What Can Asset Prices Tell Us About Historical Business Cycles?

Benjamin Chabot
Remy@umich.edu

University of Michigan and NBER

Preliminary and Incomplete – Please do not cite – comments encouraged

Abstract

Have business cycles become less volatile? Are modern recessions less frequent and more benign then their 19th century counterparts? This paper makes use of a large cross section of asset prices and the restrictions investor's first order conditions to estimate the aggregate consumption implied by asset prices. Unlike GDP, unemployment or industrial production, historical asset prices are measured with great precision and have been contemporaneously collected for hundreds of years. The methodology outlined in this paper can therefore be employed to estimate business cycles for time periods and locations that were previously obscured by data limitations.

Introduction

Have business cycles become more benign? In 1959, Arthur Burns, impressed with the previous decade of stability, predicted that the modern business cycle "is unlikely to be as disturbing or troublesome to our children as it was to our fathers"[1]. Unlike Irving Fisher's prediction a generation earlier, history has been kind to Burns' forecast. The uninterrupted expansion of the 1960s proved so stable that the Department of Commerce decided to change the name of its *Business Cycle Digest* to the *Business Conditions Digest*[2]. When recessions did return in the 1970s and 80s they were relatively mild in comparison to the era of Burns' father when, in the words of Christina Romer, "there is simply no denying that all hell broke lose in the American economy"[3].

A word of caution is in order. Although most measures of business cycle severity have decline in the latter half of the 20th century, the declines are not monotonic. It would probably have surprised Burns (as well as any modern economist unfamiliar with pre-B.E.A. data) to learn that despite the moderation of the business cycle since the 1930s the variance of annual GDP growth in the 1970s and 80s was statistically indistinguishable from the 1870s and 80s[4].

Has the business cycle really become less volatile? Are post-World War II recessions shorter in duration or less frequent then recessions 100 or 150 years ago? To answer these questions we require a consistent measure of the frequency and severity of business cycles. Unfortunately, the "*statistics that economists use to measure the severity of business cycles, such as data on the unemployment rate, real gross national product, and industrial production, have been kept carefully and consistently only since World War II. Therefore, the conclusion that government policy has smoothed business cycles is based on a*

---

[1] 1959 Presidential Address to the American Economic Association
[2] J. Bradford DeLong Dec,12,1998 "Introduction to symposium on Business Cycles". http://www.j-bradford-delong.net/Comments/bcsymposium_intro.html
[3] Romer (1999) *Journal of Economic Perspectives* p.26
[4] Burns, of course, was quite familiar with pre-B.E.A. business cycle data. He would have likely been surprised at the magnitude of 1970s GDP fluctuations.
The GDP data are available at http://www.eh.net/hmit/gdp/. The Bartlett Test for homogeneity of variances has a P-value of .17

comparison of fragmentary prewar evidence with sophisticated postwar statistics"[5]. In a series of influential papers, Christina Romer (1986a-c) persuasively argued that the apparent dampening of the post-war business cycle was a figment of the data[6].

What would we find if we could examine business cycle data that was consistently recorded? This paper seeks to answer this question by making use of a recently collected panel of cross sectional stock prices. Unlike GDP, industrial production or unemployment, stock prices are measured with great precision and have been contemporaneously recorded for hundreds of years. Economic theory links stock returns to aggregate consumption via a simple law of one price relationship. This paper exploits the theoretical relationship between asset returns and consumption growth to estimate the aggregate consumption implied by asset returns over the past 175-years.

The Link between Asset Prices and Consumption

Asset returns are linked to consumption via investors' first order conditions. Consider an investor who may buy or sell an asset at time t. A necessary condition for expected utility maximization is that the marginal cost of the asset equal its marginal benefit

$$P_t \frac{\partial U}{\partial c_t} = E_t[\frac{\partial U}{\partial c_{t+1}}(P_{t+1} + D_{t+1})] \qquad (1)$$

Here P_t denotes the price of this asset at time t, and D_{t+1} is this asset's

[5] Christina Romer "Business Cycles" in The Concise Encyclopedia of Economics. http://www.econlib.org/library/Enc/BusinessCycles.html

[6] See Balke and Gordon (1986) and Weir (1986) for dissenting views

stochastic dividend (if any) at time $t+1$. The left hand side of equation (1) is the marginal cost (measured in utility) of purchasing the asset at time t. The right hand side is the expected gain in future utility resulting from the purchase. Since price and consumption are known at time t we can rewrite (1) in terms of asset returns and intertemporal marginal rates of substitution.

$$1 = E[m_{t+1}R_{t+1}] \qquad (2)$$

$m_{t+1} = \frac{\frac{\partial U}{\partial c_{t+1}}}{\frac{\partial U}{\partial c_t}}$ is the intertemporal marginal rate of substitution (IMRS) and R_{t+1} is the asset's gross asset return at time t+1. Equation (2) is the fundamental equation of financial economics. If all investors are free to buy or sell at the same price a m_{t+1} that satisfies (2) is guaranteed to exist by the law of one price.

Equation (2) also links expected asset returns to consumption. To see this, expand $E_t[m_{t+1}R_{t+1}]$ to express price in terms of expected covariance with investors' IMRS.

$$1 = E_t[m_{t+1}]E_t[R_{t+1}] + \mathrm{cov}_t(m_{t+1}, R_{t+1}) \qquad (3)$$

Imagine a risk-free real asset granting its holder 1 real time t dollar of consumption at time t+1. If such an asset existed, its price would be equal the conditional expectation $E_t[m_{t+1}]$. Thus, $E_t[m_{t+1}]$ is equal to the reciprocal of the real gross risk-free rate.

$$E_t[m_{t+1}] = \frac{1}{R_{f,t+1}}$$

Substituting the definition of the risk free rate into (3) yields

$$1 = \frac{E_t[R_{t+1}]}{R_{f,t+1}} + risk\ adjustment$$

$$risk\ adjustment = \mathrm{cov}_t[m_{t+1}, R_{t+1}]$$

In other words, the price of a risky asset is equal to the sum of its expected payout discounted by the risk free rate and a risk adjustment. The risk adjustment is determined by the asset's covariance with the investor's consumption. Assuming investors are risk averse, the stocks with future payoffs are positive correlated with consumption will have low prices and high expected returns.

There is a large literature employing modern data to examine the links between consumption and the cross section of asset returns. Mankiw and Shapiro (1996) and Breeden, Gibbons, and Litzenberger (1989) find modest correlations between contemporaneous consumption growth and expected returns. Lettau and Ludvigson (2001a 2001b), Parker and Julliard (2004), and Bansal, Dittmar, and Lundblad (2004) document stronger links when long run consumption correlations or contemporaneous variables that correlate with long run consumption growth are used in place of contemporaneous consumption growth. These papers are motivated by the desire to explain cross-sectional differences in expected return via a link to consumption. This paper adds to the finance literature by documenting an historical link between asset price implied consumption and the business cycle.

This paper also draws on a long tradition in economic history. Historians are often interested in time periods and locations where variables of interest are either poorly measured or entirely unobservable. In these cases it is common to estimate variables of interest via theoretical links to data that is available or measured with more precision. Fogel and Engerman's (1974) use of anthropometric data and Friedman and Schwartz's (1963) used of money and prices are two prominent examples of this technique.

Estimating Consumption from a Cross Section of Asset Returns

Equation (2) links expected return to investors' IMRS in the cross-section of asset returns. This suggests a natural measure of business cycle volatility. Given a sample of asset returns we can estimate the IMRS from the moment conditions in (2) and use it to compute the consumption stream implied by a given utility specification.

Given a sample of N assets for T time periods, the gross return on asset n at time t can always be expressed as a projection on a constant and mean zero factors

$$\begin{aligned} R_{nt} &= E(R_{nt}) + b_{n1} f_{1t} + \ldots + b_{nk} f_{kt} + \varepsilon_{nt} \\ E[f_k] &= E[f_k \varepsilon_{nt}] = 0 \quad \forall n,k,t \end{aligned} \tag{4}$$

Where b_{nk} is assets n's sensitivity to the mean-zero common factor f_k, and ε_{nt} is assets n's idiosyncratic risk. Our estimation strategy is to choose a sufficiently large number of factors such that the resulting error terms ε_{nt} are diversifiable. That is, we wish to choose factors until the remaining variation in asset returns is sufficiently independent across securities that $\underset{N\to\infty}{Lim} E[\frac{1}{N}\sum_{n=1}^{N}\varepsilon_{nt}] = 0 \quad \forall t$. The moment conditions in (2) together with the factor structure in (4) imply an IMRS that is a linear function of the factors and a constant $m_{t+1} = c_0 + \sum_{k=1}^{K} c_k f_{kt+1}$.

Choosing the Factors

The common factors can be extracted from the T x T covariance matrix of asset returns via a factor analytic technique. Let $\mathbf{r}$ denote the T by N matrix of returns in excess of the first factor.[7] Unbalanced panels are common in financial asset data because some assets will not be available for some time periods.

[7] I $r_{nt} = R_{nt} - ($ return on the first factor at time t). Conner and Korajczyk (1988) take the U.S. treasury bill as a proxy for the risk free rate and use returns in excess of the risk free rate to extract factors. This technique can be applied to returns in excess of any well-diversified portfolio as long as the process of taking excess returns does not change the number of common factors. No nominally risk free asset existed for much of the 19th century sample. I therefore take returns in excess of the value-weighted index of all asset returns and use the value-weighted index as a factor to avoid the possibility of an omitted factor.

Define $n_{t\tau}$ to be the number of assets with returns in both periods t and τ. Replace the missing assets with zeros to form the matrix Ω and let $\Omega^*_{t,\tau} = (rr'/n_{t\tau})$. Where the subscript t τ denotes the t, τ element of the matrix Ω^*. Connor and Korajczyk (1988) show that as the number of assets grows the first k eigenvectors of Ω^* form valid proxies for the k unobservable factors.

This factor extraction technique assumes the factor return structure in (4) and selects factors that contain the most information about asset returns. Unfortunately, this technique is sensitive to the length of the data window used in estimation. Over long time periods the factor loadings of individual stocks may change. I therefore constrain the estimation window to 10 years and estimate consumption streams via rolling 10-year windows of data and then average the consumption estimates together using a triangular kernel.

Theory offers no guidance as to the number of common factors. Modern cross sectional stock data suggests that as few as 3 factors are sufficient to capture the common cross sectional variation in asset returns. Statistical methods of estimating the number of factors from asset returns notoriously overstate the true number of factors but seldom select more then 10 from modern data. The cost of omitting a factor is high and the while the cost of including a spurious factor is low. I therefore estimate the model with 10 factors.

Estimating m_{t+1}

With assets and factors in hand we can estimate m_{t+1} by altering the moment condition in (2) to account for an unbalanced panel of asset returns.[8]

$$E(R_{nt+1}m_{t+1}) = D_{nt+1} \qquad (5)$$

[8]This assumes that the missing asset returns are missing at random.

where R_{nt+1} is the return on asset n at time $t+1$, with zeros in place of missing observations, D_{nt+1} is a dummy variable equal to zero if the return on asset n is missing at time $t+1$ and m_{t+1} is defined as

$$m_{t+1} = c_0 + \sum_{k=1}^{K} c_k f_{kt+1} \qquad (6)$$

Define the following error model

$$\begin{aligned} u(\theta)_{nt+1} &= m(\theta, \mathbf{F}_{t+1}) R_{nt+1} - D_{nt+1} \\ \theta &= [c_0 \ldots c_K] \end{aligned} \qquad (7)$$

Our goal is to pick the free parameters $\theta = [c_0 \ldots c_K]$ to best price the observable asset returns. Let T_n be the number of observations for asset n. Then $g(\theta)_n$ is the average pricing error of n-th asset:

$$g(\theta)_n = \frac{1}{T_n} \sum_{t=1}^{T} u(\theta)_{nt} \qquad (8)$$

To estimate θ, I form the vector of average pricing errors $\mathbf{g}(\theta) = [g(\theta)_1 \ldots g(\theta)_n \ldots g(\theta)_N]'$. and choose $\theta = [c_0 \ldots c_K]$ to minimize $J = \mathbf{g}(\theta)' \mathbf{W} \mathbf{g}(\theta)$ for a positive definite weighting matrix $\mathbf{W}$.

I use a diagonal weighting matrix with the inverse of the variances of $E[R_n]$ on the diagonal and zeros elsewhere. Equation (3) tells us that assets with high expected returns are highly correlated with consumption. We don't observe expected return, however, only an estimate. The weighting matrix places more weight on the moments of assets whose expected return is measured with confidence.

The Consumption Stream Implied by Historical Asset Returns and CRRA Utility

To this point we have relied on no more than the first order condition to estimate a sample IMRS from asset prices. Given a time series of IMRS realizations what consumption stream does this imply? To answer this question we need to specify a functional form for utility. I estimate consumption streams implied by constant relative risk aversion preferences. Time-Separable constant relative risk aversion (CRRA) preferences are perhaps the most common specification in the macroeconomic literature. CRRA preferences are described by the following utility function

$$U = E_0[\sum_{t=1}^{\infty} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma}] \qquad (9)$$

Where $\beta < 1$ is the time discount factor and $\gamma > 0$ is the coefficient of risk aversion. CRRA utility implies the following IMRS

$$m_{t+1} = \beta(\frac{c_{t+1}}{c_t})^{-\gamma} \qquad (10)$$

The CRRA IMRS is not sufficient to identify consumption in levels. However, we can identify the growth rate of consumption. Given a sample IMRS time series we can compute the implied consumption path by normalizing consumption in the initial period to 1 and using (10) to compute the remaining consumption sequence.

Filtering the IMRS Estimates and Computing Implied Consumption

Before transforming the IMRS estimates into implied consumption streams I filter the data to remove the high frequency variation. Specifically, I use the following algorithm to estimate a series of overlapping filtered stock market implied consumption streams:

1. Extract common factors from the first 10 years of stock market returns.
2. Estimate a monthly IMRS via the asset moment conditions.
3. Remove the non-business cycle variation in the sample IMRS estimate with the HP-filter (w=14400)
4. Use the filtered IMRS estimates and the CRRA utility specification to compute an implied consumption stream
5. Move ahead one year (i.e. choose years 2-11) and repeat steps 1-4.
6. Repeat steps 1-5 until the end of the data is reached

The result is a series of overlapping 10-year estimates of implied consumption. The HP-filter smoothes the IMRS estimates without removing the business cycle variation. This results in better estimates but the filtering comes with a cost. The HP-filter is notoriously inaccurate near the end points of the data series. I account for this by averaging the overlapping estimates together with a diagonal kernel that places the most weight on estimates in the middle of a 10-year sample and the least weight on the end points.

Evaluating the Goodness of Fit

To evaluate the accuracy of the estimation method I compare the stock market implied consumption to non-durable personal consumption expenditures published by the B.E.A. The B.E.A. has only collected monthly non-durable consumption data since 1959. We should be reasonably confident in these modern consumption estimates. The difference between the B.E.A. estimates and the stock market implied consumption should therefore serve as a good test of the stock market based estimation technique.

Figure I graphs the stock market based consumption estimates and the B.E.A non-durable consumption expenditures.

Figure I

(c_{t+1}/c_t) Implied by Stock Returns and Non-Durable Personal Consumption Expenditures: 1958-2000

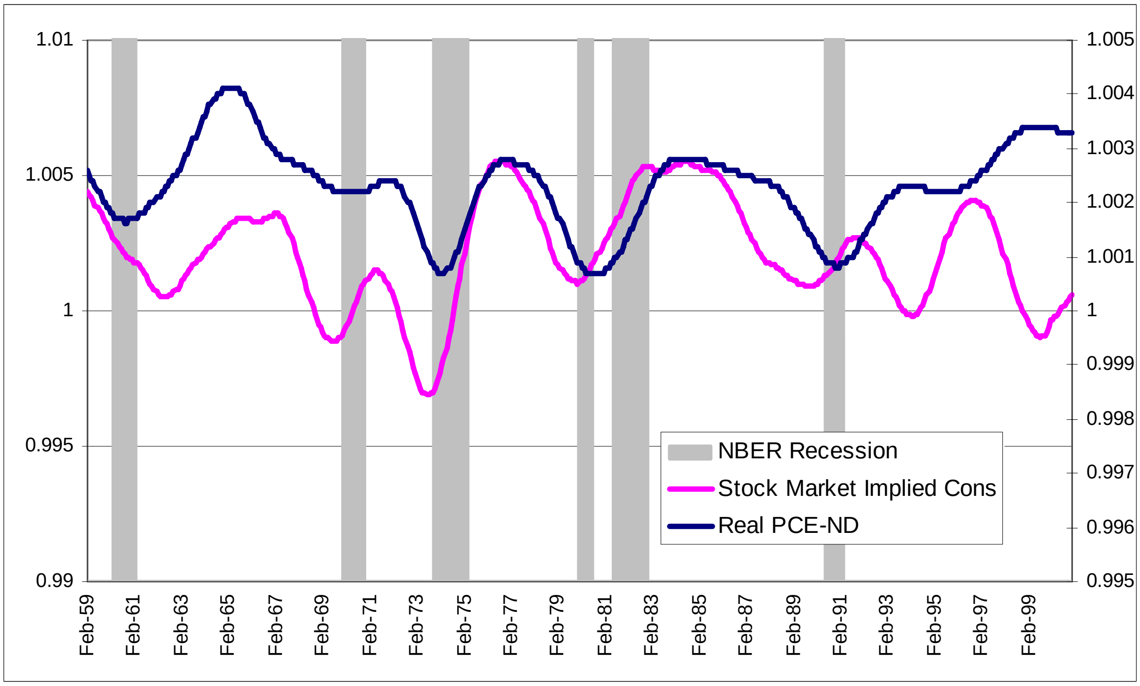

Left Scale = Gross Stock Market Implied Consumption growth rate.
Right Scale = Gross B.E.A. Non-Durable PCE growth rate.
NBER Recession = NBER business cycle contraction

The consumption stream implied by asset prices does a fairly good job of matching both the timing of the BEA consumption series and the NBER business cycle dates. Most peaks and troughs occur during NBER expansions and recessions respectively. The variation of the stock market implied consumption series is considerably larger then the BEA series, however (note the differing scales in Figure I). This excess volatility is a manifestation of the well known excess volatility in asset returns documented in Shiller (1989). As a whole, the comparison between the presumably accurate B.E.A. data and the stock market implied consumption suggests that the stock market provides a good measure of the timing of business cycle fluctuations but vastly overstates the magnitude of these variations.

The Consumption Implied by Asset Prices 1825-1958

The remainder of the paper is organized by time period and data set. I report estimated consumption streams for three distinct time periods defined by the data sets available in each period. For the 1926-1958 period, I use the University of Chicago's Center for Research in Security Prices (CRSP) data set. For 1866-1925 I use my own self-collected data and for 1825-1865 I use Goetzmann, Ibbotson and Peng's NYSE data set.

Consumption in the early CRSP Era: 1926-1958

The Center for Research in Security Prices (CRSP) data set contains the monthly holding period returns of all securities listed on the NYSE, NASDAQ or AMEX between 1926 and 1958. I employ the CRSP data to estimate stock market implied consumption via the methods outlined above. Figure II contains a graph of the stock market implied consumption estimated with the 1926-1958 data. No other monthly consumption estimates exist for this time period. For comparison I include Industrial Production filtered at the same frequencies as implied consumption.

Figure II

(c_{t+1}/c_t) Implied by Stock Returns and Industrial Production: 1926-1958

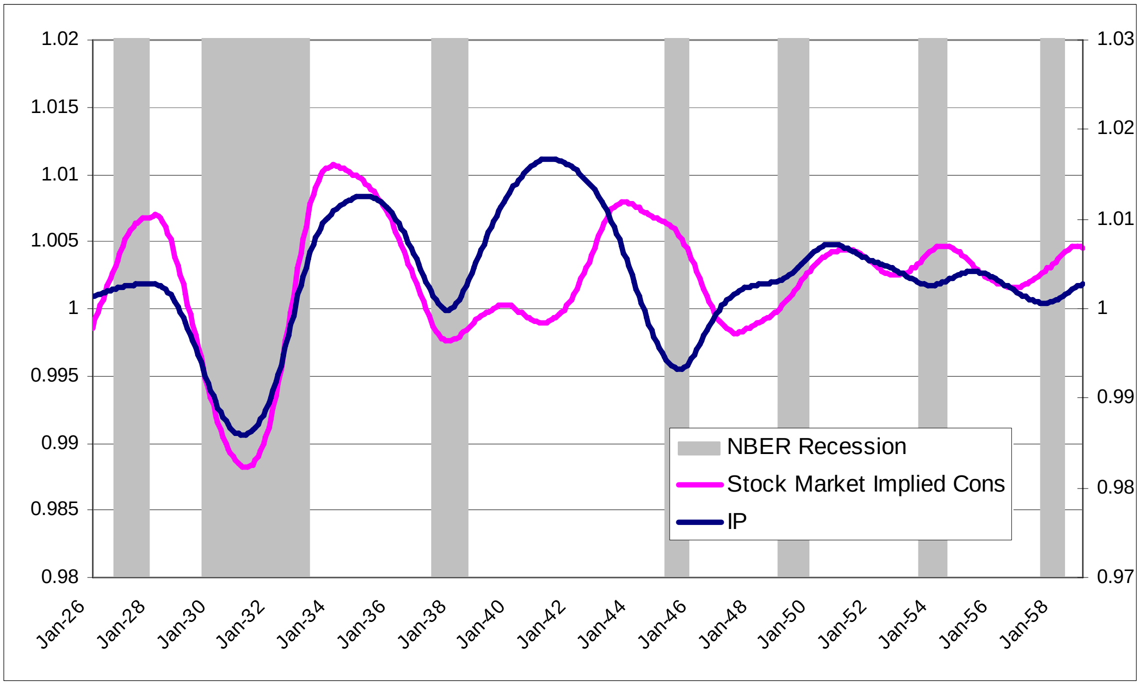

Left Scale = Gross Stock Market Implied Consumption growth rate.
Right Scale = Industrial Production
NBER Recession = NBER business cycle contraction

During the sample period, the stock market implied consumption series does a poor job of matching the NBER business cycle dates. The implied consumption series declines during the contractions of the great depression but is near its peak during the contractions of 1927 and 1946 and shows no noticeable decline during the contractions of 1953-54 or 1957.

The lack of correlation between the NBER business cycle dates and the implied consumption series is as much a statement about NBER business cycle dating conventions as a reflection on the accuracy of the stock market implied consumption series.

Although no monthly consumption data is available before 1958 we do have annual consumption estimates. Chapter 26 in Robert Shiller's (1989) *Market Volatility* draws on a number of sources to compute annual consumption estimates dating back to 1890[9].

[9] Shiller's (1989) data is available online at http://www.econ.yale.edu/~shiller/data/chapt26.xls

Table I compares the 1926-1958 annual stock market based consumption estimates to the annual consumption estimates reported in Shiller (1989).

Table I

	Stock Market Implied Consumption Annual Growth Rate	Shiller (1989) Consumption series Annual Growth Rate
All years	2.66%	1.74%
Years with NBER Contraction	0.90%	-.02%
Years without NBER Contraction	4.32%	3.56%
<u>Contraction Years</u>		
1926-27	4.6%	3.75%
1929-33	-3.81%	-3.72%
1937-38	-1.73%	0.31%
1945	6.62%	5.02%
1948-49	0.73%	0.68%
1953-54	4.87%	1.63%
1957-58	4.13%	-0.11%

Both series grow more in years without contractions. The Shiller consumption series grows at an above average rate during the recessions of 1926-27 and 1945. The stock market implied consumption estimates also grow at an above average rate during these recessions. Likewise, both series grow at a below average rate during the recessions of 1929-33, 1937-38 and 1948-49. The two measures do diverge. The 1937-38 contraction is a mild downturn in the Shiller series but a severe recession in the stock market implied consumption series. The stock market series also implies strong growth during the 1953-54 and 1957-58 recession while the Shiller consumption series is below average in the first recession and barely declines in the second.

Implied Consumption 1866-1925

I use self-collected data to estimate implied consumption between 1866 and 1925. The data consists of all stocks that traded on the New York Stock Exchange between 1866 and 1925. The data was sampled every 4-th Friday between January 1866 and December 1925.

Again I employ a rolling estimator with a 10-year window. Starting with the 1866-1875 period I estimate the consumption stream implied by asset prices. I then repeat the estimation for the 1867-1876 period and so on until 1916-1925. With overlapping estimates in hand I form one estimate by taking a weighted average where the weights are determined by a triangular kernel. Figure III contains a graph of the consumption stream implied by asset prices.

Figure III

(c_{t+1}/c_t) Implied by Stock Returns: 1866-1925

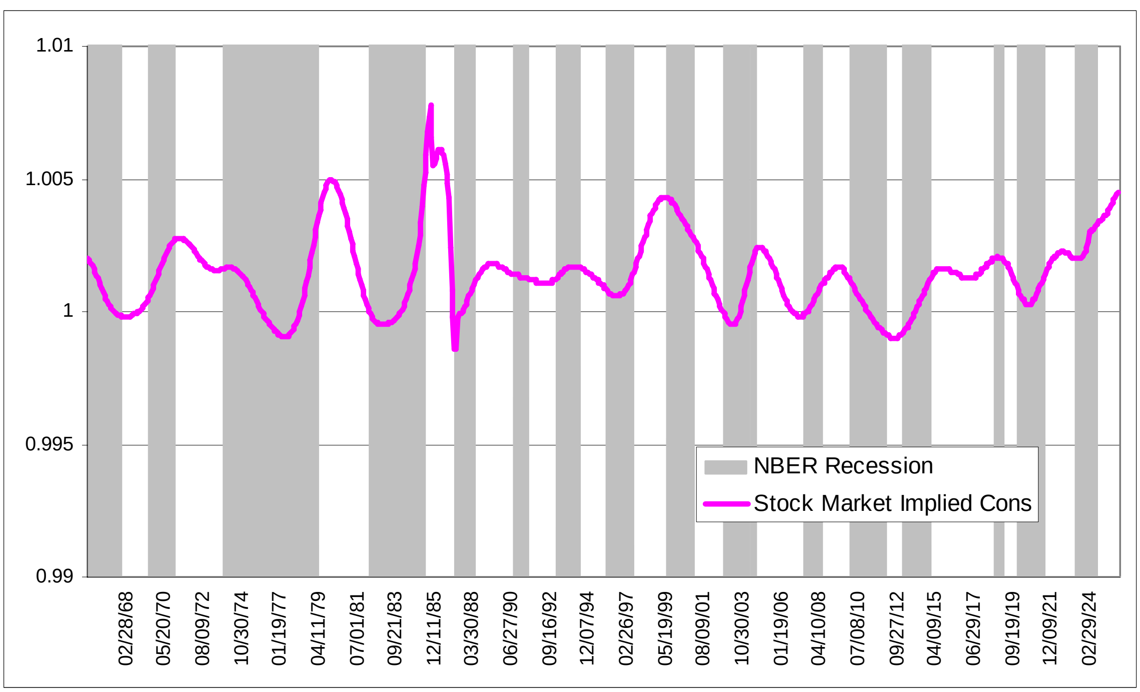

In many cases the peaks and troughs of the consumption implied by asset prices matches the NBER dates. There are some notable exceptions, however. The implied consumption series exhibits strong growth during the recessions of 1899-1900, 1918-1919 and 1923-24. The Implied Consumption series also grows at a moderate rate during the recessions of 1869-70 and 1893-94.

Table II compares the behavior of the stock market implied consumption series to Shiller's annual consumption series during the 1890-1925 NBER business cycle contractions.

Table II

	Stock Market Implied Consumption Annual Growth Rate	Shiller (1989) Consumption series Annual Growth Rate
All years 1890-1925	1.83%	1.99%
Years with NBER Contraction	1.55%	1.55%
Years without NBER Contraction	2.46%	2.99%
<u>Contraction Years</u>		
1890-91	1.80%	1.30%
1893-97	1.65%	1.43%
1899-1900	5.11%	4.39%
1902-04	0.85%	0.74%
1907-08	0.58%	-4.11%
1910-14	-0.13%	0.41%
1918-21	1.68%	2.50%
1923-24	3.37%	6.30%

The implied consumption series and Shiller's consumption series both exhibit below average growth during the contractions of 1890-91, 1893-97, 1902-04, 1907-08 and 1910-14. Both series enjoy above average growth during the contractions of 1899-1900 and 1923-24. The two series tell conflicting stories during the contractions of 1918-21 and 1907-08. During the contraction of 1918-1921 the stock market implied consumption is slightly below average while Shiller's series is well above average. During the recession of 1907-1908 both series exhibit below average growth. The Stock Market series shows poor growth but not as severe as the 1910-1914 recession. Shiller's series, on the other hand, suffers the worst decline in its history.

Implied Consumption 1825-1865

I employ Goetzmann, Ibbotson and Peng (2001) NYSE data set to estimate antebellum consumption. Goetzmann, Ibbotson and Peng's NYSE data includes the monthly prices of 300 stocks that traded on the NYSE between 1825 and 1865 and an annual dividend series for each stock. I convert the annual dividend into a monthly dividend by assigning 1/12 of the annual dividend each month.

The data is sparse in the first 10 years and completely missing in 1848 and parts of 1849. Due to the missing data I estimate consumption in the 1825-1847 period and the 1850-1865 period with rolling windows that do not overlap the missing years. Figure IV contain graphs of the 1825-1865 consumption implied by asset prices.

The implied consumption shows a sharp depression in 1825-26. The data is very sparse during the 1820s and early 1830s however, and this estimate should be viewed with caution. The well known depressions of 1837 and 1839-43 appear in data as well as a deep contraction in the 1850s.

Figure IV:

(c_{t+1}/c_t) Implied by Stock Returns: 1866-1925

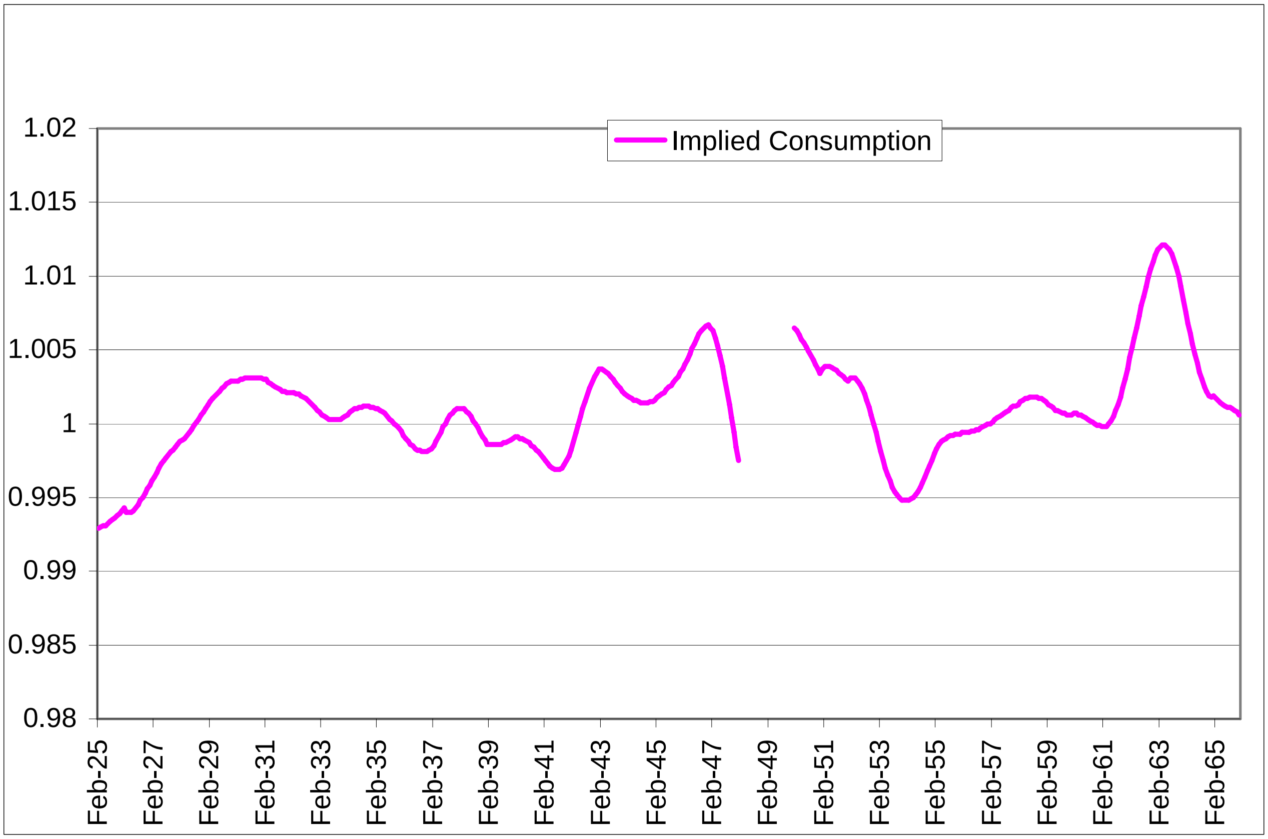

There are no NBER business cycle dates to compare too before 1854. There is one NBER contraction between June 1857 and Dec 1858 and another in 1860-61. Neither appears in the implied consumption data.

Business Cycle Dates

I began by asking if the business cycle has changed over time. Now that we have a consistent measure of business cycle activity we can ask whether economic downturns are more benign or less frequent. I follow Romer (1994) and compute business cycle dates by fitting a loss criterion to best match NBER dates during the period for which we have the most confidence in the existing data.

The business cycle dating algorithm should take into account both the severity and length of an economic contraction. Romer (1994) suggests a loss function. The loss function I choose is represented by the dashed area in Figure V. I choose a growth rate k that need not be equal to 1. The growth rate defines the aggregate consumption loss illustrated in Figure V. I choose this growth rate and aggregate consumption loss to minimize the difference between the consumption losses that occurs one defines a recession as occurring whenever consumption loss falls below this threshold and the consumption loss that occurs when one dates business cycles via the NBER business cycle dates.

For each downturn severe enough to be labeled a recession, I compute the lost consumption by taking the area bounded by the stock market implied consumption growth rate and the horizontal line at growth = k.

Figure V

Consumption Loss

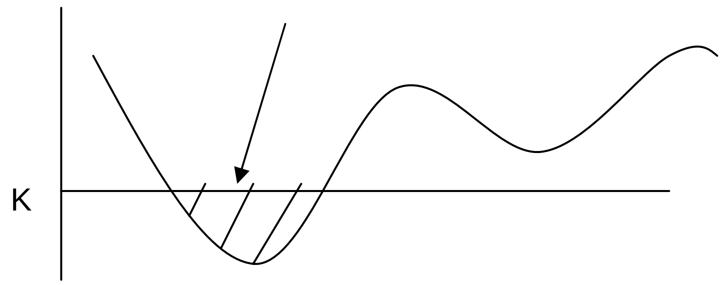

The result is a measure of consumption loss for each decline in the implied consumption series severe enough to be labeled a recession.

The optimal dating criteria sets k=1.0002 and requires an aggregate consumption loss of at least 1.87% before a downturn can be classified as a recession. This criterion successfully matches 78% of the NBER recessions dates during the CRSP era. Applying the dating rule to the entire 1825-2000 period results in the following business cycles dates.

Table III

Contractions	Intersect NBER contraction	Consumption Loss
Apr 1972 – Sep 1974	Yes	6.27%
Aug 1946 - Nov 1948	Yes*	3.51%
Feb 1937 - Aug 1939	Yes	4.30%
Apr 1929 - Aug 1932	Yes	37.30%
Feb 1911 - Dec 1913	Yes	3.09%
Dec 1875- Jun 1878	Yes	2.44%
Nov 1852 - Feb 1857	N/A	13.20%
Jul 1839-Mar 1843	N/A	8.41%
Aug 1836-Jul 1837	N/A	3.07%
Feb 1825 – Aug 1828	N/A	19.72%

* NBER contraction begins in Nov 1948

** data begins in Feb 1825

Modern consumption contractions appear to be no less severe then the contractions during the 60 years following the U.S. civil war. The past three recessions were less severe then the great depression but with the exception of the 1930s one would have to go back to the turbulent 1850s to find contractions as large.

Conclusion

This paper makes use of recently collected stock data and the restrictions of the investors' first order conditions to estimate the consumption stream implied by asset prices. The abundance of 19th century stock data allows us to extend consumption estimates back to 1825. The consumption implied by asset prices does a poor job of matching NBER business cycles but with the notable exception of the 1907-08 recession, is very similar to the 1959-2000 B.E.A. consumption series and the 1890-1958 annual consumption estimates in Shiller (1989).

The consumption implied by asset prices exhibits little evidence of a dampening business cycle. Contractions, when they occur, appear to be just as strong in the post-war period as the late 19th and early 20th Century. The 1820s-50s did appear to be more volatile, however. Likewise, the frequency of contractions has decreased since the antebellum period but there is little evidence of change since then. Using a loss criteria to best match NBER CRSP-era business cycle dates, contractions appear twice in the periods between the civil war and the 1930s and the 1930s and today.

References

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad, 2004, "Consumption, dividends, and the cross-section of equity returns", Journal of Finance (forthcoming).

Breeden, D. T., M. R. Gibbons, and R. H. Litzenberger, 1989, Empirical Test of the Consumption-Oriented CAPM, Journal of Finance, 44, 231-262.

Connor, Gregory; Korajczyk, Robert (1988) "Risk and Return in an Equilibrium APT: Application of a New Test Methodology" *Journal of Financial Economics, 21*

Davis Joseph (2004) "An Annual Index of U.S. Industrial Production, *1790-1915*" Quarterly Journal of Economics (November 2004)

Fogel Robert and Stanley L. Engerman, (1974) *Time on the Cross: The Economics of American Negro Slavery*. Boston: Little, Brown and Company, 1974. xviii + 286 pp.

Friedman Milton and Anna Jacobson Schwartz, (1963) *A Monetary History of the United States, 1867-1960.* Princeton: Princeton University Press (for the National Bureau of Economic Research), xxiv + 860 pp.

Goetzmann, Ibbotson, and Peng (2001) "A New Historical Database For The NYSE 1815 To 1925: Performance And Predictability" *Journal of Financial Markets, Volume 4, Issue 1, January 2001, Pages 1-32*

________________ Stock data online: http://icf.som.yale.edu/nyse//downloads.php

Lettau, M., and S. Ludvigson, (2001a), "Consumption, Aggregate Wealth, and Expected Stock Returns", Journal of Finance, 56, 3, 815—849.

Lettau, M., and S. Ludvigson, (2001b), "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying", Journal of Political Economy, 109, 1238-1286.

Mankiw, N. G., and M. D. Shapiro, (1986) "Risk and return: consumption beta versus market beta", Review of Economics and Statistics, 68, 3, 452-59.

Parker, J. A., and C. Julliard, (2004) Consumption Risk and the Cross-Section of Expected Returns, Journal of Political Economy, 113, 1, 185-222.

Romer, Christina (1986a) "Is the Stabilization of the Postwar Economy a Figment of the Data?" American Economic Review 76 (June 1986): 314-334.

________________(1986b) "New Estimates of Prewar Gross National Product and Unemployment." Journal of Economic History 46 (June 1986): 341-352.

________________(1986c) "Spurious Volatility in Historical Unemployment Data." Journal of Political Economy 94 (February 1986): 1-37.

_______________(1994) “Remeasuring Business Cycles.” Journal of Economic History 54 (September 1994): 573-609.

_______________(1999)“Changes in Business Cycles: Evidence and Explanations.” Journal of Economic Perspectives 13 (Spring 1999): 23-44.

Shiller Robert J.. (2003) "From Efficient Markets to Behavioral Finance," *Journal of Economic Perspectives*,

______________(1989) *Market Volatility* Cambridge, Mass.: MIT Press

______________Online Consumption data: http://www.econ.yale.edu/~shiller/data/stock%20data%20annual%201871-2003.htm