Towards a Dynamic Price Index

Ricardo Reis*

Princeton University, NBER, and CEPR

June 2005

Abstract

The existing theory on cost-of-living price indices is built on the classical static model of consumption with certainty. This paper studies the construction of a price index that measures the cost of living for a consumer who lives for more than one period and faces uncertainty. It proposes one useful definition of a price index in these conditions: the DPI, for dynamic price index. The DPI has several interesting properties. First, the structure of financial markets is a key input into determining whether higher consumer prices raise or lower the DPI. Second, asset prices should be included in the DPI, with a potentially large weight. Third, if goods are durable then the slower they depreciate, the larger their weight should be in the DPI. This paper takes a first pass at quantitatively building a DPI for the post-war U.S. economy and finds that it provides a very different account of the inflation experience during this period.

JEL classification: C43, D91, E31.

Keywords: Consumer price index; Cost-of-living index; Inflation; Intertemporal consumption choices.

*I am grateful to Stephen Cecchetti, Angus Deaton, Tim Erickson, Charles Goodhart, Robert Shiller, and Jon Steinsson for their useful comments and to David Johnson and Richard Bahr for help with the data from the Bureau of Labor and Statistics. Alisdair McKay provided excellent research assistance. Contact: rreis@princeton.edu. First draft: April 2005.

1 Introduction

Price indices are present everywhere in economic discourse. Journalists refer to them daily, policymakers monitor them closely, and much of macroeconomics is about explaining their ups and downs. While there are many alternative ways of aggregating all of the prices in the economy into a single number, it is the consumer price index (CPI) that typically gets the most attention. In most countries, the CPI measures the change in expenditure required to buy a fixed basket of goods when facing a new set of prices, as Laspeyres or Paasche suggested.

Economists have for long criticized this approach since, in general, consumers will substitute across goods in response to changes in prices. Rather than keeping *baskets* fixed, economists have proposed keeping *utility* fixed instead. The resulting cost-of-living price index dates back to Konus (1924), who proposed measuring "the relative change occurring in the monetary cost of those consumers' goods which are necessary for the maintenance of a certain standard of living." Konus (1924) coupled this definition with the assumption that people choose their consumption to maximize utility subject to a budget constraint. Static classical demand theory could then be used to compare the expenditure necessary to reach the same level of utility for two different sets of prices, while taking into account consumers' substitution across goods.

Since Konus's work, economists have made great progress in the use of classical demand theory to construct price indices. Research found that the Laspeyres and Paasche price indices provide bounds on the cost-of-living index and proposed flexible functional forms of prices that approximate the index closely. Further progress was made by characterizing the conditions under which the index is independent of the standard of living and prices may be aggregated in stages. More recently, research has focussed on incorporating changes in tastes and in the quality of the goods available to consumers. In the United States today, the CPI already incorporates many of these advances.[1]

However, the framework behind the economic price index is still the one suggested by Konus. Notably, researchers still assume that people maximize one-period utility with

[1] The International Labor Office (2004) outlines the procedures currently in place in building consumer price indices, while the National Research Council (2002) is an excellent reference for the debates and progress in this literature. The Boskin commission (Boskin et al, 1997) makes a critical assessment of the CPI and a number of suggestions to make it closer to a cost-of-living index; see also Shapiro and Wilcox (1996) and the debate in the 1998 Winter issue of the Journal of Economic Perspectives.

perfect certainty. Over the past few decades though, economists have learned how to model the behavior of consumers who live for many periods in an uncertain world. The aim of this paper is to bring the modern theory of consumption, that explicitly takes dynamics and uncertainty into account, to shed light into the old question of how to build an economic price index. The product of this work is a dynamic price index, or DPI.[2]

Taking dynamics into account in the construction of a price index unveils a new source of substitution bias that the CPI ignores. Consumers that live for more than one period react to higher prices today relative to the future by substituting away from present into future consumption. The DPI takes this intertemporal substitution into account in constructing the adequate intertemporal cost-of-living price index.

There are two further attractions in moving towards the DPI. First, it is natural to use the DPI to compare two dates, whereas the common practice of using a one-period cost-of-living price index to compare dates is theoretically awkward. Second, while all price indices compare current prices to a counterfactual set of prices, the CPI does not naturally pin down one specific counterfactual. The DPI instead, by incorporating in its theoretical foundations the presence of uncertainty on future prices, suggests as a natural counterfactual what consumers expected to happen.

The study of the DPI in this paper proceeds as follows. Section 2 introduces the model of consumption and addresses the conceptual issue of how to define the DPI. The model in section 2 is stated with some generality, so that the definition is widely applicable. Even at this general level, one can characterize some basic properties of the DPI.

Further understanding of the properties of the DPI requires specializing to more restricted versions of the general model. Sections 3 to 5 continue the theoretical investigation, with each section focussing on one of three different classes of prices.

Section 3 focusses on non-durable consumer goods prices. One might expect that the static CPI theory would accurately reflect movements in these prices, but this turns out not to be the case. Section 3 shows that depending on which financial markets are available, higher non-durable goods prices which unambiguously raise the CPI, may raise, leave unchanged, or even lower the DPI.

Section 4 considers asset prices. Assets allow consumers to transfer funds across time

[2]A more appropriate, though more lengthy nomenclature might be DSCPI, for **D**ynamic **S**tochastic **C**onsumer **P**rice **I**ndex.

and their prices reflect the relative price of consumption today relative to consumption in the future. Section 4 shows that asset prices enter the DPI by affecting consumer's willingness and ability to substitute consumption over time. It finds that asset prices' weight in the DPI depends on people's impatience, their risk aversion, and on the persistence of shocks to prices.

Section 5 turns to durable goods prices. These goods combine features of both non-durable goods and assets: a durable good gives well-being in the present and is also a vehicle for investing for the future. Section 5 shows that the more durable a good is in the sense of depreciating at a slower rate, the larger is its weight in the DPI. A good that is very durable may receive a much larger weight than a good that is non-durable but is otherwise identical.

Section 6 turns to applying these theoretical lessons in practice. I construct a DPI for the post-war U.S. economy considering the broad consumption categories in the CPI together with equity and bonds. While this is only a rudimentary first pass at the data, it allows a peek at the main features of the U.S. DPI. Section 6 compares the CPI and the DPI in the post-war and finds that they differ widely on average, especially since the prices of housing and bonds have a much larger weight in the DPI. The DPI provides a provocative alternative account of recent U.S. history.

Finally, section 7 concludes and discusses three related issues: the possible uses of the DPI, its limitations, and suggestions for future research bringing dynamics and uncertainty into price index theory.

A Review of the Literature

The literature using the static Konus approach to building cost-of-living indices is long and distinguished. I refer the reader to the book treatments by Fisher and Shell (1972), Diewert and Montmarquette (1983), and Pollack (1989), and to the excellent survey article by Diewert (2001).

The consideration of intertemporal trade-offs in the context of price indices was, to my knowledge, first articulated by Alchian and Klein (1973).[3] They considered an infinitely-lived consumer who has access to complete financial markets so that there is an Arrow-

[3] Irving Fisher (1906, ch. 10) briefly considered intertemporal trade-offs when he noted that for indexing long-term loan contracts, the right price index should include the price of durable goods.

Debreu futures price associated with all possible future states of the world. They proposed to define the price index in terms of keeping total utility fixed (including utility in the present and in the future), and noted that since the prices of futures contracts give the prices of future claims on consumption, they should be included in the price index. I will consider a similar scenario briefly in section 3. While I will ultimately conclude that in a complete markets Arrow-Debreu world, the DPI is not a very useful measure, Alchian and Klein's insight that asset prices belong in the DPI will be strongly supported by the analysis that follows. I will move further by showing precisely how asset prices enter the DPI and with what weight.

Pollack (1989, chapter 3) states a definition of the intertemporal cost-of-living index similar to that in Alchian and Klein (1973). The comments in the previous paragraph therefore apply. Pollack studies at length whether it is possible to form period sub-indices in an intertemporal context. Many of his results apply to the DPI and should play an important role in its construction in practice. Shibuya (1992) is the only article that I am aware of that considered both uncertainty and dynamics in building a price index. He used very restrictive assumptions though. His model corresponds to one of the special cases in section 4 of this paper.

Goodhart (2001) persuasively argues that asset prices, and especially house prices, should belong in a welfare-based price index. This paper provides a theoretical foundation to many of his comments. More recently, Bajari, Benkard and Krainer (2004) study the impact of a change in house prices on welfare. Some of the features of their analysis will appear in the study of durable goods in this paper. While they modelled dynamics, they did not consider uncertainty however.

Finally, a recent literature has focussed on the role of asset prices in measuring inflation. The starting assumption in this work is that the right measure of welfare is the CPI and asset prices may be useful insofar as they forecast future CPI inflation. This paper instead suggests that regardless of their statistical relation with the CPI, asset prices enter directly the correct welfare measure, the DPI.[4]

[4] Cecchetti et al (2000) provide a good summary of the state of knowledge in this literature. Goodhart and Hofmann (2000) and Stock and Watson (2003) are useful references on the role of asset prices in forecasting inflation, while Bryan, Cecchetti and O'Sullivan (2001) use the prices of assets and other goods to statistically measure the common component in price increases. Finally, a closely related literature studies a different question from that of constructing a cost-of-living index: whether optimal monetary policy should react to asset prices. See Gilchrist and Leahy (2002) and the references therein.

2 Defining the Dynamic Price Index

2.1 Verbal description of the model

I consider a relatively standard model of consumer behavior. An infinitely lived-consumer cares about lifetime utility which is equal to the expected discounted sum of period-utilities. Each period, she obtains utility from consuming non-durable and durable goods.

The consumer allocates her wealth to several uses. She can acquire consumption of non-durables at the price vector P_t, or invest in durables at the price vector R_t. She can also buy or sell several one-period financial assets, which may include bonds which pay \$1 next period for sure, and risky equities which pay some random amount that depends on the state of the world next period. These assets trade at the price vector Q_t. The sources of wealth are the payoffs from these financial assets plus the market value of the stock of durables after depreciation, which can be sold at the re-sale price R_t^S. The consumer may also receive some other income from for instance supplying labor.

At a given date, the consumer faces three types of state variables that influence her plans. First, she enters the period with a stock of wealth W_t accumulated from past savings. Second, she faces the current prices of goods and assets P_t, R_t, and Q_t. And finally, she must take into account the calendar date t if preferences or the depreciation of durables changes over time, and must consider the information that allows her to forecast future prices represented by some sufficient statistic h_t. I assume that prices are the only source of uncertainty (or that the other sources of uncertainty can be traded away in financial markets) in order to focus on price changes, what a price index should be measuring.[5]

The consumer maximizes total lifetime utility subject to a sequence of budget constraints. Her optimal behavior can be summarized by a value function of the state variables: $V(W_t, P_t, R_t, Q_t, t, h_t)$. This function is equal to the consumer's expected lifetime utility at date t, conditional on behaving optimally. It defines the standard of living on which the price index will be based.

The description of the model is complete. Note that one does not have to describe the way markets function or even the behavior of other agents in the economy. The reason is that the problem of constructing a cost-of-living price index is by definition a partial

[5] This is not a limitation of the analysis though; at the end of this section, I will extend the definition of the DPI to allow for non-price sources of uncertainty.

equilibrium problem. The index is supposed to measure the impact on the consumer of price changes that are exogenous to her. The cost of living price index is not a measure of overall welfare in an economy; it does not consider the impact of the changes in prices on firm's profits or on wages paid, nor does it account for their source. To measure the cost of living, all that is necessary is consumer behavior.

To understand most of what follows, the verbal description of the model above suffices. For completeness, the next sub-section will write down the consumer problem mathematically. Readers are invited to skip this and proceed directly to the definition of the DPI.

2.2 The formal model

At date t, the consumer chooses $\{C_{t+i}, S_{t+i}, B_{t+i}\}$ to solve:

$$\max \mathbb{E}_t \left[\sum_{i=0}^{\infty} \beta^i u_{t+i}(C_{t+i}, S_{t+i}) \right] \quad \text{subject to:} \tag{1}$$

$$P_{t+i}C_{t+i} + R_{t+i}S_{t+i} + Q_{t+i}B_{t+i} \leq W_{t+i}, \tag{2}$$

$$W_{t+1+i} = D_{t+1+i}B_{t+i} + R^S_{t+1+i}\Delta_{t+1+i}S_{t+i}, \tag{3}$$

$$W_{t+1+i} \geq 0 \quad \text{all for } i = 0, 1, 2, ... \tag{4}$$

The notation refers to:

C_t - vector of consumption of non-durable goods,

S_t - vector of consumption of durable goods,

B_t - vector of holdings of financial assets,

$\mathbb{E}_t[.]$ - expectations operator conditional on the information captured in h_t,

h_t - vector of variables that are the sufficient statistic for the information at date t,

β - discount factor, which lies between zero and 1,

$u_t(.)$ - felicity function at date t, strictly increasing in both arguments and concave,

P_t - vector of prices of non-durables,

R_t - vector of prices of durables,

Q_t - vector of prices of financial assets,

W_t - wealth entering date t, defined in (3),

D_{t+1} - vector of payoffs at date $t+1$ of assets purchased at date t,

R^S_t - re-sale price of durable at date t,

Δ_t - diagonal matrix with $1 - \delta_{j,t}$ in each element,

$\delta_{j,t}$ - physical depreciation rate of durable j from $t-1$ and t.

Expression (1) states the objective of the consumer, to maximize lifetime utility. At each date, she faces the budget constraint (2) stating that consumption of goods plus the purchase of financial assets cannot exceed the consumer's wealth. Wealth his defined in (3) and equals the payoff from financial assets plus the market value of the remaining stock of durables. The consumer enter date t with some wealth W_t, which may include the market value of some labor income. Finally, (4) is a constraint preventing Ponzi schemes, which in this setting is equivalent to wealth never being negative.

Evaluating total utility at the optimal solution gives the value function $V(W_t, P_t, R_t, Q_t, t, h_t)$. It depends on the stochastic prices as well as on calendar date t to capture the fact that tastes and the durables technology may change over time.

2.3 The DPI and its features

I define the DPI as follows:

Definition 1 *The dynamic price index π_t is the scalar that solves:*

$$V(\pi_t W_t, P_t, R_t, Q_t, t, h_t) = \mathbb{E}_{t-1}\left[V(W_t, P_t, R_t, Q_t, t, h_t)\right] \tag{5}$$

That is, the DPI is the relative increase in wealth that the consumer requires in response to news on prices in order to maintain her well-being at the level which she expected.

There are two important ingredients that go into this definition. The first one is that the consumer's well-being is measured not by period-utility, but rather by *lifetime* utility. If prices are higher today, this does not only affect consumption and utility now, but it also changes the accumulation of durables and wealth and thus it affects consumption and utility in the future. And it is lifetime utility that the consumer ultimately cares about.

Ignoring the effect of a change in prices on future utility may severely bias the measurement of the cost of living. For instance, if the price of durable housing increases but at the same time the price of perishable food falls, it may turn out that the increase in food consumption exactly offsets the fall in the stock of housing so that period utility is unchanged. However, the lower investment in housing will lead to a lower stock of housing in the future and thus lower utility. Focussing on period rather than lifetime utility would mistakenly lead to the conclusion that the cost of living is unchanged, when in fact it has increased.

Another virtue of the lifetime perspective is that it makes it transparent that asset prices must belong in the price index. A change in asset prices today affects the desire to transfer funds across time through savings. In general, it has an impact on consumption today, as well as on consumption at all future dates. Well-being is affected and an adequate cost-of-living index should reflect it.

The second new ingredient is the statement of the counterfactual situation to which we compare current prices. The DPI defines this counterfactual as what the consumer expected the period before. One advantage of this definition is that it gives a natural temporal interpretation to the price index; the DPI measures the impact of the news on prices. Take for instance the example of an announcement today that some price will be higher at some future date. (This would show up as a realization of h_t.) According to the definition above, the DPI would move at the announcement date, but not at the future date when the change takes place. This is appropriate because the announcement leads to change in behavior not just in the future, but also today. Consumers respond today by consuming less and saving more in order to be able to smooth consumption, and perhaps by adjusting their mix of durables and non-durables. They feel the impact of this shock today, so today's price index should reflect it.

Alchian and Klein (1973) and Pollack (1989) instead proposed cost-of-living indices that compared current prices to some imaginary counterfactual. It is difficult to imagine what this counterfactual could be. It is by definition completely unanticipated by people, who are assumed to have not even considered the possibility that such a change could occur. The definition of the counterfactual in terms of what was expected seems more plausible and it is more consistent with the underlying intertemporal model of consumption.

In practice, price indices are often used to compare prices this period with prices last period. Insofar as prices last period are one useful indicator of prices this period, this approach may be consistent with the definition of the DPI. However, if this is justified from thinking of the counterfactual as utility in the previous period, this is inappropriate. Well-being today and well-being last period are not comparable since circumstances (e.g., tastes) are most likely different today that they were before. The price index would be capturing these differences rather than the impact of prices. It is therefore the last period's expectation of today's well-being that belongs on the right-hand side of (5) rather than last period's value function.[6]

The practical difficulty with collecting information each period has led statistical agencies to compute the CPI relative to a reference period that is only infrequently updated. Likewise, it may prove burdensome to update information on expected future conditions at each date. One could then define the DPI with respect to a base date n periods ago by modifying the right-hand side of (5) to have the expectations conditional on the information n periods ago. By taking as the reference date the last period, the definition in (5) becomes akin to the chained price indices that are increasingly popular.

2.4 Some basic properties of the DPI

There are some basic properties of the DPI that hold even at this very general level. It is easy to show that if the consumption of all goods is positive, then the value function falls with P_t and R_t. Since it is also easy to show that well-being increases with wealth, it follows that the DPI increases with the realization of consumer prices. The intuition is straightforward. For a fixed nominal wealth, the higher are prices, the less goods people can afford to buy. Thus the more they need to be compensated to maintain their well-being.

As for financial assets, it is easy to show that the DPI increases with the price of the asset if it is being held in a positive amount. As long as the consumer is saving, higher asset prices (or lower interest rates) imply that the cost of purchasing a lifetime stream of consumption is now higher. Thus the DPI must rise.

A third property of the DPI is that if all prices today turn out to be equal to what was expected last period, the DPI will generally not equal one. The reason is that the DPI keeps actual and expected well-being in line, not actual and expected prices. The DPI will only happen to be one if the value function happens to be linear in prices, which is unlikely if people are averse to risk. With aversion to risk, consumers will take future uncertainty into account in their choices so that if prices turn out to be what was expected, this has an impact on well-being.

[6] If the consumer is mortal and knows for how long she will live, the inadequacy of comparing value functions at different dates is even more stark. Well-being in the last period of life includes only utility that period, whereas well-being one period before includes utility for two periods; the two are not comparable.

There is another property of the DPI that while somewhat specific, is often useful. Samuelson and Swamy (1974) emphasized that it would be desirable to have an index that did not depend on the current level of wealth. For the static cost-of-living index, they showed that homotheticity of preferences is required for this to happen. For the DPI, it is homotheticity of the value function that is required. If the value function is homogeneous of degree γ in wealth, then the DPI is independent of wealth and equal to:[7]

$$\pi_t = \left(\frac{V(1, P_t, R_t, Q_t, t, h_t)}{\mathbb{E}_{t-1}\left[V(1, P_t, R_t, Q_t, t, s_t)\right]} \right)^{-\frac{1}{\gamma}}. \tag{6}$$

Having established what affects the DPI and with what sign, the remainder of this paper is concerned with determining the weights of each price in the DPI. Sections 3 to 5 examine this question in theory, while section 6 constructs in practice a DPI for the United States. Most of what follows will assume that the only source of uncertainty is prices. It is worth taking a short detour to extend the definition of the DPI to consider uncertainty from other sources.

2.5 Extension to non-price uncertainty

Non-price uncertainty could be driven by shocks to tastes, shocks to the wear and tear of durables, or shocks to income, among many other possibilities. Let all of these shocks be captured by some vector z_t.

The modified DPI is then defined as:

$$\mathbb{E}_z\left[V(\pi_t W_t, P_t, R_t, Q_t, t, h_t, z_t)\right] = \mathbb{E}_{t-1}\left[V(W_t, P_t, R_t, Q_t, t, h_t, z_t)\right]. \tag{7}$$

The operator $\mathbb{E}_z[.]$ integrates over the conditional distribution of the z_t given the $t-1$ information and the realization of prices at t, while $\mathbb{E}_{t-1}[.]$ integrates over both z_t and prices. This way, the only difference between the left-hand side and the right-hand side is that the realization of prices at t is known in the left-hand side, but integrated over in forming expectations in the right-hand side.

[7]Whether the value function is homogeneous or not depends on the details of the consumption problem. Typically, homotheticity of the utility function is enough to guarantee homotheticity of the value function.

3 Non-durables prices and the financial market structure

This section focusses on non-durables prices and on the role of financial markets in the DPI. To make the intuition underlying the results clearer, I will use a simpler version of the general model. As will be clear as I proceed, the main lessons do not depend on the specifics of the model.

The details of the model are the following: There is only one consumption good, which is non-durable. The consumer starts at date 0 with some positive wealth W_0. The only source of uncertainty is that at the beginning of date 1, the price of non-durables from date 1 onwards is realized and it can take two values P^H or P^L such that $P^H > P^L$, with probabilities p and $1-p$ respectively. These two states of the world correspond to a bad state in which prices are high forever, and a good state in which prices are low forever. From period 1 onwards there is no uncertainty and the consumer can save by buying a bond which costs β and returns \$1 in the next period.

The question that I ask is when is the DPI at date 1 high or low, in the sense of being larger or smaller than one. Focusing on the DPI when prices are high π^H, when it is larger than one, the consumer is worse off when prices are higher; if it equals one, she is indifferent; and when it is smaller than one, the consumer is better off with higher prices.

The solution of the model from period 1 onwards is straightforward. The consumer faces the same problem at every date, so she chooses to always consume the same amount, C^i. However, this may be different in the bad state $i = H$, or in the good state $i = L$. It is easy to show that $P^i C^i = (1-\beta)W^i$, where W^i is the wealth with which the consumer enters date 1 if state i is realized. Since higher consumption makes the consumer better off, this implies that there are only three possible cases: (i) If $P^H/W^H > P^L/W^L$, then $\pi^H > 1$. (ii) If $P^H/W^H < P^L/W^L$, then $\pi^H < 1$. (iii) And if $P^H/W^H = P^L/W^L$, then $\pi^H = 1$.

The DPI therefore depends on the real wealth of the consumer in each state. Even when prices are high, the DPI may be low if nominal wealth happens to be even higher. Since the only source of uncertainty in this economy are prices, if nominal wealth is different between the two states, this must be because financial assets have made wealth contingent on the realization of prices. It is through this channel that the financial markets available to the consumer affect the DPI. By allowing for nominal wealth to respond to changes in prices, assets allow the consumer to amplify, attenuate, or even reverse the impact of these price changes on welfare and on the DPI.

Consider then how different financial market structures affect the DPI. The first natural case is that of complete markets. In this model, market completeness only requires that there are two assets traded at date 0, one costing Q^H and paying one dollar in state H at date 1, and the other costing Q^L and paying one dollar in state L. In this case, the problem of the consumer at date 0 is:

$$\max_{c_0, W^H, W^L} \left\{ u(C_0) + \frac{\beta}{1-\beta} \left(pu((1-\beta)W^H/P^H) + (1-p)u((1-\beta)W^L/P^L) \right) \right\}$$
$$s.t.\ C_0 + Q^H W^H + Q^L W^L \leq W_0 \tag{8}$$

The first-order conditions for optimality imply that:

$$\frac{Q^H}{Q^L} \times \frac{(1-p)P^H}{pP^L} = \frac{u'((1-\beta)W^H/P^H)}{u'((1-\beta)W^L/P^L)} \tag{9}$$

Using the previous result, it then follows that there are three possible cases: (i) If $Q^H/Q^L > pP^L/\left[(1-p)P^H\right]$, then $\pi^H > 1$. (ii) If $Q^H/Q^L < pP^L/\left[(1-p)P^H\right]$, then $\pi^H < 1$. (iii) And if $Q^H/Q^L = pP^L/\left[(1-p)P^H\right]$, then $\pi^H = 1$.

Starting with the last case, if asset prices are "fair," the consumer fully insures herself at date 0 against the period 1 price shock. Her welfare at period 1 does not depend on the shock and the DPI is one. If the asset that pays when consumer prices are high is too expensive relative to the asset that pays in the good state, then the consumer will under-insure, so that when prices happen to be high, she will be worse off. Conversely, if the asset that pays in the high price state is cheap, she will over-insure. In this case, when consumer prices turn out to be high, the DPI will actually be low.

This leads to the first lesson:

Lesson 1: *If consumers can hedge all consumer price fluctuations in the financial market, then higher consumer prices can lead to a higher, unchanged, or even lower DPI, depending on whether asset prices are fair or not.*

This reversal of common intuition can be even more extreme if financial markets are incomplete. Consider a second scenario in which the only asset available to the consumer pays in the high-price state. In this case, the consumer cannot transfer any wealth into the low-prices state, so her consumption will be zero forever if consumer prices turn out to be low. If consumer prices are high though, as long as the price of the asset is finite, she will purchase some positive quantity of this asset to sustain a positive consumption into the future. It then follows that regardless of the relative consumer prices, of their relative probability, and of the price of the asset, high prices means good times: $\pi^H < 1$.

This example may be extreme. But it shows that when markets are incomplete, the DPI can move in seemingly paradoxical ways. Even if we know all asset prices and can conjecture all of the possible consumer prices that may occur as well as the chances that they will happen, still, without knowledge of the financial market structure, we will be unable to predict how the DPI will respond to consumer goods prices. There are many ways in which people enter contracts with nominal payoffs that vary with consumer prices. Indexation clauses written into labor contracts are one example as are indexed bonds or the provisions in many life insurance and pension contracts. If tracking different people's tastes was already a challenge for the static cost-of-living index, tracking their financial contracts makes it even more daunting to construct an accurate dynamic cost-of-living index.

Lesson 2: *If consumers can only trade a limited set of assets that do not span all possible price contingencies, then only with a precise knowledge of which assets are available can one predict how the DPI responds to consumer prices.*

There is one case of incomplete markets though for which there are more definite conclusions. This is the third and final scenario that I consider. It is the case in which consumers do not have access to any assets whose payoffs depend on consumer prices. From the perspective of building a DPI for a whole economy, this may be the relevant scenario. Changes in the price of consumer goods are an aggregate shock, which the economy cannot insure against within itself. Trying to insure against this macro risk in the world is currently almost impossible, since financial markets for macro risk-sharing are missing (Shiller, 1993).

In terms of the model, this translates into the consumer being able to trade in only one asset, which returns the same payoff regardless of the realization of consumer prices in period 1. The consumer is then constrained to enter period 1 with the same nominal wealth, regardless of the state of the world. This then implies that higher consumer prices necessarily lead to a rise in the DPI, i.e., $\pi^H > 1$. The third lesson is:

Lesson 3: *If assets' payoffs do not depend on consumer prices, then higher consumer prices necessarily lead to a higher DPI.*

4 Asset Prices

Asset prices played a role in the previous section in determining how random goods prices affect the DPI with complete financial markets. If asset prices are random as well, they affect the DPI directly by changing the relative price of consumption between dates. This section studies the weight of random asset prices in the DPI.

Again to focus on the object of study, I consider a simpler version of the general model. I assume that the consumer has access to only two assets, bonds and equity. This assumption is less restrictive than what it may seem, since we can see equity as a market portfolio including many assets held in optimal proportions. The bond pays one dollar next period with certainty and trades at the stochastic positive price $Q_{0,t}$. Note that holding this bond carries risk: first because its real payoff is uncertain, and second because holding it for more than one period involves a re-financing risk as bond prices change. Equity trades at the price $Q_{E,t}$ and pays a stochastic dividend D_{t+1} that is i.i.d. over time. For both asset prices, it is useful to separate common from idiosyncratic movements. Letting $Q_{0,t} = Q_t \hat{Q}_{0,t}$ and $Q_{E,t} = Q_t \hat{Q}_{E,t}$, then Q_t is the common component and $\hat{Q}_{0,t}$ and $\hat{Q}_{E,t}$ are the idiosyncratic independent components.

A further simplification is to assume that there is only one non-durable good, trading at the stochastic positive price P_t, and that utility equals the natural logarithm of consumption of this good. The consumer's problem is therefore, in dynamic programming form:

$$\begin{aligned} V(W_t, ...) &= \max_{C_t, B_{E,t+1}} \{\log(C_t) + \beta \mathbb{E}_t [V(W_{t+1}, ...)]\} \\ s.t.: P_t C_t + Q_{0,t} W_{t+1} &+ (Q_{E,t} - Q_{0,t} D_{t+1}) B_{E,t+1} \leq W_t. \end{aligned} \tag{10}$$

I consider a few possibilities for the stochastic properties of prices. The first case is when prices are i.i.d. The appendix solves this problem and finds that, focussing only on the terms in consumer prices and the common component of asset prices:

$$\ln(\pi_t) = (1 - \beta) \ln(P_t) + \beta \ln(Q_t). \tag{11}$$

The DPI therefore increases less than proportionally with an increase in consumer prices. If the discount factor is close to one, this increase may be very small. In contrast, the static CPI moves proportionally one-to-one with consumer prices. Behind this attenuation is the ability of the consumer to substitute consumption intertemporally. Higher prices today have a smaller impact on welfare if the consumer is able to respond by borrowing from the future to sustain a smooth consumption path.

Asset prices not only receive a positive weight in the DPI, but as long as the discount factor is close to one, this weight is much higher than that of consumer prices. For $\beta = 0.9$, a reasonable discount factor for one year, asset prices receive a weight that is 9 times larger than goods prices in the DPI. By ignoring asset prices, the static CPI misses out on the largest share of the appropriate dynamic cost-of-living index.

Missing from the full expression for the DPI in (11) is a messy term involving the idiosyncratic component of asset price shocks. The appendix shows that, conditional on the optimal holdings of bonds and equity, B_0^* and B_E^*, the elasticity of the DPI with respect to idiosyncratic asset price shocks is: $\partial \ln(\pi_t)/\partial \ln(\hat{Q}_{0,t}) = B_{0,t}^* Q_{0,t}/W_t$ and $\partial \ln(\pi_t)/\partial \ln(\hat{Q}_{E,t}) = B_{E,t}^* Q_{E,t}/W_t$. Each individual asset is therefore weighted according to its weight in the consumer's portfolio. The more of an asset the consumer is optimally holding, the larger is the impact of a change in its price on welfare.

The lesson from the i.i.d. case is therefore that:

Lesson 4: *If asset and goods prices are equally persistent, asset prices receive a considerably larger weight than goods prices in the DPI. The price of each asset receives a weight proportional to its portfolio share.*

The second scenario that I consider allows for persistent shocks to prices. Consumer prices in the post-war U.S. approximately follow an ARIMA(1,1,0) process in logs. That is, in the post-war, $\ln(P_t) - \ln(P_{t-1}) = \eta(\ln(P_{t-1}) - \ln(P_{t-2})) + \varepsilon_t^P$, and the least-squares estimate of η is about 0.84. In this case, the appendix shows that the elasticity of the DPI with respect to the news on consumer prices equals $1/(1 - \beta\eta)$.

If consumer prices follow a random walk ($\eta = 0$), their weight in the DPI is therefore exactly equal to one. This confirms the previous intuition that the weight of consumer prices in the DPI is affected by the ability to intertemporally smooth consumption. With $\eta = 0$, shocks to consumer prices are permanent so the permanent income hypothesis dictates that consumption changes at all dates by the same amount without any intertemporal borrowing or lending. The intertemporal substitution effect is not present so the source of substitution bias in the CPI is absent and it coincides with the DPI.

As long as prices do not exactly follow a random walk though, the CPI and the DPI differ. The more persistent are shocks to consumer prices, the larger their impact on the DPI.[8] For the statistical process that consumer prices follow in the post-war U.S., and if the discount factor is about 0.9, then a 1% increase in consumer prices raises the DPI by about 4%. This leads to the new conclusion:

Lesson 5: *If goods prices follow a random walk, then their weight in the DPI is the same as in the CPI. If the persistence of shocks to goods prices is larger (smaller), their weight in the DPI is larger (smaller) than in the CPI, because the consumer responds to these shocks by intertemporally substituting consumption.*

The final issue that I investigate is the expected trend in the DPI. To focus on this issue, I abstract away from idiosyncratic asset price shocks, but I do not restrict the stochastic process followed by goods prices and the common component of asset prices. In this third and last scenario, the appendix shows that if at date t the news on goods prices from t onwards exactly coincide with was expected at $t-1$, then $\pi_t > 1$. If goods prices turn out to be precisely what was expected, the DPI therefore rises. An alternative way to state this property is that the DPI is concave in goods prices.[9] This leads to the last lesson:

Lesson 6: *If prices turn out to equal what they were expected to be, the DPI is above one.*

5 Durable goods

The last category of prices is those referring to durable goods. These goods are particularly interesting, because they combine features of both goods and assets; they yield utility and they also transfer wealth across time. This section studies their role in the DPI.

To focus on the role of durables, consider the following simpler consumer problem:

$$\begin{aligned} V(W_t, P_t, R_t, Q_t) &= \max_{C_t, S_t} \{(1-\alpha)\ln(C_t) + \alpha \ln(S_t) + \beta \mathbb{E}_t \left[V\left(W_{t+1}, P_{t+1}, R_{t+1}, Q_{t+1}\right)\right]\} \\ s.t. &: P_t C_t + \left(R_t - Q_t(1-\delta)R^S_{t+1}\right) S_t + Q_t W_{t+1} \leq W_t \end{aligned} \tag{12}$$

Relative to the general DPI problem, here I specialize to: only one of each of durable and

[8] From the perspective of policy, lowering the persistence of goods prices will reduce the impact of price shocks on welfare.

[9] The concavity of the DPI implies that an increase in the variability of prices lowers the DPI. That is, more variable prices on average raise welfare—an old result in price theory.

non-durable goods, a Cobb-Douglas utility function, a bond as the only financial asset, and i.i.d. prices P_t, R_t, Q_t, and R_t^S. As before, these simplifying assumptions are used to make the intuition behind these results more transparent.

Since the purpose is to understand how durables differ from non-durables, I focus attention on a particular quantity: the relative marginal weight of durables relative to non-durables in the DPI, defined as $M_t = (\partial \ln(\pi_t)/\partial \ln(R_t)) / (\partial \ln(\pi_t)/\partial \ln(P_t))$. This should measure whether, up to a first order, the DPI is more responsive to changes in durables or in non-durables prices.

The consumer problem does not have an analytic solution in general. In a steady state in which all asset prices equal their expected value, denoted by upper bars, the appendix shows that:

$$\bar{M} = \frac{\alpha}{(1-\alpha)(1-\beta(1-\delta)\bar{R}^S/\bar{R})}. \tag{13}$$

This expression shows that the larger is the weight of durables in period utility (α), the larger their weight on the DPI. This is natural, since the more durables matter to the consumer, the more changes in their price affect her welfare and so the DPI. The expression also shows that the more patient is the consumer (β), the more durables matter in the DPI. A more patient consumer saves more, part of which on durables, which increases the importance of durables for consumer welfare.

Expression (13) also shows that $\bar{M}$ increases with $\bar{R}^S/\bar{R}$. The larger is the expected re-sale price of durables relative to their current price, the larger is the return from holding them. Durables become a more attractive investment so the consumer buys more of them and is therefore more sensitive to fluctuations in their price. The appendix further shows that in a neighborhood of this steady state, if the return on durables exceeds the return on financial assets, M_t is higher. In this case, durables not only yield utility but also dominate assets as a vehicle for savings.

These results lead to the conclusion:

Lesson 7: *The relative marginal weight of a durable relative to a non-durable increases with the weight of durables in period utility, consumer's patience, and the expected re-sale relative to the purchasing value of the durable.*

The parameter δ measures how durable a good is. In the special case when it equals 1, the good is non-durable, while if it equals zero it lasts forever; in between, the smaller is δ

Figure 1: Relative marginal weight of durables in the DPI as a function of depreciation

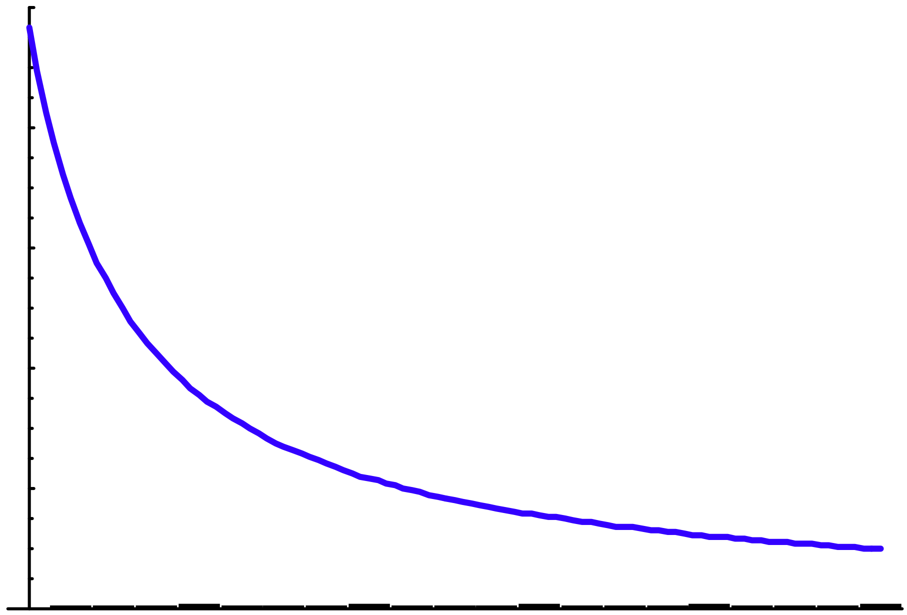

the more durable is the good. It is easy to see from (13) that the relative marginal weight of durables increases with durability.

Figure 1 investigates this result without resorting to approximations. It plots M_t against δ under the assumption that R_t^S is log-normally distributed with mean R_t and standard deviation 0.1, and for the parameter values: $\alpha = 0.5$, $\beta = 0.9$, and $Q_t = \beta$.

At the extreme right of the figure is the case where there is full depreciation so that S_t is a non-durable. In this case, the relative weight between durables and non-durables is $(1 - \alpha)/\alpha$, which equals 1 for the chosen parameter values. As the durability of S_t increases though, its marginal impact on the DPI becomes increasingly larger than that of non-durables. Durability makes S_t become closer to an asset and consequently its weight in the DPI increases towards the weight that assets receive. When δ gets close to 0, the weight of non-durables becomes close to that of assets, and at the extreme when the stock of durables does not depreciate, durables prices receive a larger weight than asset prices do. Intuitively, if $(1 - \delta)R_{t+1}^S/R_t$ becomes close to or larger than $1/Q_t$, investing in a durable gives a higher return that investing in assets does. The consumer then holds a large stock of durables both for consumption and as a vehicle for savings, so a marginal change in their price has a large effect on welfare.

Lesson 8: *Durables goods prices receive a larger weight than non-durables in the DPI. The more durable a good is, the larger its weight in the DPI.*

One of the most important durable goods purchased by consumers is housing. There are some issues that arise specifically with housing. For instance, a consumer may own or rent a home. The discussion so far has considered the perspective of a homeowner. For a household that instead rents, housing takes the role of a non-durable good and the analysis in section 3 applies. Another important issue that arises with housing is that one household's purchase is often another household's sale. While this does not affect the analysis above when it comes to the DPI for each household, it does bring additional complications when it comes to aggregating across consumers towards a single price index. Tackling these issues is beyond the scope of this paper, though see Goodhart (2001) and Bajari, Benkard and Krainer (2004) for progress on this front.

6 A first pass at the post-war U.S. DPI

This section takes a first step at building a DPI for an actual economy. Considering in detail all of the prices in an economy is well beyond the scope of this section. Moreover, there are many practical complications that the model in (1) -(4) ignores.[10] In this section, I take a simple first pass at the problem that considers only a few broad categories of goods. This serves the purpose of demonstrating how to go about building a DPI, as well as detecting some of the broad trends in the U.S. DPI.

6.1 The model of the U.S. economy

I consider the following version of the general DPI model:[11]

[10] Dealing with taxes is a particularly important one. Triplett (1983) has a very lucid discussion of this and related issues.

[11] I calculate the DPI as if there was a representative agent of the U.S. economy. In general, the DPI does not require this assumption. In the same way that a recent literature has evaluated welfare in economies with heterogenous agents, so can one calculate the DPI in these circumstances. This is an interesting area for future research, but it is beyond the scope of this paper.

$$V(W_t, ...) = \max_{C_t, S_t, B_{E,t}} \mathbb{E}_t \left[\sum_{i=0}^{\infty} \beta^i \left(\sum_{j=1}^{4} \alpha_{j,t+i} \ln(C_{j,t+i}) + \sum_{j=5}^{6} \alpha_{j,t+i} \ln(S_{j,t+i}) \right) \right] \quad \text{subject to:}$$

$$\sum_{j=1}^{4} P_{j,t+i} C_{j,t+i} + \sum_{j=5}^{6} R_{j,t+i} S_{j,t+i} + Q_{0,t+i} B_{0,t+i} + Q_{E,t+i} B_{E,t+i} \leq W_{t+i}, \tag{14}$$

$$W_{t+1+i} = B_{0,t+i} + D_{t+1+i} B_{E,t+i} + \sum_{j=5}^{6} R^S_{j,t+1+i} (1 - \delta_j) S_{j,t+i}, \quad \text{for } i = 0, 1, 2, ...$$

Within non-durables, I consider four sets of goods: food, energy, services (without shelter or energy), and other non-durables. Durables comprise two goods: shelter and other durables. These 6 categories include 100% of the goods in the CPI.

A period refers to one year. The depreciation rates for durables come from the Bureau of Economic Analysis' fixed assets tables. I obtain the depreciation rate per year by dividing the amount which the asset depreciates in one year by its net stock. For consumer durables, the average depreciation since 1957 is 0.211. For housing, the average depreciation rate using the available data from 1987 to 2003 is 0.016. In both cases, depreciation rates vary very little from year to year, justifying the assumption of a constant rate.

The weight $\alpha_{j,t}$ equals the expenditure share on good j. I measure these using the relative importance of each good in the last revision of the CPI—the appendix describes the details. These weights are allowed to vary over time, to take into account the main trends in consumption in the last 50 years. For instance, the weight of shelter plus services in the consumer basket has risen from 29% in 1964 to 56% in 2004. The Bureau of Labour and Statistics has revised these weights six times since 1964; table 1 shows the revisions for the categories of goods that I consider. In the model, I assume that the consumer learns about these taste changes only when they occur.[12]

[12] There are several alternatives in describing what the consumer knows about taste changes. One possibility is to assume that the consumer at the beginning of the sample is perfectly aware of the future changes in the $\alpha_{j,t}$ and at the end of the sample she presumes that from then on into the infinite future the $\alpha_{j,t}$ will remain unchanged. This leads to very similar results to the ones in the text. Another possibility is to statistically model uncertainty on tastes and have the consumer take this uncertainty on her future tastes into account when forming expectations. This complicates the optimization problem considerably and requires taking into account non-price sorces of uncertainty as in section 2.5.

Table 1. Relative weights of different components in the CPI

CPI weights	Food	Energy	Services	Other non-durables	Shelter	Other durables
1964-1977	0.224	0.067	0.185	0.210	0.202	0.115
1978-1982	0.177	0.086	0.190	0.131	0.291	0.124
1983-1986	0.190	0.124	0.213	0.131	0.213	0.128
1987-1997	0.162	0.074	0.229	0.139	0.277	0.120
1998-2001	0.153	0.070	0.238	0.125	0.298	0.116
2002-2003	0.147	0.062	0.237	0.118	0.315	0.121
2004	0.144	0.071	0.234	0.110	0.329	0.113

The time series for the prices of different goods are those produced by the Bureau of Labor and Statistics. It is difficult to obtain reliable re-sale prices for durables, so I assume that they equal the market price for a new purchase. The growth rate of prices is plotted in figure 2.

There are two assets, equity and bonds. Returns on equity equal D_{t+1}/Q_t^E in the model and I measure them using the value-weighted index of the stocks in the NYSE, AMEX and NASDAQ compiled by the Center for Research in Security Prices. In the model, the price of bonds equals the inverse of bond returns. I measure them using the Fama and Bliss (1989) series on the yield of zero-coupon one year Treasury bonds. Figure 3 plots the two series.[13]

The key input for the DPI is not the level of prices, but rather the unexpected shocks from the perspective of one year before. To obtain these, I model the demeaned first-difference of the log of the six goods prices and the two asset returns as following a first-order VAR. There may be some scope for improvement in this forecasting model. Using annual post-war U.S. data though, a VAR(1) already involves estimating 64 different parameters; it is challenging to estimate even more sophisticated models of the joint dynamics of these eight time-series.

I set the discount factor equal to 0.94 to match the inverse average return on Treasury bonds. The average returns on equity are calibrated to match equity's portfolio share in the U.S.'s portfolio. Since 1966, the value of equities and the value of outstanding Treasury

[13]I tried a few alternative measures. Measuring the returns on equity on the S&P 500, and the return on bonds using different maturities led to very similar results.

Figure 2: Growth in the prices of separate categories of goods 1964-2004

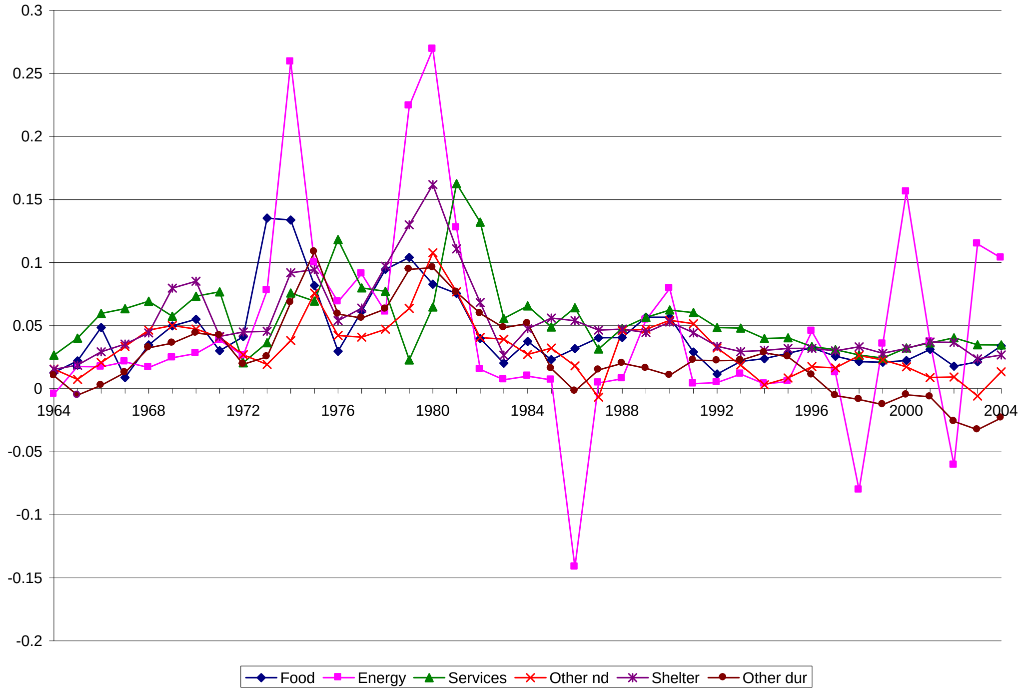

Figure 3: Asset returns 1964-2004

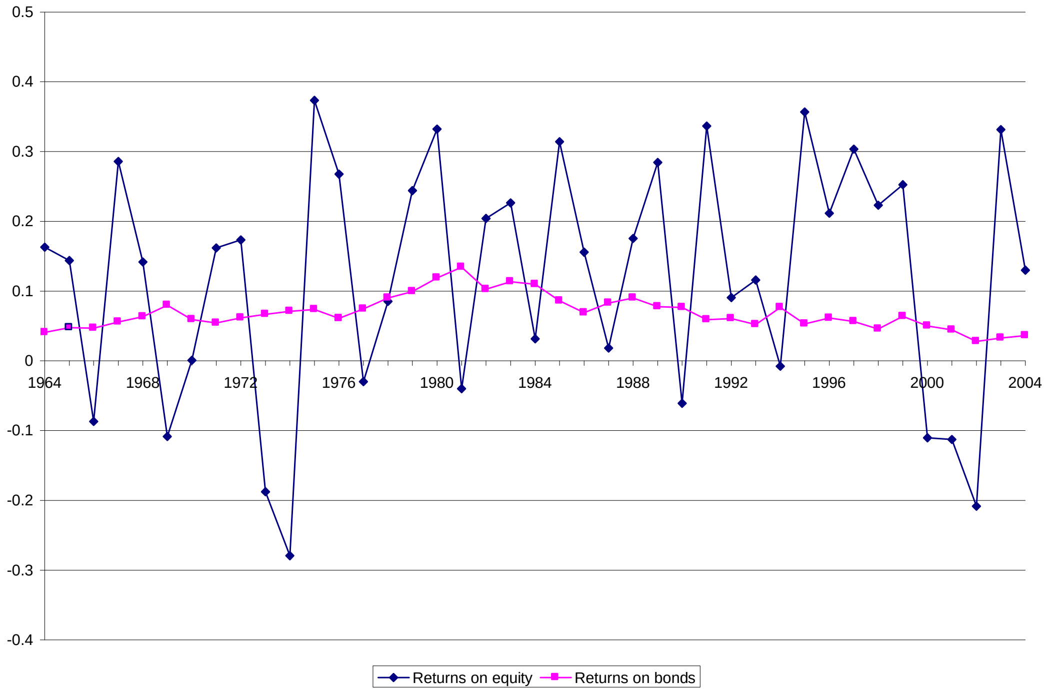

securities imply that 59% of financial investments are on equity. Calibrating equity returns to match this value implies average returns below what we have historically observed. The high return on equity in the data would imply a portfolio share of equity above 100%—the well-known equity premium puzzle. I calibrate the parameters in this way because while the portfolio share of equity is important for the dynamics of the DPI, the average returns are not.

Given these assumptions, the appendix shows how to compute a first order log-linear approximation of the DPI. Feeding in the actual realizations of prices in the United States from 1964 to 2004 gives a time series for the DPI.

6.2 The U.S. DPI from 1964 to 2004

Figure 4 plots the log of the U.S. DPI. Recall that the DPI measures the percentage increase in nominal wealth that consumers would require to compensate them for the increase in prices. Figure 4 also plots the log of the CPI at each date relative to the previous date.[14]

There are three clear negative periods according to the DPI. The first is the period from 1973 to 1974 in which the DPI peaked at 22%. These were particularly bad years due to the combination of high consumer prices following the oil price shocks and high asset prices (or low returns). The next large increase in the DPI takes place in 1990, when it reached 14%. Again high oil prices associated with the first Iraq war are the main culprit. The final negative episode are the three years from 2000 to 2002. The source of this rise is different. The increase in the price of housing and the low returns on equity are the main driving forces now. In 2002, consumers required a 22% increase in their wealth to compensate them for the unexpectedly high prices.

As for good times, figure 4 shows that for the DPI was particularly low in the 1990s. The DPI was negative in 8 years in this decade, confirming the common impression that the cost of living was low in the 1990s.

Figure 4 shows that in general, the DPI and the CPI move in different directions. The correlation coefficient between the two series is only 0.11. An important difference visible in figure 4 is that whereas the DPI is close to uncorrelated, the CPI is very persistent. This

[14]Note that the DPI is defined in (5) using the previous year as the base year. The CPI is typically presented using a fixed date as the base period instead. The conceptual equivalent of the DPI is the change in the CPI from one year to the next.

Figure 4: The DPI and the CPI in the United States 1964-2004

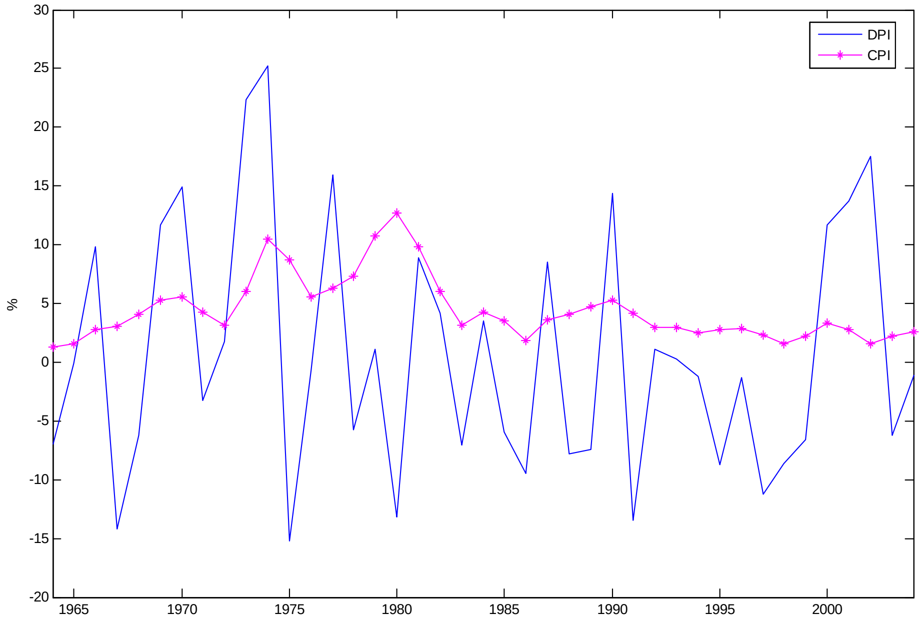

reveals a crucial distinction between the two price indices. Movements in an individual price series $p_{j,t}$ show up in the CPI when they occur. That is, decomposing $p_{j,t} = \sum_i (p_{j,t-i} - p_{j,t-i-1})$, the CPI responds at date $t-i$ to the $p_{j,t-i} - p_{j,t-i-1}$ component of the sum. The relevant decomposition for the DPI is instead Wold's decomposition of the price series into its innovations (its moving average decomposition). The DPI responds to movements in prices when they are *discovered*, not when they actually take place. Because a sequence of surprises all in the same direction is unlikely, the DPI will tend to be serially uncorrelated.

The behavior of the price indices in the late 1970s illustrates this difference. During this period, prices kept on increasing so the CPI was high for the entire period. However, while people were surprised by the increase in prices following the first oil price shock, if they formed expectations optimally, they should have anticipated further increases in prices. Even though prices were high in 1978 or 1979, they were lower than expected. Relative to their expectations, people were positively surprised, and so the DPI fell.

Goodhart (2001) and Bryan, Cecchetti and O'Sullivan (2001) expressed the concern that any price index that included asset prices would be very volatile. Figure 4 partially justifies this prediction: the DPI is considerably more volatile than the CPI. (Though less so than what one might have expected.) The standard deviation of the DPI is 10% whereas the CPI's standard deviation is only 3%.[15]

Table 2 investigates the contribution of each component to the DPI's volatility. The first row shows the standard deviation of each component of the DPI. As expected, equity returns are extremely volatile, as is the price of energy.

To calculate the impact of news on each component of the DPI on the index would require identifying these structural shocks. In a VAR with 8 time-series, this would require a large number of identifying assumptions. The second row in table 2 instead shows the marginal impact on the DPI, averaged over the sample, of a 1% shock to each component of the DPI, if all goods prices are assumed to follow independent random walks and all returns are i.i.d. This corresponds to all of the coefficients in the forecasting VAR being zero and the variance-covariance matrix of innovations being diagonal.

In this case, the marginal impact of a non-durable good's price equals the sample average of its expenditure share in table 1 (recall lesson 5). The table reveals that while energy prices are volatile, they have a small impact on the DPI, because they account for little in people's expenditures. Among durables, shelter has a much larger impact on the DPI than its expenditure share. Indeed, a 1% raise in the price of shelter has almost the same impact on the DPI as a 1% increase in all of the other 5 goods prices. The reason for the importance of shelter was uncovered in section 5: shelter is very durable, depreciating by little every year.

The next row in table 2 shows the marginal impact of news to each price if the first difference in prices and returns all follow independent AR(1)'s. The residuals of each of the autoregressions can still be interpreted as structural shocks to each price because they are independent. Section 4 showed that the more persistent are shocks to prices, the larger their weight in the DPI. The impact of all prices but one increases since price shocks are typically very persistent. The exception is equity prices, which as is well know follow a random walk. This result shows that while equity prices are very volatile, they actually have a modest impact on the DPI. Rather, among asset prices, it is bond prices that dominate. Overall, a 1% increase in the prices of bonds and shelter raise the DPI by three times more than a

[15] Even if the DPI is volatile, a pertinent question is: if the DPI is the right measure of the cost of living, then why is stability a desirable property?

1% increase in the price of the other 6 components. It is these two prices that drive most of the movements in the DPI.

Table 2. The volatility of the components of the DPI and their average impact

	Food	Energy	Services	Other non-durables
Standard deviation	0.031	0.075	0.028	0.023
Marginal impact with $N = 0$	0.186	0.077	0.209	0.158
Marginal impact with diagonal N	0.496	0.137	0.410	0.467
	Shelter	Other durables	Equity	Bonds
Standard deviation	0.032	0.032	0.173	0.027
Marginal impact with $N = 0$	0.737	0.091	0.590	1.341
Marginal impact with diagonal N	3.043	0.786	0.556	5.679

Note: N is the matrix of coefficients in the VAR for the first difference of goods prices and returns.

Table 2 also reveals the reason why the DPI is so volatile. Note that the marginal impacts in the third row sum across goods to 11.5. That is, a 1% increase in all prices raises the DPI by over 11%. For the VAR used in figure 5, an increase in 1% in all prices raises the DPI by 12%. Because all prices are very persistent, high prices today imply high prices in the future. Consumers therefore require a large increase in wealth in response to news on prices in order to be able to afford consumption in the future.

It is customary in building price indices to normalize the weights of each component so that they sum to one. Doing this for the DPI results in a price index that is almost perfectly correlated with the actual DPI, but has a standard deviation equal to 0.8%, less than a third of the CPI's volatility. If one is very worried about the volatility of the DPI, this normalized price index provides a more stable alternative which tracks very closely the movements in the DPI. However it no longer measures the percentage increase in wealth required by consumers.

6.3 Exploring some alternatives

Another price index that typically receives much attention is the deflator for personal consumption expenditures. The prices for its components are collected by the Bureau of Economic Analysis for the national income and product accounts (NIPA) independently of the work for the CPI. Figure 5 shows the log of the DPI using as expenditure shares and goods prices the series from the NIPA tables. In the figure is also the DPI using the CPI series, already displayed in figure 4.

Figure 5: The DPI using the NIPA (PCE deflator) or the BLS (CPI) data

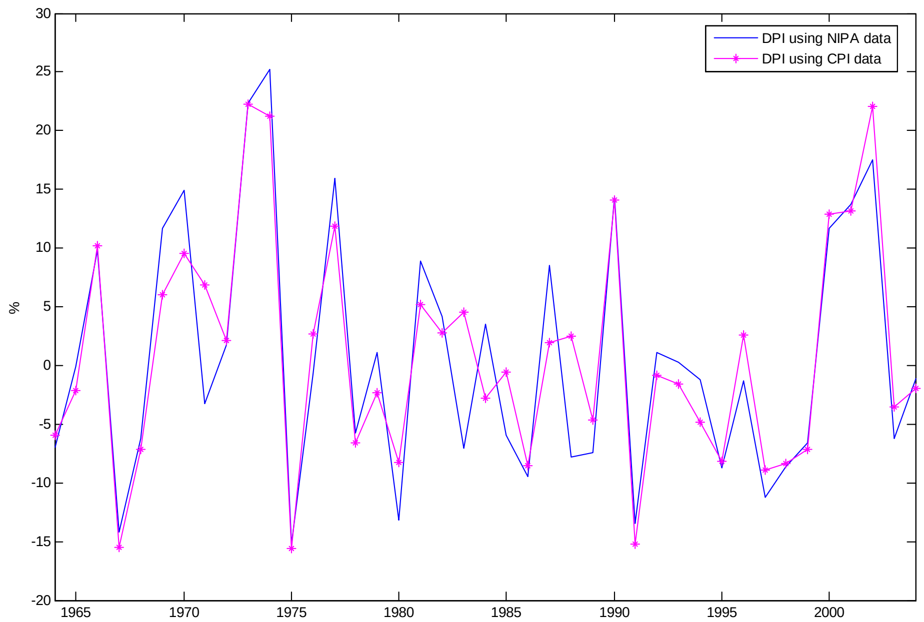

Clearly, the two series in figure 5 are very similar. Their correlation coefficient is 0.92. The conclusions for the DPI are not sensitive to using one or the other source of data.

The analysis of the U.S. DPI showed that, together with bond prices, the price of housing is the crucial driver of the price index. There have been many criticisms of the shelter series in the CPI. For one, there was an important change in the methodology used by Bureau of Labor and Statistics in 1982. Before this date, shelter prices referred to the price of recently purchased houses; afterwards, rents have been used to assess the cost of shelter. Moreover, while rents and house prices should be tightly linked, in the data they have moved in apparently disconnected ways (see Gillingham, 1983, or Verbrugge, 2005). Gordon and vanGoethem (2004) make a further criticism using the following insightful observation: according to the shelter price series, shelter should be 5.1 times more expensive in 1999 relative to 1925, yet the asking price for single-family homes in Washington, DC is 22.5 times higher. At least over long horizons, the shelter series seems to be biased downwards.

Figure 6 uses the CPI data but replaces the shelter price series by a series for house prices made available by the Office of Federal Housing Enterprise Oversight (OFHEO), based on weighted repeat-sales indices. The data is only available from 1975, so the DPI can only be computed from 1977 to 2004. In the figure is also the DPI using the CPI data for housing. The two are significantly correlated (with a coefficient of 0.63), but there is at least one significant difference. The price index using OFHEO house prices is more volatile, with a standard deviation that is about twice as large. House prices are significantly more volatile than rents, which is difficult to justify with standard economic models. This excess volatility puzzle has been noted before in the literature.

Figure 6: The DPI using the OFHEO's house price or the CPI's shelter price data

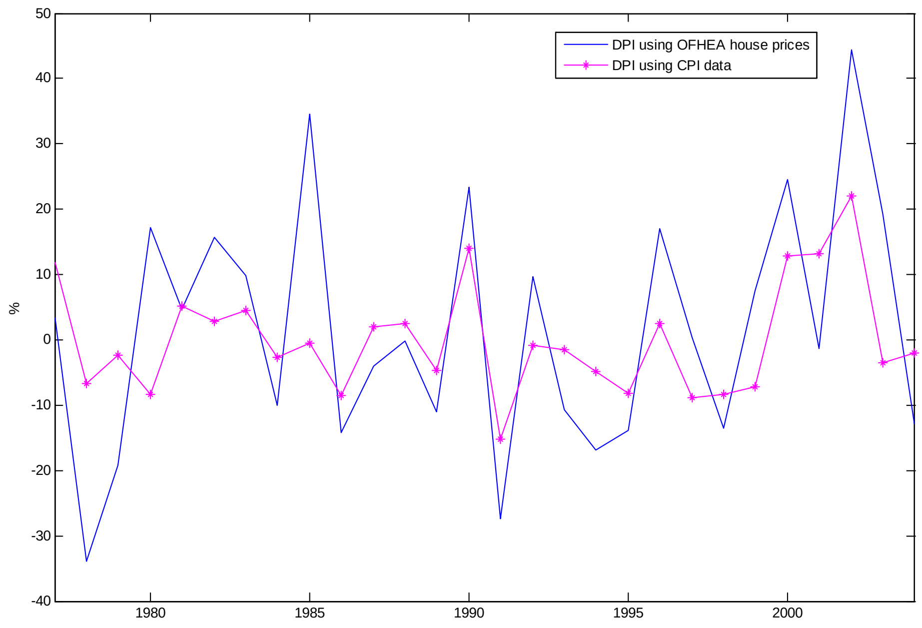

7 Conclusion

During the 1970s, macroeconomics went through a revolution. Static models that ignored uncertainty were for the most part abandoned and replaced by their dynamic counterparts with uncertainty incorporated using the assumption of rational expectations. In the study of consumption, following Hall's (1978) seminal study, a new vision of reality emerged. Consumption was no longer a predictable function of current income; instead future changes in consumption were unpredictable. Income was no longer the source of variations in consumption, but it was unexpected shocks to income that mattered. Consumers were now forward-looking and announcements on the future had an impact on the present.

This paper brought all of these ingredients into the theory of cost-of-living price indices. Perhaps not surprisingly, the cost of living in a dynamic stochastic framework leads also to a new vision of the world. The cost of living is now close to serially uncorrelated; it is news on prices that matter; and news that prices will increase in the future have an impact today. The resulting price index, the DPI, depends on a host of new interesting factors: the structure of financial markets, asset prices, and the durability of goods, among others. In practice in the United States, the DPI is driven mostly by changes in the prices of bonds and housing.

There is still a long way ahead in bringing modern models of consumption dynamics under uncertainty into price index theory. Since Hall's (1978) work there has been much progress in understanding consumption by considering the role of limitations on borrowing, habits, temptations, distorted expectations, and inattentiveness, among other realistic features of behavior. This work has revealed other surprising features of consumption choices that will surely have counterparts in future dynamic price indices.

At a more modest level, there are also many opportunities to improve on the estimates of the U.S. DPI in this paper. Research could consider more goods and prices as well as more involved utility functions that better fit the cross-sectional patterns of demand. In terms of the study of dynamics, the most important next task is probably to find better forecasting models of individual prices in order to isolate news on prices.

Before this work proceeds, it is important to be clear about what the DPI measures and what it does not. The DPI is the appropriate intertemporal measure of the cost of living. It measures by how much a single consumer would have to be compensated as she learns about changes in the prices she faces. For instance, the DPI is the appropriate measure to compare the cost of living in two different areas for a consumer with a fixed income. Another important use of the DPI would be as a compensation index to which income or pensions can be tied in order to keep the cost of living of their recipients unchanged.

Shifting from the CPI to the DPI as a compensation measure may explain a long-standing puzzle. Economists have wondered why is it that people do not index their nominal income to the CPI. This paper suggests that the reason may be that indexing to the CPI is not in the consumer's best interest, since it is almost uncorrelated with the appropriate dynamic price index. If this is the case, improving the measurements of the DPI and developing the institutions that would allow people to index to it, should be a priority for policy.

What the DPI does not measure is for instance overall welfare in an economy. Rather, it isolates solely the impact of prices on the well-being of consumers. In a time of low unemployment and high growth people may be very well off; yet, if this comes with higher than expected prices, the DPI will rise. Moreover, there is no justification anywhere in this paper for using the DPI as a target for monetary policy. A central bank may not have much control over many of the components of the DPI. Nor is it clear that this is desirable. If asset prices are prone to bubbles or irrational movements, then targeting the DPI might even be dangerous.[16]

[16] In current work with Mark Watson, we are exploring the construction of a price index that the Central Bank can control. Mankiw and Reis (2003) made a first attempt at computing the price index that the central bank should target.

Appendix

This appendix contains some calculations omitted from the text.

Consumer behavior with random asset prices

Starting with the i.i.d. case, define b_t to equal $Q_{E,t}B_{E,t}/(W_t - P_tC_t)$, the share of financial investments on equity. The consumer's problem is:

$$V(W_t, ...) = \max_{C_t, b_t} \left\{ \log(C_t) + \beta \mathbb{E}_t \left[V\left(\left(\frac{1-b_t}{Q_{0,t}} + \frac{b_t D_{t+1}}{Q_{E,t}} \right) (W_t - P_t C_t), ... \right) \right] \right\}.$$

The Euler equation for consumption is:

$$1 = \beta \mathbb{E}_t \left[\left(\frac{1-b_t^*}{Q_{0,t}} + \frac{b_t^* D_{t+1}}{Q_{E,t}} \right) \frac{P_t C_t^*}{P_{t+1} C_{t+1}^*} \right],$$

where asterisks denote optimal choices. The optimality condition for b_t is:

$$\mathbb{E}_t \left[\frac{W_{t+1}}{P_{t+1} C_{t+1}^*} \times \frac{\hat{Q}_{E,t} - \hat{Q}_{0,t} D_{t+1}}{(1-b_t^*) \hat{Q}_{E,t} + b_t^* \hat{Q}_{0,t} D_{t+1}} \right] = 0.$$

It is easy to see that the solution $P_t C_t^* = (1-\beta) W_t$ satisfies these two conditions, while the portfolio share solves:

$$\mathbb{E}_t \left[\frac{\hat{Q}_{E,t} - \hat{Q}_{0,t} D_{t+1}}{(1-b_t^*) \hat{Q}_{E,t} + b_t^* \hat{Q}_{0,t} D_{t+1}} \right] = 0,$$

and so depends only on the distribution of dividends and on the idiosyncratic shocks to asset prices.

The value function at date t equals the expected sum of discounted utility obtained by behaving optimally. Using the optimal consumption choices and the evolution of wealth implied by the budget constraint, evaluating expectations and summing over time, gives the value function, which up to a constant equals:

$$\frac{\ln(W_t)}{1-\beta} - \frac{\beta \ln(Q_t)}{1-\beta} - \ln(P_t) + \frac{\beta}{1-\beta} \mathbb{E}_t \left[\ln \left(\frac{1-b_t^*}{\hat{Q}_{0,t}} + \frac{b_t^* D_{t+1}}{\hat{Q}_{E,t}} \right) \right].$$

The definition of the DPI then implies that:

$$\ln(\pi_t) = (1-\beta) \ln(P_t) + \beta \ln(Q_t) - \beta \mathbb{E}_t \left[\ln \left(\frac{1-b_t^*}{\hat{Q}_{0,t}} + \frac{b_t^* D_{t+1}}{\hat{Q}_{E,t}} \right) \right].$$

The term involving idiosyncratic asset price shocks is more complicated. Still, the envelope theorem implies that $\partial \ln(V(W_t, ...)/\partial \ln(\hat{Q}_{E,t}) = -b_t^*$ while $\partial \ln(V(W_t, ...)/\partial \ln(\hat{Q}_{0,t}) = -(1-b_t^*)$. Using the definition of b_t together with the definition of the DPI, the result claimed in the text follows immediately.

The case when consumer prices follow an ARIMA(1,1,0) is very similar. The optimal solutions for consumption and asset holdings are precisely the same. The only difference is that the term in log consumer prices in the value function is now divided by $(1-\beta)(1-\beta\eta)$; the new DPI follows immediately.

Finally, lesson 3 focusses only on goods prices. In general, $\ln(\pi_t)$ depends on goods prices only through the term:

$$(1-\beta)\sum_{s=t}^{\infty}\beta^s\left(\mathbb{E}_t\left[\ln(P_s)\right]-\mathbb{E}_{t-1}\left[\ln(P_s)\right]\right).$$

If $\mathbb{E}_t[P_s] = \mathbb{E}_{t-1}[P_s]$, Jensen's inequality implies that $\mathbb{E}_t[\ln(P_s)] > \mathbb{E}_{t-1}[\ln(P_s)]$ for all s. All of the terms in brackets in the sum are positive so $\ln(\pi_t) > 0$ and the DPI is above one.

Consumer behavior with random durables prices

The optimality conditions for the problem in (12) imply the Euler equations for the consumption of both goods are:

$$\frac{Q_t}{\beta} = \mathbb{E}_t\left[\frac{P_t C_t^*}{P_{t+1}C_{t+1}^*}\right] \tag{15}$$

$$\frac{\alpha}{R_t S_t^*} = \frac{1-\alpha}{P_t C_t^*}\left(1-\beta(1-\delta)\mathbb{E}_t\left[\frac{R_{t+1}^S P_t C_t^*}{R_t P_{t+1}C_{t+1}^*}\right]\right). \tag{16}$$

Under the assumption of i.i.d. shocks, the definition of the DPI implies that $\partial \ln(\pi_t)/\partial \ln(R_t) = (R\partial V(.)/\partial R)/(W\partial V(.)/\partial W)$. A similar expression holds with respect to P_t, so the relative marginal weight of durables relative to non–durables, equals $(R\partial V(.)/\partial R)/(P\partial V(.)/\partial P)$. Using the envelope theorem on the dynamic program above to evaluate the derivatives of the value function shows that $M_t = R_t S_t^*/P_t C_t^*$.

At the steady state, $Q_t = \beta$, $R_t = \bar{R}$, and $R_{t+1}^S = \bar{R}^S$. The budget constraint and equation (15) imply that wealth and consumption are unchanged. Equation (16) gives the solution for $\bar{M}$.

Log-linearizing equations (15)-(16) around the steady state shows that:

$$m_t = \frac{\beta(1-\delta)\bar{R}^S/\bar{R}}{1-\beta(1-\delta)\bar{R}^S/\bar{R}} E_t[r^S_{t+1} - r_t + q_t].$$

Small letters denote log deviations of the respective capital letter from the steady state. Since R^S_{t+1}/R_t is the return on holding durables and $1/Q_t$ is the return on financial assets, the term in brackets is the expected difference between the return on durables and the return on financial assets.

Finally, the numerical results are obtained by guessing that $P_t C^*_t = (1-\alpha)(1-\beta)W_t$ and noting that for the parameter values, equations (15)-(16) verify this guess if there is an M_t that solves:

$$\begin{aligned} 1 &= \int \frac{\beta \exp[-50x^2]}{[1-0.05(1+M_t(1-0.9(1-\delta)e^x))]\, 0.1\sqrt{2\pi}} dx, \\ 1 &= \int \frac{\beta M_t(1-0.9(1-\delta)e^x)\exp[-50x^2]}{[1-0.05(1+M_t(1-0.9(1-\delta)e^x))]\, 0.1\sqrt{2\pi}} dx. \end{aligned}$$

Figure 1 plots the M_t found by numerically integrating these equations.

The U.S. DPI - details on Table 1

The Bureau of Labor Statistics (BLS) reports the "relative importance" of each item whenever it releases the CPI. These are calculated using the current price of the good and the quantity purchased at the most recent past date in which the BLS revised the consumption expenditure weights using the Consumption Expenditure Survey. Expenditure shares were revised at the end of 1963, 1977, 1982, 1986, 1997, 2001, and 2003. Therefore, the relative importance of item j in December of 1963 measures the actual expenditure share of this item $\alpha_{j,t}$ in December of 1963. Since these shares are not revised until December 1976, I keep $\alpha_{j,t}$ fixed for the period 1964-1977.

There have been frequent adjustments to how the BLS groups items into categories. The most significant one for this paper was the change in 1983 of the treatment of shelter. Whereas before 1983, housing was treated as a durable good purchase and its price was measured using prices of houses sold, since then housing has been treated as a service and its price has been measured using market rental values. In this paper, all shelter expenditures are always treated as durable to ensure consistency in the time-series. Further details are available from the author.

The U.S. DPI - solving the model

A few steps simplify the problem in (14). First, note that the intratemporal optimality conditions for allocating consumption between non-durables are:

$$\frac{\alpha_{i,t}}{P_{i,t}C_{i,t}} = \frac{\alpha_{1,t}}{P_{1,t}C_{1,t}},$$

for i from 1 to 4. Second, define the share of non-durables $\alpha^{ND} = \sum_{i=1}^{4} \alpha_i$. Third, perform a change of variables so that the consumer chooses $\tilde{C}_{1,t} = P_{1,t}C_{1,t}/W_t$ and $\tilde{S}_{i,t} = R_{i,t}S_{i,t}/W_t$, while $\tilde{W}_{t+1} = W_{t+1}/W_t$. Fourth, let $I_{0,t} = 1/Q_{0,t}$ and $I_{E,t+1} = D_{t+1}/Q_{E,t}$, the returns on bonds and equity respectively. Fifth and finally, let $b_t = Q_{E,t}B_{E,t}/(W_t - \sum_{i=1}^{4} P_{i,t}C_{i,t} - \sum_{i=5}^{6} \delta_i R_{i,t}S_{i,t})$.

This new notation permits writing the consumer's problem as a dynamic program:

$$\begin{aligned}
V(W_t, ...) &= \max_{\tilde{C}_{1,t},\tilde{S}_{5,t},\tilde{S}_{6,t},b_t} \left\{ \alpha^{ND} \ln(\tilde{C}_{1,t}) + \alpha_5 \ln(\tilde{S}_{5,t}) + \alpha_6 \ln(\tilde{S}_{6,t}) + \ln(W_t) + \beta \mathbb{E}_t \left[V(W_{t+1}, ...) \right] \right\}, \\
\tilde{W}_{t+1} &= \left[(1 - b_t) I_{0,t} + b_t I_{E,t+1} \right] \left(1 - \frac{\alpha^{ND} \tilde{C}_{1,t}}{\alpha_1} - \delta_5 \tilde{S}_{5,t} - \delta_6 \tilde{S}_{6,t} \right) \\
&\quad + \tilde{S}_{5,t}(1 - \delta_5)(R_{5,t+1}/R_{5,t} - I_{0,t}) + \tilde{S}_{6,t}(1 - \delta_6)(R_{6,t+1}/R_{6,t} - I_{0,t}), \\
W_{t+1} &= \tilde{W}_{t+1} W_t.
\end{aligned}$$

Note that the consumer treats the α's as fixed.

The Euler equations are:

$$\begin{aligned}
\frac{\alpha_1}{\tilde{C}^*_{1,t}} &= \frac{\beta}{1-\beta} \mathbb{E}_t \left[\frac{(1 - b_t) I_{0,t} + b_t I_{E,t+1}}{\tilde{W}_{t+1}} \right], \\
\frac{\alpha_5}{\tilde{S}^*_{5,t}} &= \frac{\delta_5 \alpha_1}{\tilde{C}^*_{1,t}} - \frac{\beta(1-\delta_5)}{1-\beta} \mathbb{E}_t \left[\frac{R_{5,t+1}/R_{5,t} - I_{0,t}}{\tilde{W}_{t+1}} \right], \\
\frac{\alpha_6}{\tilde{S}^*_{6,t}} &= \frac{\delta_6 \alpha_1}{\tilde{C}^*_{1,t}} - \frac{\beta(1-\delta_6)}{1-\beta} \mathbb{E}_t \left[\frac{R_{6,t+1}/R_{6,t} - I_{0,t}}{\tilde{W}_{t+1}} \right].
\end{aligned}$$

As for the portfolio choice, the optimality condition is

$$0 = \mathbb{E}_t \left[\frac{I_{E,t+1} - I_{0,t}}{\tilde{W}_{t+1}} \right].$$

Following Campbell and Viceira (2002), I employ a log-linearization that becomes exact as the time period shortens. It implies that the portfolio share equals the Sharpe ratio:

$$b_t = \frac{\mathbb{E}\left[\ln(I_{E,t+1})\right] - \ln(I_{0,t}) + 0.5Var\left[\ln(I_{E,t+1})\right]}{Var\left[\ln(I_{E,t+1})\right]}.$$

Having characterized the solution, I now describe the steady state in which all returns are close to β^{-1}. At this steady state, $\tilde{C}^*_{1,t} = \alpha_1(1-\beta)$, $\delta_5\tilde{S}^*_{5,t} = \alpha_5(1-\beta)$, $\delta_6\tilde{S}^*_{6,t} = \alpha_6(1-\beta)$, and $\tilde{W}_{t+1} = 1$.

The next step is to log-linearize the model around this steady state. Let ($c_{1,t}$, $s_{5,t}$, $s_{6,t}$, $r_{5,t}$, $r_{6,t}$, $i_{E,t}$, $i_{0,t}$, w_t) denote log-deviations from the steady state of the respective capital letter variables with a tilde. Then group variables in the 3x1 and 8x1 vectors:

$$\begin{aligned}
y_t &= (c_{1,t}, s_{5,t}, s_{6,t})' \\
z_t &= (p_{1,t} - p_{1,t-1}, p_{2,t} - p_{2,t-1}, p_{3,t} - p_{3,t-1}, p_{4,t} - p_{4,t-1}, r_{5,t} - r_{5,t-1}, r_{6,t} - r_{6,t-1}, i_{E,t}, i_{0,t})'.
\end{aligned}$$

Recall that the dynamics of z_t are modelled as a VAR(1): $z_t = Nz_{t-1} + \varepsilon_t$, where N is an 8x8 matrix of estimated coefficients and ε_t is the residual.

Log-linearizing the Euler equations gives the system:

$$Ay_t + (B\iota_8 + C[\iota_5;\ \iota_6]N)\, z_t = 0,$$

where ι_i denotes the i$^{\text{th}}$ row of an 8x8 identity matrix, and:

$$\begin{aligned}
A &= \begin{pmatrix} 1 + \alpha^{ND}(\beta^{-1} - 1) & \alpha_5(\beta^{-1} - 1) & \alpha_6(\beta^{-1} - 1) \\ -\delta_5 & \delta_5 & 0 \\ -\delta_6 & 0 & \delta_6 \end{pmatrix}_{3\text{x}3} \\
B &= \begin{pmatrix} (\beta^{-1} - 1)\left[\alpha_5(\delta_5^{-1} - 1) + \alpha_6(\delta_6^{-1} - 1)\right] \\ \delta_5 - 1 \\ \delta_6 - 1 \end{pmatrix}_{3\text{x}1} \\
C &= \begin{pmatrix} -(\beta^{-1} - 1)\alpha_5(\delta_5^{-1} - 1) & -(\beta^{-1} - 1)\alpha_6(\delta_6^{-1} - 1) \\ 1 - \delta_5 & 0 \\ 0 & 1 - \delta_6 \end{pmatrix}_{3\text{x}2}.
\end{aligned}$$

Log-linearizing the budget constraint gives the equation:

$$\begin{aligned} w_t &= (1-\beta^{-1})e_1 y_{t-1} + \bar{b}\iota_7 z_t + (1-\bar{b})\iota_8 z_{t-1} + D\varepsilon_t \\ D &= \begin{pmatrix} 0 & 0 & 0 & 0 & (\beta^{-1}-1)\alpha_5(\delta_5^{-1}-1) & (\beta^{-1}-1)\alpha_6(\delta_6^{-1}-1) & 0 & 0 \end{pmatrix}_{1\text{x}8}, \end{aligned}$$

where e_1 is the first row in a 3x3 identity matrix.

These two equations combined with the VAR for prices then imply the solution:

$$y_t = Hz_t \tag{17}$$

$$w_t = Fz_{t-1} + G\varepsilon_t \tag{18}$$

$$z_t = Nz_{t-1} + \varepsilon_t, \tag{19}$$

where $H = -A^{-1}\left(B\iota_8 + C[\iota_5;\ \iota_6]N\right)$, $F = (1-\beta^{-1})e_1 H + \bar{b}\iota_7 N + (1-\bar{b})\iota_8$, and $G = \bar{b}\iota_7 + D$. Given a set of parameter values and the estimated N from the data all of these matrices are known.

I now solve for the value function. By definition:

$$V(W_t, ...) = \mathbb{E}_t\left[\sum_{j=0}^{\infty}\beta^j\left(\sum_{i=1}^{4}\alpha_i\ln(C_{i,t+j}) + \sum_{i=5}^{6}\alpha_i\ln(S_{i,t+j})\right)\right].$$

With the changes of variables that I performed, up to terms that do not involve prices or wealth, the value function equals:

$$\begin{aligned} &\mathbb{E}_t\left[\sum_{j=0}^{\infty}\beta^j\left(\alpha_{t+j}^{ND}c_{1,t+j} + \alpha_5 s_{5,t+j} + \alpha_6 s_{6,t+j}\right)\right] + \frac{\ln(W_t)}{1-\beta} \\ &+\mathbb{E}_t\left[\sum_{j=1}^{\infty}\beta^j\sum_{i=1}^{j}w_{t+i}\right] - \mathbb{E}_t\left[\sum_{j=0}^{\infty}\beta^j\left(\sum_{i=1}^{4}\alpha_i p_{i,t+j}) + \sum_{i=5}^{6}\alpha_i r_{i,t+j})\right)\right] \end{aligned}$$

Consider each of these terms in turn. Let $\bar{\alpha} = (\alpha^{ND}, \alpha_5, \alpha_6)$, a 1x3 vector. Then, using the solution in (17)-(19), the first term in the value function becomes:

$$\left(\sum_{j=0}^{\infty}\beta^j\bar{\alpha}HN^j\right)z_t.$$

The third term likewise becomes:

$$\left(\sum_{j=1}^{\infty}\beta^j\sum_{i=1}^{j}FN^{i-1}\right)z_t.$$

Finally, consider the fourth term. Defining the 1x8 vector $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6, 0, 0)$ the fourth term equals:

$$\left(\sum_{j=1}^{\infty}\beta^j\alpha\sum_{i=1}^{j}N^i + \frac{\alpha}{1-\beta}\right)z_t.$$

Combining all of these terms gives the solution for the value function:

$$\frac{\ln(W_t)}{1-\beta} + Jz_t,$$

where $J = \bar{\alpha}H - \alpha/(1-\beta) + \sum_{j=1}^{\infty}\beta^j\left(\bar{\alpha}HN^j + (F-\alpha N)\sum_{i=1}^{j}N^{i-1}\right)$.

The definition of the DPI then implies that:

$$\ln(\pi_t) = -\left(G + (1-\beta)J\right)\varepsilon_t.$$

Given a time-series for news on prices, this formula generates a time-series for the DPI.

References

[1] Alchian, Armen and Benjamin Klein (1973) "On a Correct Measure of Inflation," *Journal of Money, Credit and Banking*, vol. 5 (1), pp. 173-191.

[2] Bajari, Patrick, C. Lanier Benkard, and John Krainer (2004) "House Prices and Consumer Welfare," Duke University.

[3] Boskin, Michael, Ellen Dulberger, Robert Gordon, Zvi Griliches, and Dale Jorgenson (1997) "The CPI Commission: Findings and Recommendations," *American Economic Review Papers and Proceedings*, vol. 87, pp. 78-83.

[4] Bryan, Michael, Stephen Cecchetti and Roisin O'Sullivan (2001) "Asset Prices and the Measurement of Inflation," *De Economist*, vol. 149, 405-431.

[5] Campbell, John and Luis Viceira (2002) *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford: Oxford University Press.

[6] Cecchetti, Stephen, H. Genburg, J. Lipsky, and S. Whadwani (2000) *Asset Prices and Central Bank Policy*, Geneva Reports on the World Economy, no. 2, International Center for Monetary and Banking Studies and Centre for Economic Policy Research.

[7] Diewert, W. Erwin (2001) "The Consumer Price Index and Index Number Theory," University of British Columbia Department of Economics discussion paper no. 01-02.

[8] Diewert, W. Erwin and Claude Montmarquette (1983) *Price Level Measurement*, Ottawa: Statistics Canada.

[9] Fisher, Irving (1906) *The Nature of Capital and Income*, New York.

[10] Fisher, Franklin and Karl Shell (1972) *The Economic Theory of Price Indices: Two Essays on the Effects of Taste, Quality, and Technological Change*, New York: Academic Press.

[11] Gilchrist, Simon and John Leahy (2002) "Monetary Policy and Asset Prices," *Journal of Monetary Economics*, vol. 49, pp. 75-97.

[12] Gillingham, Robert (1983) "Measuring the Cost of Shelter for Homeowners: Theoretical and Empirical Considerations," *Review of Economics and Statistics*, vol. 65 (2), pp. 254-265.

[13] Goodhart, Charles (2001) "What Weight Should Be Given to Asset Prices in the Measurement of Inflation?" *Economic Journal*, vol. 111, pp. F335-F356.

[14] Goodhart, Charles and Boris Hofmann (2000) "Do Asset Prices Help to Predict Consumer Price Inflation," *The Manchester School Supplement*, vol. 68, pp. 122-140.

[15] Gordon, Robert J. and Todd vanGoethem (2004) "A Century of Downward Bias in the Most Important Component of the CPI: The Case of Rental Shelter, 1914-2003," Northwestern University.

[16] International Labor Office (2004) *Consumer Price Index Manual: Theory and Practice*, International Labor Office: Geneva. (Available at http://www.ilo.org/public/english/bureau/stat/guides/cpi/)

[17] Konus, A. (1924) "The Problem of the True Index of the Cost of Living," translated in *Econometrica* (1939), vol. 7, pp. 10-29.

[18] Mankiw, N. Gregory and Ricardo Reis (2003) "What Measure of Inflation Should a Central Bank Target?" *Journal of the European Economic Association*, vol. 1 (5), pp. 1058-1086.

[19] National Research Council (2002) *At What Price? Conceptualizing and Measuring Cost-of-Living and Price Indexes*, Panel on Conceptual, Measurement, and Other Statistical Issues in Developing Cost-of-Living Indexes, Charles L. Schultze and Christopher Mackie, editors, Committee on National Statistics, Division of Behavioral and Social Sciences and Education, Washington, DC: National Academy Press.

[20] Pollack, Robert (1989) *The Theory of the Cost-of-Living Index*, Oxford: Oxford University Press.

[21] Samuelson, Paul and Subramanian Swamy (1974) "Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis," *American Economic Review*, vol. 64, pp. 566-593.

[22] Shapiro, Matthew D. and David W. Wilcox (1996) "Mismeasurement in the Consumer Price Index: An Evaluation," *NBER Macroeconomics Annual 1996*, Cambridge MA: MIT Press.

[23] Shibuya, Hiroshi (1992) "Dynamic Equilibrium Price Index: Asset Prices and Inflation," *Monetary and Economic Studies*, vol. 10, pp. 95-109.

[24] Shiller, Robert (1993) *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*, Oxford: Oxford University Press.

[25] Stock, James and Mark Watson (2003) "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, vol. 41, pp. 788-829.

[26] Triplett, Jack (1983) "Escalation Measures: What is the Answer? What is the Question?" in *Price Level Measurement*, edited by W. Erwin Diewert and Claude Montmarquette, Ottawa: Statistics Canada.

[27] Verbrugge, Randal (2005) "The Puzzling Divergence of Rents and User Costs, 1980-2002," Division of Price and Index Number Research, Bureau of Labor Statistics.