}_{t-1}, \\ \bar{c}\tilde{c}_t + \hat{b}_t &= \left(\frac{1}{\beta} - \psi\frac{\bar{b}}{y}\right)\hat{b}_{t-1}. \end{aligned}$$

Its trace and determinant are $tr = 1 + \sigma\beta\psi\frac{\bar{c}}{\bar{y}} + \frac{1}{\beta} - \psi\frac{\bar{b}}{\bar{y}}$ and $det = \frac{1}{\beta} - \psi\frac{\bar{b}}{\bar{y}}$. Since I have $tr = 1 + det + \sigma\beta\psi\frac{\bar{c}}{\bar{y}} > 1 + det$, it follows immediately that the system contains one stable and one unstable root, so that the equilibrium is unique for all parameter values.

With internalization of the risk premium, the linearized equation system is given by

$$\begin{aligned}\tilde{c}_t &= \tilde{c}_{t-1} - 2\frac{\sigma\beta\psi}{\bar{y}}\hat{b}_{t-1}, \\ \tilde{c}_t + \hat{b}_t &= \left(\frac{1}{\beta} - \psi\frac{\bar{b}}{\bar{y}}\right)\hat{b}_{t-1}.\end{aligned}$$

Its trace and determinant are $tr = 1 + 2\sigma\beta\psi\frac{\bar{c}}{\bar{y}} + \frac{1}{\beta} - \psi\frac{\bar{b}}{\bar{y}}$ and $det = \frac{1}{\beta} - \psi\frac{\bar{b}}{\bar{y}}$. Since I have $tr = 1 + det + 2\sigma\beta\psi\frac{\bar{c}}{\bar{y}} > 1 + det$, it follows immediately that the system contains one stable and one unstable root, so that the equilibrium is unique for all parameter values. This concludes the proof of the proposition.

REFERENCES

- Barbeau, Edward J. 1989. *Polynomials*. New York, NY: Springer-Verlag.
- Baxter, Marianne, and Mario J. Crucini. 1995. "Business Cycles and the Asset Structure of Foreign Trade." *International Economic Review* 36 (4): 821–54.
- Dotsey, Michael, and Ching Sheng Mao. 1992. "How Well Do Linear Approximation Methods Work? The Production Tax Case." *Journal of Monetary Economics* 29 (1): 25–58.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86 (6): 971–88.
- LaSalle, Joseph P. 1986. *The Stability and Control of Discrete Processes*. New York, NY: Springer-Verlag.
- Nason, James M., and John H. Rogers. 2006. "The Present-Value Model of the Current Account Has Been Rejected: Round Up the Usual Suspects." *Journal of International Economics* 68 (1): 159–87.
- Neumeyer, Pablo A., and Fabrizio Perri. 2005. "Business Cycles in Emerging Economies: The Role of Interest Rates." *Journal of Monetary Economics* 52 (2): 345–80.

- Schmitt-Grohé, Stephanie. 1997. "Comparing Four Models of Aggregate Fluctuations Due to Self-Fulfilling Expectations." *Journal of Economic Theory* 72 (1): 96–147.
- Schmitt-Grohé, Stephanie, and Martín Uribe. 2003. "Closing Small Open Economy Models." *Journal of International Economics* 61(1): 163–85.
- Senhadji, Abdelhak S. 2003. "External Shocks and Debt Accumulation in a Small Open Economy." *Review of Economic Dynamics* 6 (1): 207–39.
- Sims, Christopher A. 2002. "Solving Linear Rational Expectations Models." *Computational Economics* 20 (1–2): 1–20.