

The Three Epochs of Oil¹

Eyal Dvir (Boston College) and Kenneth S. Rogoff (Harvard)

August 11, 2009

¹We thank Jim Anderson, Deepa Dhume, Gita Gopinath, Lutz Kilian, Brent Neiman, and the participants of the macro seminar at BC for helpful comments on earlier versions, as well as Susanto Basu and Fabio Ghironi for helpful conversations. An earlier version of this paper (Dvir and Rogoff Feb. 2008) circulated under the title "Oil and Global Growth." Dvir: dvire@bc.edu, Rogoff: krogoff@harvard.edu.

Abstract

We test for changes in price behavior in the longest crude oil price series available (1861-2008). We find strong evidence for changes in persistence and in volatility of price across three well defined periods. We argue that historically, the real price of oil has tended to be highly persistent and volatile whenever rapid industrialization has coincided with uncertainty regarding access to supply. We present a modified commodity storage model that fully incorporates demand, and further can accommodate both transitory and permanent shocks. We show that the role of storage when demand is subject to persistent growth shocks is speculative, instead of its classic mitigating role. This result helps to account for the increased volatility of oil price we observe in these periods.

Keywords: Oil Price, Oil Shocks, Storage, Structural Change.

JEL classification: E0, Q4, L7, N5.

1 Introduction

Much has been written on the oil shocks of the 1970's as watershed events that have transformed the energy market; that, together with the lack of high frequency data from earlier periods, have led to an almost complete concentration on the post oil shock period among economists¹. However, much can be learned about oil price behavior from the less recent past. The crude oil price time series illustrated in Figure 1 goes back to 1861²; even a cursory look reveals stark differences in the behavior of the series at different periods. First, from 1861 until about 1878, there was a period of extremely high volatility and generally high prices. Then came a much less volatile period, approximately between 1878-1972, in which prices were also generally lower. Finally, from about 1972 onwards, we see a second period of high volatility accompanied again by higher prices.

Our first task in this paper is to document these differences and formally test for changes in behavior. We run two such tests, for changes in persistence and for changes in volatility. We find striking empirical similarities between the periods 1861-1878 and 1972-2008, in that oil prices were both significantly more persistent and significantly more volatile in these periods, both relative to the long period that separates them, i.e. 1878-1972. We also estimate that a further break in oil price volatility, but not in persistence, occurred around 1934, so that oil price in the period 1878-1934, while much less volatile than in 1861-1878, was still significantly more volatile than in the period 1934-1972³.

¹Pindyck (1999) is a notable exception.

²The series is taken from British Petroleum's "Statistical Review of World Energy", revised annually and available at www.bp.com/statisticalreview. Prices are in 2007 \$US per barrel. The series is comprised of three consecutive price series: US average price in 1861-1944, Arabian Light in 1945-1983, and Brent dated in 1984-2007. The 2008 datapoint was added to the BP series using oil prices from the U.S. Energy Information Agency and U.S. GDP deflator data from the BEA.

³All of our tests reject the null of no break with very high levels of confidence. However, the confidence intervals around the exact break dates are large enough to suggest caution regarding the interpretation of individual historical events. Our emphasis will therefore be on the broad characteristics of the periods in question, rather than on the exact date

What can explain the concurrence of price persistence and price volatility? We offer an informal, historical narrative, as well as a formal model. Our approach in this paper is to look for a unifying framework which is flexible enough to allow for the very different price behavior across periods that we observe. We find striking historical similarities between the two end-periods mentioned, 1861-1878 and 1972-2008, in terms of supply and demand factors affecting the market for oil. On the demand side, as we explain in greater detail in Section 3, both periods were years of intense industrialization in what was then becoming a major engine of the global economy: the U.S. in 1861-1878, and East Asia in 1972-2008. We see these as periods in which the demand side was characterized by persistent growth shocks. On the supply side, meanwhile, both periods featured uncertainty regarding the continued access of consumer markets to oil. This was due to the monopoly of railroads on transportation in the former period, and to OPEC's ability to restrict access to easily exploitable reserves in the latter period (see Section 3 for details). Despite the remarkable difference in the scale of the oil industry between the two periods, the effect was similar: in periods of rising demand, these agents were able to restrict access to additional oil supplies, thereby causing prices to rise.

We argue in this paper that this confluence of supply and demand factors can explain why we observe large changes in oil price persistence over the years: persistent growth shocks to demand, if occurring in periods in which key players in the market had the ability to restrict access to supplies, were translated to very persistent price behavior. In other words, the monopolistic industry structure in these periods coincided with uncertainty regarding demand trends to produce uncertainty regarding price trends. We show in Section 3 that the change-points between periods that we identify in our formal testing correspond to dramatic shifts in the structure of the oil industry. During periods when there were no effective restrictions on access to sup-

of change from one period to the next.

plies, even large and persistent shocks did not cause more persistent price behavior. Rather, the shocks were accommodated through a relatively quick supply response⁴. Historically, this was by far the more prevalent pattern: for the most part the history of oil has been characterized by relatively easy access to needed oil. Consequently the market trend was quite stable from 1879 to 1971.

The historical narrative suggests reasons why certain periods would exhibit substantially more oil price persistence than others, but cannot by itself explain the observed concurrence of high persistence with high volatility. A theoretical framework for oil should be able to accommodate both phases of the market, and to explain both the observed persistence and volatility behavior of oil prices. Our third contribution in this paper is to present a model which does just that: it is an extension of the canonical commodity storage model à la Deaton and Laroque (1992, 1996), in which we introduce growth dynamics to their well-known framework. This results in a model that can accommodate both $I(0)$ and $I(1)$ stochastic processes, so that periods of stable and stochastic trends can both be considered. The model can explain our main empirical findings: it predicts that in the presence of uncertainty regarding the trend, rational storage behavior will act to *enhance volatility*. In the standard commodity storage framework, where uncertainty is in regard to deviations from trend, but the trend itself is viewed as stable, storage acts to *reduce volatility*. This feature of the storage model is in itself novel, as well as useful in terms of explaining the observed patterns in the data. We present simulations in which this behavior can increase price volatility following growth shocks significantly.

The large literature on commodity price behavior falls broadly into two

⁴Examples abound: the shocks to demand imposed by the needs of the two world wars, or by the postwar reconstruction of Europe, were large, persistent, and open-ended. However none of these major upheavals affected the price of oil much. When agents tried to restrict access to supplies their efforts were futile, as U.S. producers learned in 1918, when the Federal government threatened to draft their workers if they did not comply (Yergin [1991], page 179).

major strands, depending on whether the commodity in question is perceived to be renewable. On the one hand, models of storage have been used mostly to study renewable commodities such as corn and wheat; see Wright (2001) for a survey of theory and evidence⁵. The study of non-renewable commodities, a definition which includes oil, has followed an altogether different path, strongly influenced by the seminal contribution of Hotelling (1931). Krautkraemer (1998) surveys the theory and evidence. In the current paper we choose storage as the more useful of the two strands. In this we are motivated primarily by the empirical evidence, which shows quite clearly that finite availability of oil - a separate issue from that of free access to currently available supplies - is not of first order significance in explaining oil price behavior. In particular, proven world oil reserves have been increasing in recent decades, in spite of ever increasing production⁶. As a result, it may well be that technological advances in oil exploration and utilization will be enough to satisfy demand in the foreseeable future. That is the assumption that we make in this paper.

Our work is also related to the ongoing debate on oil and the macroeconomy (see the two recent surveys by Hamilton [2008] and Kilian [2008a]). This literature focuses on correctly identifying the source of shocks to oil prices, dealing exclusively with the post-1973 period, usually with higher-frequency data. We argue in this paper that a long-term view is essential to this debate: shocks to the oil market have had remarkably different effects on the real price of oil across historical periods, not because of their origin on the supply or the demand side, but rather because of the ability (or lack thereof) of key players in the market to restrict access to supplies. In particular, in periods when the ability to restrict access to supplies was lacking, the oil

⁵A notable exception is the paper by Routledge et al. (2000), who find that oil prices exhibit the strong mean reversion associated with storage models, as well as a permanent factor. However these authors focus only on the very recent past (1992-96), and do not offer a theoretical justification for the permanent factor.

⁶See BP Statistical Review (2008) for proven reserves and production data from 1980.

market showed remarkable flexibility and relative price stability, even in the face of massive disturbances, and regardless of their source in the supply or demand side of the market.

The paper proceeds as follows: section 2 presents our empirical findings on oil price behavior over time. Section 3 puts these findings in the context of the history of supply and demand for oil. Section 4 introduces the model, and Section 5 examines the model behavior under both transitory and permanent shocks. Section 6 concludes.

2 Behavior of the Real Oil Price: Then and Now

Table 1 presents some sample statistics regarding each of the three periods delineated in the introduction, as well as the entire series. The differences between the sub-samples are quite clear: the mean price between the years 1861-1877, at \$50.9 (measured in 2007 U.S. dollars), was almost triple the mean price between 1878-1972, at \$17.2 . In contrast, there is only a 13% difference (again in 2007 U.S. dollars) between the mean price between the first period, 1861-1877, and the third, 1973-2008, at \$44.3 . A similar pattern holds for differences in the unconditional standard deviation of annual prices across these periods: at \$25.3, the standard deviation of price in the period 1861-1877 is five times higher than that of the period 1878-1972 (\$5.1), but only 12% higher than the standard deviation of price in the years 1973-2008 (\$22.3). The next lines of the table compare the rates of change across periods. Note in particular that the unconditional standard deviation of absolute price change (a common measure of volatility) behaves in an analogous manner: it is quite similar in the first and third periods (27.5% between 1861-1877 compared with 22.7% between 1973-2008), but remarkably lower in the intervening period. In sum, Table 1 shows that there seems to be much in common in terms of the behavior of real oil prices between the periods

1861-1877 and 1973-2008, while both periods look markedly different in most respects from the intervening period of 1878-1972.

Studies of the time series properties of real oil prices have taken one of the following approaches in the face of these clear non-linearities in the series: either analyzing the series as a whole, or, much more commonly, treating the series as composed of separate series "pasted together" (in the words of Hamilton, 2008), and proceeding to analyze them in isolation. In one important category, that of determining whether or not oil prices exhibit a unit root, these different approaches have led to opposite conclusions. Pindyck (1999), an example of the former approach, ignores the aforementioned differences, and judges the entire series to be mean-reverting to a moving quadratic trend. At the opposite end, Hamilton (2008) notes the abrupt change in the series, and proceeds to analyze the third period only (with quarterly data). He accordingly determines that the real price of oil follows a random walk with no drift⁷.

In what follows we will treat both the assumption of a pure $I(0)$ process and the assumption of a pure $I(1)$ process as our null hypotheses, and test whether the series exhibits a shift from $I(0)$ to $I(1)$ (or vice versa) against *both* of these assumptions. In other words, instead of trying to decide whether the series as a whole or in part exhibits a unit root, we aim to determine whether it shows clear transitions from a stochastic trend to a deterministic one, and vice versa. In order to do that, we employ a series of tests proposed by Harvey, Leybourne, and Taylor (2006, HLT henceforth). This series of tests is a modified version of tests for change in persistence proposed earlier by Kim (2000), Kim et al. (2002), and Busetti and Taylor (2004), all of whom build on the unit root testing method of Kwiatkowski *et al.* (1992)⁸. What makes the HLT test of change in persistence appealing in our context is that

⁷Studies of higher frequency data from the 1990s (daily, weekly) do find a mean-reverting factor as well as a permanent factor. See for example Routledge et al. (2000), Schwartz and Smith (2000).

⁸Appendix A of Dvir and Rogoff (2009) provides details.

it maintains its properties of consistency and appropriate size both under the $I(0)$ null and under the $I(1)$ null. This allows us to test for structural change without taking an a-priori stand regarding the null hypothesis. A disadvantage of the HLT modified test is that it does not take account of possible heteroskedasticity of the stochastic process. This is important, since oil price volatility does in fact exhibit structural breaks, as we shall see later in this section. However, applying a bootstrap method to the tests can correct this issue: following Cavaliere and Taylor (2008), we use a wild bootstrap⁹ to arrive at heteroskedasticity - robust bootstrap p-values for our test statistics¹⁰.

Table 2 presents the results of the HLT change of persistence test, using the real oil price series. Since the test is designed to find a single change-point, whereas the series exhibits two obvious break candidates, we conducted the test separately for periods 1 and 2, and for periods 2 and 3. The exact end points shown in the table (1881 and 1965) were chosen arbitrarily; the qualitative results are robust to small changes in these end points¹¹. The table reports three sets of test statistics, all functions of Kim's (2000) ratio-based test statistic K . The null hypothesis in each test is that there is no change in persistence present; the alternative is that a single change in persistence occurs. The testing procedure first computes the value of K

⁹This technique was introduced by Hansen (2000). Rather than re-sampling, it generates a bootstrap sample by multiplying the residuals from regressing the series on the exogenous variables, with random numbers from a normal distribution. In this way the pattern volatility present in the original shocks can be replicated in the bootstrap sample.

¹⁰The HLT modified test-statistics are sized correctly only in the homoskedastic case, and are therefore not well suited for the wild bootstrap procedure. However these statistics are computed directly from, and are uniformly lower than, the non-modified test statistics. We bootstrap the non-modified statistics, and compute the robust p-values using the lower modified statistics. These p-values should therefore be viewed as biased upwards.

¹¹As a ratio-based test, the HLT change in persistence test is designed to reject the null if a point can be found after which the behavior of the series is statistically different than its behavior before that point. When two such points exist, and moreover the behavior of the series before the first point is similar to behavior after the second point, this test will lose power.

for each potential break-point. As the true break-point is unknown, the procedure then computes the mean value (MS), the mean-exponential value (ME), and the maximum value (MX) taken over all these values of the test statistic. To test the alternative hypothesis of a change in persistence from $I(1)$ to $I(0)$, the same procedure is performed, with the statistics MS^R , ME^R , and MX^R computed in the same way, taken over the reciprocal values of K . As is common in this literature, we test for a change in persistence from a trend - stationary process to a difference - stationary process (and vice - versa). These tests are robust to the inclusion of higher - order polynomial trends, such as a quadratic trend as in Pindyck (1999). It should also be noted that by construction, the tests can be performed regardless of the type or order of time-dependence in the stationary process, as long as some regularity conditions hold (see Kwiatkowsky *et al.*, 1992 for original proofs). Estimates for the break-points follow Kim (2000).

Testing for change in persistence in the years 1861-1965, then, we find very strong evidence for a significant change in from a high-persistence, local to $I(1)$ process, to a low-persistence $I(0)$ process, where the point of change is estimated to be in 1877. This is shown by the very high values of the relevant test statistics, MS^R , ME^R , and MX^R , which are all significant at the 1% level. In the period 1881-2008, we find strong evidence for a change in persistence from a low-persistence $I(0)$ process to a high-persistence, local to $I(1)$ process, with the point of change estimated at 1972. Again, all three of the relevant test statistics, MS , ME , and MX , point to the same conclusion. Two of the test statistics have bootstrap p-values that are less than 0.001, while the third's (MS) bootstrap p-value is slightly higher at 0.021¹². We use

¹²Note that for our test to reject either null, price persistence on one side of a tentative break should be statistically different from price persistence on the other side of the break, but need not meet any particular critical value. Because of this, Caner and Kilian's (2001) criticism of KPSS tests, namely, that they may suffer from size distortions and therefore produce spurious rejections, does not apply to our paper. Moreover, we get extremely high values for our test statistics, associated with very low bootstrap p-values. Therefore the likelihood of a Caner-Kilian type spurious rejection in our data is low.

simulation results from Kim (2000) to calculate the 95% confidence interval around our change-points. As can be seen from Table 2, these are estimated somewhat imprecisely, with confidence intervals of 8 and 10 years for the first and second change-points respectively¹³. However the rejection of both the null of a pure $I(0)$ process and the null of a pure $I(1)$ process is quite clear, implying that a more nuanced view of the series is in order, a view which takes into account potential changes in persistence behavior.

Apart from the rate of persistence, another time series aspect of real oil prices that has attracted much attention in the literature is their volatility. As already mentioned, volatility (as measured by the standard deviation of absolute rates of growth) was high before 1878, low from around that time until the early 1970s, then high again until the end of our sample in 2008. We therefore conducted a test for multiple breaks in oil price volatility, using the methods of Bai and Perron (1998, 2003). We define volatility as the mean absolute residual from a regression of oil price growth on its lagged value. The null hypothesis is that volatility is constant throughout. The test results are shown in Table 3, and illustrated in Figure 2. We identify three potential breakpoints: 1878, 1934, and 1972. All three test statistics against the null of no break are highly significant, implying that the series contains at least one breakpoint. However, in deciding how many breakpoints there are, the various criteria explored by Bai and Perron do not agree: their sequential procedure selects only one breakpoint, in 1878, whereas the Bayesian Information Criterion selects all three¹⁴. A look at the coefficients denoting mean volatility in the different periods can explain this discrepancy. We see that the coefficients of Periods II (1879-1934) and IV (1973-2008) are very similar, and both are quite different from the coefficient for Period III

¹³Kim (2000) finds in simulations that for change-points that occur at the 25th or the 75th percentile of a given series, his procedure for estimating the change-point location has a maximum standard deviation of 1.6893 for $T=100$. We accordingly use this value, scaled to our sample length, to calculate our 95% CIs.

¹⁴The LWZ information criterion also chooses 1878 as the only break; however, this criterion is known to perform badly when breaks are present (i.e. the alternative is true).

(1935-1972). As Bai and Perron recognize (2003, pp. 15-16), in these cases the sequential procedure can be improved upon: the number of breaks should be chosen according to the last significant test statistic, instead of the usual practice of choosing according to the first insignificant test statistic. In the current case, as seen in Table 3, this improved sequential procedure puts the number of breaks at three, similarly to the BIC. Oil price volatility then has gone down by about half sometime in the last quarter of the nineteenth century (with 1878 as our best estimate), then gone down again by about two thirds around 1934. When it increased again, according to our estimate in 1972, it regained its level of the early twentieth century, but did not reach the heights set by oil prices before 1878. Note that 95% confidence intervals for the change-points are quite large; as in the change in persistence test, our confidence in the occurrence of breaks in the series' behavior is far stronger than our confidence in the exact dates of these breaks. Nevertheless, these years will be useful as anchors in our historical narrative in Section 3.

We can sum up our empirical findings as follows: real oil price from 1861-1877 (or 1878) was highly persistent and volatile, from 1878-1934 was not as persistent and less volatile, from 1934-1972 it was still not very persistent and displayed even lower volatility. Finally, from 1972 on the real price of oil returned to being highly persistent and volatile, though not as volatile as in the pre-1878 period. Later on, in Section 4, we present a model of the oil market that ties together these patterns of price behavior. But first we need to put our model in historical context, which is the purpose of the next section.

3 Industrialization and Market Structure: Transition Points in Context

In Section 2 we identify three points of transition. In 1877-8 and again in 1972, oil price persistence and volatility both changed, while in 1934 we find

a change in volatility, but no change in persistence. Of these three points, only 1934 can be linked to a major oil discovery, that of the East Texas Oil Field a few years earlier. The other two points of transition, we will argue, had to do with technological and geographic factors that enabled changes in market structure. In 1878 began the construction of Tidewater, the first long-distance pipeline, which eventually ended railroad monopoly over the transportation of oil. In 1970 the East Texas Oil Field peaked, ending U.S. control over excess exploitable reserves, and signalling the rise to prominence of OPEC¹⁵.

These three points also show striking similarities and differences from a demand point of view as well. Of these three points, 1934 is special also in that it occurs in the midst of a worldwide recession. Both 1877-8 and 1972 were years in which the global economy, and with it demand for oil, were booming, driven by the large-scale industrialization of the United States and of East Asia, respectively. Rapid industrialization is by definition a transitional stage, and as such it features growth rates that are on the one hand unsustainably high and on the other hand quite persistent, since the process of industrialization often stretches over decades. In other words, periods of rapid industrialization are characterized by persistent growth shocks to income.

We will argue that the two observed periods in which oil prices were both highly persistent and highly volatile occurred because two conditions were simultaneously met in each of these periods: access to supplies was restricted and demand was unsustainably high. It is important to note that U.S. industrialization was far from over when the first such period ended in 1878, and that post-war industrialization in East Asia was well underway by

¹⁵The oil industry was, of course, much bigger in the late 20th century compared to its size a hundred years earlier; oil is also no longer used mainly for illumination as it did at first. It is notable however that prominent features of the industry remain relatively unchanged: oil was internationally traded from the very beginning, demand for it being global. Moreover, its efficiency as a source of energy made it indispensable to consumers from the earliest stages, a feature of the industry that remains crucial to this day.

1972, our estimate of the beginning of the second period of high persistence and high volatility. These years were not turning points in oil demand, rather they signified major structural changes in the petroleum industry, in which key players with the ability to restrict access to supplies either emerged or declined in importance. When only one of the conditions was met, as for example happened during both World Wars (when demand was unsustainably high but supply was unrestricted), the market was significantly less persistent and less volatile. This necessary confluence of demand and supply factors has been relatively rare looking back all the way to 1860, but of course has been the reality in the oil market in recent decades¹⁶.

It is worth emphasizing that we do not focus on the source of shocks to the oil market, nor do we attempt to identify these shocks. Our main objective is to account for the radically different behavior of oil prices across different regimes, rather than determining the genesis of shocks to oil prices within a given regime. Studies such as Kilian (2008b, 2008c) which focus on the post-1973 experience exclude this type of analysis by virtue of their intentionally limited scope. Their results by construction do not extend to environments in which the oil industry's structure is radically different.

By 1865, the Oil Regions in northwestern Pennsylvania, where all commercial oil production was located, were well served by three different railroads¹⁷. These firms enjoyed an oligopolistic position, as both production and refining were highly competitive. For illustration purposes, in 1877, the year before Standard Oil consolidated its control of refining, the "open fare" for rail transport of crude oil from the Oil Regions to New York was \$1.40 per

¹⁶We stress supply restrictions rather than reserve depletion since there is no evidence that "running out of oil" was ever a real danger. Oil security, the danger of not having access to existing oil, was on the other hand very real. In this regard, "capacity constraints" must be viewed as mechanisms to restrict supply, since by their nature these constraints - derricks, storage tanks - can be loosened in the medium run, whereas the amount of extractable oil in any given field cannot be increased beyond a certain point.

¹⁷What follows draws much from Yergin (1991). For a more detailed account, see the NBER Working Paper version (Dvir and Rogoff, 2009).

barrel (Bentley, 1979, page 28); this amounted to 58% of the average price of a barrel of crude oil in that year according to our data. Williamson and Daum (1959, Chap. 17) estimate the per barrel cost of carriage by rail at no more than \$0.40 around that time, giving us an idea of the margins involved. Rockefeller's vision was a large refining concern that could bargain effectively with the railroads; Standard's business advantage was well understood at the time to consist of the special "rebates" that it was in a position to demand from the railroads¹⁸.

In 1878, oil producers who were trying to break the joint monopoly of transportation and refining started construction on the world's first long-distance pipeline, the Tidewater. It was completed in May of 1879, and was intended to by-pass both the railroads and Standard Oil's refineries. In the face of this technological breakthrough, Rockefeller proceeded to construct Standard's own long-distance pipelines, choosing to destroy the railroads' monopoly on transportation in order to strengthen Standard's monopoly on refining. This spelled the end of attempts to increase profitability by restricting access to oil. Having invested in his own infrastructure, and given the very low transportation costs it afforded him¹⁹, Rockefeller's strategy now was to sell as much oil as possible.

These events occurred against a backdrop of ever rising demand, both domestic and foreign. The United States was going through rapid industrialization at the time, eventually overtaking Britain as the world's leading center of manufacturing. During the period, the share of world industrial output made in the U.S. rose spectacularly, from 7.2% in 1860, to 14.7% in 1880, to 23.6% in 1900. In absolute numbers, U.S. manufacturing production rose by a factor of three in the two decades between 1860 and 1880, and by

¹⁸Standard's business advantage over independent refiners as a result of its strong bargaining position is estimated by Bentley (1979) to have been \$1.00 per barrel of refined oil.

¹⁹Williamson and Daum (1959, p. 458) estimate that per barrel transportation costs in Standard Oil's own pipelines were between \$0.12 - \$0.20, less than half the rail cost of carriage.

a factor of eight between 1860 and 1900²⁰. U.S. population more than doubled from 1860-1900, rising from 31.8 million to 76.4 million, while GDP per capita rose almost as fast, from \$2,445 in 1860 to \$4,091 in 1900 (constant 1990 dollars)²¹. As a result domestic consumption of illuminating oil rose from 1.6 million barrels (mb) in 1873-75 to 12.7 mb in 1899, while that of machine lubricating oil rose even more, from 0.2 mb in 1873-75 to 2.4 mb in 1899 (Williamson and Daum [1959], pp. 489, 678). Even as urban communities in the United States and Europe shifted to gas or electric lighting, kerosene remained in high demand in other parts of the world. By the turn of the twentieth century, there was increasing demand for gasoline, from the burgeoning auto industry.

The transition point we identify in 1877-8 was therefore the starting point of sweeping changes in market structure, brought about in an environment of rapidly growing demand. Before 1878 the railroads were using their monopolistic position to limit the supply of crude to the markets in the interest of rent extraction. After 1878 that power was slipping away from them at a fast clip; by 1884 Rockefeller's network of long-distance pipelines was essentially complete, and the railroads were sidelined. Moreover, since Standard Oil owned the vast majority of long-distance pipelines, and with demand expected to continue unabated, there was no player in the market who had both the interest and the capability of limiting supplies²².

This state of affairs continued until the early 1930's, when the discovery of the East Texas Oil Field created an oil glut of proportions heretofore unknown. Against the backdrop of Depression, the result was a slump in price that threatened the entire industry, eventually leading to Federal regu-

²⁰Bairoch (1982) is the source for the U.S. absolute and relative industrial output numbers.

²¹Figures for U.S. population and GDP per capita are from Maddison (2003).

²²The producers in the Oil Regions were relatively small-scale, and had repeatedly failed in their attempts to control production. There was one exception: in 1887-8, there was a willingness to cooperate on the part of Standard Oil, and production reduction was achieved for a short while.

lation. The U.S. government came to effectively control supplies: since East Texas production was far below its potential, and given the authority to raise and lower production quotas as circumstances required, the U.S. government (both Federal and state, in particular the Texas Railroad Commission) had the power to increase or decrease oil supply almost at will. Over the decades since, while it still had that power, the U.S. government would use it to stabilize the market on numerous occasions. It increased production enormously during World War II, as well as during supply crises involving the Middle East, in 1953 (Iran), 1956 (Suez), and 1967 (Six-Day War). When the surge of oil was no longer needed, it had the power to reduce production once more. U.S. regulation thus acted as an automatic stabilizer (Yergin, 1991, page 259). This had the effect of reducing the standard deviation of supply and demand shocks, and accords well with the observed reduction in volatility that we date to 1934, around the time that this mechanism went into effect. Quite the opposite from the railroads' rent extraction strategy before 1878, U.S. government agencies aimed to stabilize price by adjusting quantity as needed. The supply of oil, far from limited, was in effect quite flexible.

Our third transition point is 1972, where we find that oil price persistence and volatility both increased. In 1970 U.S. oil production reached its peak. In March 1971 the Texas Railroad Commission, for the first time since World War II, allowed production at 100% capacity; the ability of U.S. government agencies to increase production in times of need was gone (Yergin, 1991, pp. 567-8). Excess capacity existed now only in the Middle East, giving the rulers of these countries the same kind of market power enjoyed by the railroads almost a century earlier: the ability to extract large rents from consumers by limiting access to oil supplies. In the Gulf, 1972 saw an agreement of "participation" of oil producers, i.e. the transfer of some ownership rights of the oil resources located on their land from the international oil companies to the governments. These developments changed fundamentally the nature

of the market: the oil producing countries were now owners (whole or part) of their reserves, the only easily-exploitable oil reserves left in the world. In November 1973, of course, OPEC member states acted to dramatically restrict the West's access to oil supplies²³.

As in the early years of the oil industry, these events were occurring in a period of increasing demand. The demand for oil is driven, first and foremost, by income. In recent decades, world GDP and global oil production have moved in lockstep; the International Energy Agency estimates long-run income elasticity of world oil demand at about 0.5, i.e. each percentage point increase in world GDP is accompanied by a 0.5% increase in the global demand for oil (IEA [2006,2007]). This may in fact be an underestimate: Gately and Huntington (2002) find that income elasticity of oil demand in OECD countries is 0.55, but for non-OECD countries the income elasticity may be as high as 1. The IEA estimates that in 2001-2005, China's higher-than-average propensity to consume oil may have raised the global income elasticity to 0.8 (IEA [2007]).

This time it was East Asia that was industrializing fast: first Japan, then Taiwan and South Korea, and finally China. Japan's GDP per capita, for example, more than tripled in two decades, rising from \$3,986 in 1960 to \$13,428 in 1980. Japanese GDP by itself already equaled 37% of U.S. GDP by 1980 (all comparisons in 1990 international dollars, Maddison [2003]). China industrialized slightly later, but the pace of its industrialization in the final decades of the twentieth century was as rapid as that of the U.S. a century earlier, if not more: between 1980 and 2000, Chinese industrial production rose by a factor of 9; between 1970 and 2000, it grew by a factor of 21. In relative terms, the Chinese share of world industrial output was only

²³There is a debate in the literature on whether OPEC can be shown to have acted collusively to withhold supplies from the market. Clearly OPEC's degree of control over prices has been inconsistent over the years, however that in itself does not settle the issue: Smith (2005, 2008) and Almoguera and Herrera (2007) are recent references. In our context, what matters is only OPEC's ability to restrict access to the world's only easily exploitable reserves of oil. This ability is undisputed (Smith, 2008).

0.7% 1970; it has increased to 6.3% in 2000²⁴. The IMF's World Economic Outlook (2008) projects that Asia's share of global trade and manufacturing will continue to soar in the coming decades, despite the short term dislocations caused by the current financial crisis. Overall, the average growth rate in Asia (excluding Japan) between 1973-2001 was 5.4%, compared with 2.1% for Western Europe, 3.0% for the United States, and 2.7% for Japan during the same years (Maddison [2003]). It seems clear that Asia, outside of Japan, was still very much in the midst of an era of industrialization at the onset of the present century, with no end in sight. This is similar to the situation of the U.S. economy about a hundred years earlier.

We see therefore a repeat, on a much broader scale, of important features from the market environment that prevailed before 1878: a combination of supply limits and ever rising demand²⁵. As in the earlier period, with supply limited in this fashion, shocks to demand would be fully incorporated into the price. Since these shocks are very persistent, in an era where the trend in demand is uncertain, we would expect the price of oil to be more persistent, relative to a period where these limits on supply are not binding. This persistence in the price of oil can be reasonably expected to continue as long as demand shocks are persistent, or until the ability of OPEC to effectively limit access to supplies no longer exists, either due to an independent source of oil, or to an alternative source of energy.

4 An Extended Commodity Storage Model

Our model is an extension of the classic commodity storage framework. Chambers and Bailey (1996) and Deaton and Laroque (1996) extend the

²⁴Statistics for China are from the World Bank's World Development Indicators database.

²⁵OECD growth was also high at various points during this time period, a fact that no doubt has been important in the timing of the first and second oil shocks (see Barsky and Kilian [2002, 2004]). However this fact cannot account for the very high price persistence that we observe in this period.

model to allow for autoregressive shocks. We extend it further to explicitly incorporate demand, and to allow for growth shocks²⁶.

4.1 Availability and Storage

Time is discrete, indexed by t . The market for oil consists of consumers, producers, and risk neutral arbitrageurs. The latter have at their disposal a costly storage technology which may be used to transfer any positive amount of oil from period $t - 1$ to period t . Storage technology is limited by a non-negativity constraint, i.e. the amount stored at any period cannot drop below zero. This implies that intertemporal arbitrage, although potentially profitable, cannot always be achieved. In these cases, where the non-negativity constraint is binding, the market is "stocked out".

Define *oil availability*, denoted A_t , as the amount of oil that can potentially be consumed at time t . In other words, this is the amount of oil that has already been extracted from the ground, either in period t or at some point in the past, and has not been consumed before period t . It is given by

$$A_t = X_{t-1} + Z_t, \tag{1}$$

where X_{t-1} denotes the stock of oil transferred from period $t - 1$ to t , and Z_t denotes the amount of oil that is produced at time t . For simplicity, we assume that no oil is lost due to storage²⁷. Decisions concerning both variables - how much to store, how much to produce - are assumed to have been made before period t began. In period t agents must decide how to divide A_t between current consumption Q_t and future consumption, so that demand - the sum of current consumption and the amount stored for the

²⁶Other papers which extend the storage model in various ways include among others Alquist and Kilian (2008), Ng and Ruge-Murcia (2000), and Routledge et al. (2000). These papers do not seek to explain the different behavior of oil prices across historical periods as we do here, nor do they incorporate growth shocks explicitly into the framework.

²⁷Alternatively, we could have specified storage costs by a given loss percentage, as in Deaton and Laroque (1996).

future - must always equal current availability:

$$A_t = Q_t + X_t. \tag{2}$$

4.2 Demand for Oil

Let Y_t denote a demand parameter, which can be thought of as some known function of world GDP. Y_t therefore represents the income effect on the demand for oil.

We can then write an inverse demand function for oil as follows:

$$P_t = P(Q_t, Y_t), \tag{3}$$

which is decreasing in its first argument, and increasing in its second. This inverse demand function constitutes a departure from the canonical model, where demand for the commodity is a function of price only, effectively assuming no income effects. This departure is a natural one to make, however, in the context of oil, where oil consumption and income are very highly correlated (see references in Section 3). Indeed, we posit an inverse demand function in which only the *ratio* of consumption to income matters, so that if both variables grow at the same rate, the price of oil will remain constant. We assume therefore that the inverse demand function (3) is homogeneous of degree zero:

$$P_t = P(Q_t, Y_t) = P\left(\frac{Q_t}{Y_t}, 1\right) = p(q_t), \tag{4}$$

where lowercase letters denote variables normalized by Y_t . We think of the normalized variables as "effective" amounts, in the sense that a growing income leads to higher energy needs, spreading any given amount of oil more thinly. A rise in Y_t would therefore, *ceteris paribus*, decrease the effective amount of oil available for consumption and cause a rise in current price²⁸.

²⁸A disadvantage of using normalized quantities is the difficulty in directly calibrating the model to actual observable quantities. This would be an important issue if we had

We will use a CES inverse demand function:

$$P_t = q_t^{-\gamma} = (a_t - x_t)^{-\gamma}, \quad (5)$$

where $\gamma > 1$ is the inverse elasticity of demand, and a_t , x_t denote effective availability and storage in period t , respectively. It is natural to assume that the effective demand for oil is inelastic with respect to price. As equation (5) makes clear, for a given supply of oil, price is a function of the competing demands of current and future consumption. If the desire to consume more in the future grows (driven by expectations of future conditions), more oil is stored rather than consumed today, resulting in a price rise today even though supply has not changed.

We now turn to the specification of Y_t , the income effect. We consider two alternative assumptions regarding the particular stochastic process that Y_t will follow. First, we consider a simple AR(1) process, analogous to the stochastic process that Deaton and Laroque (1996) consider for supply. In particular, under this assumption we have

$$\frac{Y_{t+1}}{\bar{Y}_{t+1}} = \left(\frac{Y_t}{\bar{Y}_t} \right)^\rho e^{\varepsilon_{t+1}}, \quad (6)$$

where $\varepsilon_{t+1} \sim N(0, \sigma_\varepsilon^2)$ is an iid shock, and \bar{Y}_t is trend income, i.e. the level of income that would prevail at time t in a world without income shocks. Trend income may be constant, i.e. $\bar{Y}_t = \bar{Y}$ for all t , or increasing over time at rate $\bar{\mu} > 0$. The analysis of both cases will be quite similar. We think of this case as more closely relevant to income shocks in developed economies, where the economy exhibits business cycles around a stable trend.

Second, we will consider the case where income is subject to growth shocks, and therefore has a stochastic trend. Specifically, we now assume

the ability to perform such calibration. Unfortunately, quantity data for the oil industry - production, consumption, stocks - are not readily available for the full period (1861-2008). Our model is highly stylized as a result.

that

$$Y_{t+1} = e^{\mu_{t+1}} Y_t, \quad (7)$$

such that

$$\mu_{t+1} = (1 - \phi)\bar{\mu} + \phi\mu_t + v_{t+1}, \quad (8)$$

where $\phi \in (0, 1)$ is a persistence parameter and $v_{t+1} \sim N(0, \sigma_v^2)$ is an iid shock. It is also possible here to express the stochastic process in terms of the ratio between Y_t and trend income \bar{Y}_t as in (6). Dividing both sides of (7) by \bar{Y}_{t+1} we get:

$$\frac{Y_{t+1}}{\bar{Y}_{t+1}} = e^{\mu_{t+1} - \bar{\mu}} \frac{Y_t}{\bar{Y}_t}. \quad (9)$$

We think of this case as more relevant to income shocks in some developing countries, in particular quickly industrializing economies such as the U.S. in the second half of the nineteenth century or China in the last decades of the twentieth century. In these economies very high growth rates can be extremely persistent. In principle the world price should be affected by developments in both types of economies, depending on the relative intensity of oil use. Specifically, a positive income shock in an advanced economy such as the twentieth century U.S., where the trend is stable, is expected to disappear over time. In contrast, a positive income shock in an emerging economy such as the late twentieth century China is perceived to have a more lasting effect, as China continues its long march towards becoming a high income economy.

4.3 Supply of Oil

In the canonical commodity storage model, supply Z_t varies according to some stochastic process ψ_t around a predetermined mean \tilde{Z}_t , and it is this variability in supply that creates an incentive for inter-temporal smoothing by the large pool of risk neutral arbitrageurs. As the literature has long recognized, demand and supply shocks in the canonical model are isomorphic:

one can think of a negative realization of ψ_t as representing an especially cold winter (demand) or a breakdown in a major pipe (supply). For this reason, since we model demand explicitly, it would be redundant to model supply shocks separately. Our choice to model demand and not supply explicitly has to do, of course, with the argument of Section 3.

However, supply in our model is not constant. Rather, it grows at the trend income rate $\bar{\mu}$. That is,

$$Z_{t+1} = \tilde{Z}\bar{Y}_t. \quad (10)$$

We include trend income \bar{Y}_t in equation (10) to capture the effects of technological progress: next period's supply depends on current technology. Our assumption here is that overall technological progress, which drives global GDP growth, applies to the oil extraction and exploration sectors as well. Accordingly, while the total amount of oil existing in the earth's crust is indeed finite, technological progress is key to exploiting an increasing fraction of it over time. Note importantly that our assumption is that oil supply depends on the *technology* driving income growth, and not on income growth itself. Therefore shocks to demand will drive a wedge between supply and demand, causing a shift in equilibrium price. When these shocks are lasting, as in the growth shocks case, the changes in equilibrium price will also be lasting.

4.4 Storage of Oil

The defining characteristic of the canonical model is the availability of storage technology. Private storage is essentially arbitrage: agents will buy oil and store it if the expected future price, adjusted for finance and storage costs, is higher than the current price. As is common in the literature, we assume free entry into the storage sector as well as risk neutrality, implying that the actions of arbitrageurs will raise the current price until it is high enough to

render the strategy unprofitable in expectation. Conversely, if the expected future price is low relative to the prevailing price, agents will reduce their stocks, causing current price to drop. Note however that stocks cannot be negative, limiting the efficacy of inter-temporal smoothing in that case.

The amount being stored X_t and the price of oil P_t are determined together in equilibrium, given the realization of the exogenous parameter Y_t . When equilibrium at time t is fully optimal, i.e. when the storage non-negativity constraint is not binding, the price of oil must obey the following arbitrage condition:

$$P_t = \beta E_t[P_{t+1}] - C, \quad (11)$$

where $\beta = 1/(1+r)$ is the discount factor, and $r > 0$ is the exogenously given interest rate. The parameter $C > 0$ denotes the per barrel cost of storage²⁹. Equilibrium price P_t must be such that there is no incentive to increase or decrease X_t .

The inter-temporal price condition (11) does not hold in the case of a stockout, i.e. the case where $X_t = 0$ because the storage non-negativity constraint is binding. In this case arbitrageurs expect the future price of oil to be sufficiently lower than the current price that they would sell any amount of oil they had, except that they have nothing left to sell; every barrel of extracted oil is being used for consumption. As a result, current price is above its unconstrained level:

$$P_t > \beta E_t[P_{t+1}] - C. \quad (12)$$

4.5 The Rational Expectations Equilibrium

The canonical commodity storage model is a rational expectations model with one state variable - availability of oil A_t - and one choice variables

²⁹The cost of storing a barrel of oil have most likely decreased over time. We ignore this for simplicity, since accounting for a downward trend in storage cost cannot explain the observed changes in price persistence or volatility.

- storage of oil X_t . A solution of the model - the rational expectations equilibrium - consists of a storage rule, which specifies the level of storage for every possible value of the state variable. Determination of price and consumption follows immediately from this rule³⁰. In our extended version of the model the rule retains its salient characteristics, well known from the literature (see below). But in the extended version, similarly to the AR(1) case considered by Chambers and Bailey (1996), storage is also the function of one (or two) exogenous variables, depending on assumptions regarding the income process. Relative income Y_t/\bar{Y}_t - how far above or below its mean is the current level of income - serves as the second state variable of the model when we assume that incomes follows a stable trend. For the case where income is subject to growth shocks, we need a third state variable: the current growth rate of income, denoted by μ_t .

In order to solve the model and arrive at the correct storage rule, we express all quantity variables in their normalized forms. The model can be then be summarized by two (or three) transition functions (for the state variables) and one response equations (for the decision variable). Agents in the model observe all the state variables every period, and decide on storage accordingly, taking into consideration expectations regarding the next period's price.

The transition functions for the stable trend case are:

$$a_{t+1} = \frac{x_t + z_{t+1}}{(Y_t/\bar{Y}_t)^{\rho-1} e^{\bar{\mu} + \varepsilon_{t+1}}}, \quad (13)$$

$$\frac{Y_{t+1}}{\bar{Y}_{t+1}} = \left(\frac{Y_t}{\bar{Y}_t} \right)^\rho e^{\varepsilon_{t+1}}, \quad (14)$$

³⁰Williams and Wright (1991) show that one can add an intended production rule to the model, whereby suppliers decide one period in advance, depending on their price expectations, on their desired production. Appendix B of Dvir and Rogoff (2009) shows how this extension can be included in our model. However adding this second choice variable is computationally cumbersome and moreover does not add much to the insights of the model.

where equation (13) is derived by normalizing equation (1) by Y_{t+1} and using (6). Effective supply z_{t+1} is arrived at by dividing equation (10) through by Y_t .

For the stochastic trend case, there are three transition functions:

$$a_{t+1} = (x_t + z_{t+1})/e^{\mu_{t+1}}, \quad (15)$$

$$\frac{Y_{t+1}}{\bar{Y}_{t+1}} = e^{\mu_{t+1} - \bar{\mu}} \frac{Y_t}{\bar{Y}_t}, \quad (16)$$

$$\mu_{t+1} = (1 - \varphi)\bar{\mu} + \varphi\mu_t + v_t, \quad (17)$$

where the transition function (15) is derived by normalizing equation (1) by Y_{t+1} and using (7) instead.

The response equation for both cases is:

$$(a_t - x_t)^{-\gamma} = \beta E_t[P_{t+1}] - C. \quad (18)$$

Note importantly that equation (18), which determines optimal storage, holds only when the state variables are such that the optimal storage is non-negative. If the state variables dictate negative storage, this response condition breaks down and we have simply $P_t = a_t^{-\gamma}$.

The existence and uniqueness (under certain general conditions) of the rational expectations equilibrium, as well as its important properties, have been proven in the literature. In particular, Chambers and Bailey (1996) prove these properties for the case of auto-correlated supply shocks. However, commodity storage models generally cannot be solved analytically even in their most simple form, due to the non-negativity constraint. We therefore follow the literature since Gustafson's (1958) original contribution and proceed to solve the model numerically. This can be done using a variety of methods³¹. For computational reasons, we choose to use the spline collocation method

³¹Williams and Wright (1991), Chap. 3, survey the numeric methods applied to commodity storage models in the literature.

(see Judd [1998], Miranda and Fackler [2002] for a discussion)³².

The storage rules in our extended model are identical in form to the ones that result from the canonical model. The difference is that in the extended model these rules hold for the *normalized* variables instead of the original quantities. In other words, effective storage has a relationship with effective availability in the extended model, under both sets of assumptions regarding demand, that is qualitatively similar to the relationship between actual storage and actual availability in the canonical model.

Figure 3 exhibits the optimal storage rule as well as the corresponding equilibrium price, both as functions of effective oil availability a_t (on the horizontal axis), with the other state variable(s) held constant³³. The figure is qualitatively similar regardless of our assumption on income's stochastic process.

Figure 3 Here.

In the figure, points on the curves that correspond to a particular level of effective availability represent the rational expectations equilibrium - effective storage and the resulting equilibrium price - that would prevail if effective oil availability were indeed at that level. As illustrated in the figure, effective storage is zero when the effective amount of oil available is low, then after a kink at \bar{a} it rises monotonically³⁴. The marginal propensity to store is always less than one; that is because a rise in storage must lower the expected future

³²Details can be found in Appendix C of Dvir and Rogoff (2009).

³³Certain assumptions need to be made regarding the model's parameters in order to solve the model numerically. Demand elasticity $-1/\gamma$ is set at -0.2. The cost of storage C is 0.02 per barrel. The discount factor β is set at 0.97. Effective supply capacity \tilde{Z} is set at 1.3. The trend income growth rate $\bar{\mu}$ is set at 0.02, while the persistence parameter ρ is set at 0.5. Lastly, the income shock's standard deviation σ is set at 0.01.

³⁴See Deaton and Laroque (1992), Theorem 1. When availability is relatively low (oil is temporarily scarce), agents will sell off all existing inventories of oil while the equilibrium price is high, expecting it to fall in the following period. Storage will be therefore zero, and indeed would have been negative had it been possible - agents would want to hold a short position.

price, as it raises future availability of oil. The kink in the storage rule occurs when an additional barrel of stored oil will generate an expected profit of zero:

$$\bar{a}^{-\gamma} = \beta E[z^{-\gamma}] - C, \quad (19)$$

where z , recall, is normalized production (which is a function of current and trend income; see below).

Importantly, the kink at \bar{a} is also seen in the equilibrium price function p : when oil is relatively abundant, i.e. $a > \bar{a}$, the price function is more elastic. That is because once storage kicks in, a rise in oil availability causes a less than proportionate rise in the amount available for consumption, since there is now competing demand from arbitrageurs who are keen to increase their stocks. As a result, equilibrium price is less sensitive to changes in availability when storage is positive relative to a stockout situation where this competing demand is non-existent.

As shown by Chambers and Bailey (1996), the rational expectations equilibrium is generally dependent not only on current availability, but also on the current shock. Figure 4 therefore examines how the storage rule and the equilibrium price change with current income, given by Y_t/\bar{Y}_t , when income is subject to shocks around a deterministic trend. This exercise is very similar to Chambers and Bailey, with similar results: since shocks exhibit positive auto-correlation, a higher than average income today leads to a higher expected future income, implying that at any level of effective availability, arbitrageurs will raise the optimal amount of storage in expectation of higher prices next period. This, in turn, leads to a higher equilibrium price today wherever storage is positive.

Figure 4 Here

When income is subject to growth shocks the model behaves in a very similar way: when current income is high relative to its trend, or when the current growth rate is relatively high, we would see the storage rule shifting

up and to the left, with the equilibrium price rising accordingly. Intuitively, the logic is straightforward: when income is subject to growth shocks rather than autocorrelated level shocks, it is still the case that a current income that is high relative to trend predicts relatively high income in the future. Therefore the storage rule and the equilibrium price should respond in a similar way. When the current growth rate is relatively high, this implies a higher-than-average growth rate in the future, given our assumption of autocorrelation in growth rate, transition equation 17. That again leads agents to expect higher income in the future, which works in the same direction. We see then that the introduction of growth shocks does not change the characteristics of equilibrium in the model. It will, however, change its dynamic behavior, an issue to which we turn next.

5 Dynamic Behavior of Storage and Price in the Extended Model

We can now examine the dynamic behavior of the model following shocks to the income process. We first examine the dynamic behavior of our extended model under the assumption that shocks to income are AR(1), and show that our extended model behaves very similarly to the model analyzed and estimated by Deaton and Laroque (1996). In particular, storage in these conditions serves its classic purpose, to mitigate shocks: arbitrageurs transfer stocks from times of plenty to times of want. Therefore storage cannot explain price volatility. We then proceed to show that the model's dynamic behavior is quite different when we assume that demand is subject to growth shocks. Storage will act to magnify shocks in the model, thereby increasing price volatility³⁵.

³⁵Kilian (2008c) argues that oil price movements that cannot be explained by either supply or industrial demand shocks should be thought of as shocks to precautionary demand. The endogenous storage response in our model is also separate from the direct effects of

5.1 AR(1) Shocks to Income

Consider first a positive and persistent shock to income, when the stochastic process is AR(1). The shock's effects on effective availability, equilibrium price, effective production, and effective storage are depicted in Figure 5. The figure exhibits the results of the following simulation: we let the system run for 30 periods, where each period a new value of ε_t is drawn from the appropriate distribution. We perform 100,000 repetitions of this simulation, with the figure showing mean values for each period. This produces the baseline case. We then repeat this exercise with one change, namely that in period 2 there occurs a three standard deviation positive shock to ε_t . The mean results of 100,000 repetitions of this simulation are shown in dashed lines.

The shock to income results in *effective* availability dropping sharply, leading therefore to an immediate rise in current equilibrium price, as a fixed amount of oil must satisfy a larger thirst for it. These twin effects on effective availability and price subside gradually over time, as the income shock dissipates, and the system returns to its steady state. Oil production is inelastic in our model, implying that effective oil supply will drop with the shock, only to recover slowly as the shock dissipates³⁶.

The shock's effects on storage are more complex. Arbitrageurs are caught between two contradictory forces following this type of shock: on the one hand, the rise in current prices and drop in current effective availability dictate a drop in optimal effective storage according to the storage rule (see

the shock, however it responds to the expected mean price, rather than to its expected volatility as is the case with precautionary demand.

³⁶If production were elastic, as in Williams and Wright (1991), producers would respond to these developments by increasing their planned production in expectation of higher prices. As seen in the analytically solvable no-storage case (see Appendix B of Dvir and Rogoff, 2009), the elasticity of planned production with respect to current income is less than one, due to the positive slope of producers' marginal cost curve. Effective production would respond therefore in a qualitatively similar way if production were not perfectly inelastic.

Figure 3). On the other hand, due to the shock's persistence future income is also expected to be higher than average, implying higher expected future prices and therefore an increased incentive to store at any level of effective availability (see Figure 4). However, the former effect must dominate the latter in the case of an AR(1) income process, as future income is expected to be *less* than current income, implying that future effective availability is expected to be higher than its current value, and accordingly that the future price is expected to be lower than its current, post-shock level. As a result, the effect of a positive shock to income would be to reduce storage until the system reverts back to its steady state.

Figure 5 Here.

We see in the system's response to a temporary and persistent demand shock the underlying reason for the disappointment expressed by Deaton and Laroque (1996) regarding the storage model's inability to account for the auto-correlation seen in commodity prices. When shocks are transitory, i.e. when the system is stationary, storage acts as a countervailing force: in Figure 5, when the equilibrium price is above its steady state level, storage is smaller than its own steady state level, in a partial compensation for the shock. This in itself does contribute to a higher equilibrium price in the next period, as observed in the data. However, persistence of the shock only serves to reduce the magnitude of this response, since the connection between current and future conditions formed by the shock's persistence substitutes in part for the inter-temporal connection that is due to arbitrage. Therefore an AR(1) shock does not deliver the added persistence that Deaton and Laroque (1996) are looking for.

5.2 Growth Shocks to Income

Our extended model allows us, as we have seen, to incorporate growth shocks into the storage framework. In this case, storage does not act as a counter-

vailing force anymore; indeed, immediately following the shock storage tends to *magnify* the shock's effect on equilibrium price. Figure 6 demonstrates the effects of a positive and persistent shock to income growth, in this case a three standard deviation positive shock to ν_t . As in the AR(1) case, a positive demand shock lowers effective availability and raises the equilibrium price. However, in this case the shock brings about a transition to a new steady state in which effective availability is expected to be at a lower level permanently, accompanied by a permanently higher price level, and a permanently lower effective production level. Importantly, due to positive auto-correlation in the stochastic process, this transition is spread over several periods. This provides a role for storage. As in the AR(1) case, arbitrageurs are subject to contradictory forces: the current rise in price induces a corresponding drop in storage, while the prospect of higher prices in the future induces a storage increase. However, the crucial difference between the two cases is that here, due to the shock's persistence, equilibrium price in the future is expected to *increase* relative to the current, post-shock price. As a result, in the stochastic trend case the storage-increasing effect of future prices is stronger than the storage-decreasing effect of the current price. Storage in the transition period is therefore higher in expectation relative to the expected path it would follow had the shock not occurred. In the stochastic trend case, the shock's persistence magnifies the storage response instead of diluting it as in the stable trend case³⁷. Note that if growth shocks were iid, there would be almost no storage response: the price determined by the current period shock would be expected to persist in the future, so there would be no incentive to change the amount of optimal storage chosen before agents gained knowledge of the shock.

Figure 6 Here.

³⁷The same logic applies to the opposite case, where income growth suffers a negative shock, but with a caveat. Storage response, in this case a decrease, is stronger the more persistent the growth shock. However, since storage cannot be non-negative, this effect is bounded in the negative growth shock case.

It is revealing to compare our simulated response to the case where storage is not allowed. This is done in Figure 7, where the post-shock expected paths of the endogenous variables are compared with the paths these variables would take if storage were not possible. Due to the shock to income growth, storage spikes sharply upwards, leading to a slightly *higher* equilibrium price relative to the no-storage case. Effective availability in the periods of transition to the new steady state is high relative to the no-storage case, since storage remains positive throughout the transition (in expectation of higher prices in the future). Therefore we see a slower expected convergence to steady state price relative to the no-storage case.

Figure 7 Here.

We see then that in the presence of growth shocks, a rise in current price may be associated with an increase in optimal storage, rather than a decrease as would always be the case in the canonical model. Storage in this case does not "lean against the wind", as is its customary role; it actually magnifies the shock somewhat, by increasing demand exactly when it is already high, in preparation for even higher demand in the future. This behavior could act to increase price volatility above and beyond what it would otherwise be without storage. In fact, when we simulate the model with and without storage, we find that price volatility following a growth shock is indeed higher when storage is present. Figure 8 shows the results of these simulations. As before, we let the model run for 30 periods, simulating the evolution of prices and quantities after a positive growth shock to income. We then repeat the process 100,000 times. Each time the model is run, we regress the growth rate of price on its lag, and use the absolute value of the residuals as our measure of volatility (the same measure that we used in testing for differences in volatility of real oil price data in Section 2). Figure 8 presents the mean value of this measure, per period, for the case where storage is allowed, and for the case where it is constrained at zero. We see that the initial positive

growth shock results in much higher price volatility - roughly double with the parameters assumed here - when storage is allowed, and that the effect lasts for several periods until price volatility converges back to its long run mean³⁸.

Figure 8 Here.

In our extended commodity storage model with growth shocks, we have arrived at a result that accords well with our empirical findings: periods in which persistent growth shocks are dominant should be periods in which price exhibits extra volatility, relative to periods in which AR(1) shocks are more prevalent. In quantitative terms, our simulations show that storage alone can double, on average, the volatility of prices immediately after a positive growth shock. This difference is of a similar order of magnitude to the differences in price volatility that we observe in the data. This tells us that the mechanism which we emphasize in this paper can go a long way towards accounting for the differences across periods in the complete time series of real oil price.

6 Conclusion

We argue in this paper that a long-term view is essential to understanding the dynamic behavior of oil prices. We show that shocks to the oil market can have remarkably different effects on the real price of oil across historical periods, not because of their origin on the supply or the demand side, but rather due of the ability (or lack thereof) of key players in the market to restrict access to supplies. With effective restrictions on access to excess supplies, growth shocks can generate oil prices that are both highly persistent and, through an endogenous storage response, highly volatile. On the other

³⁸Recall that our persistence parameter here is set at $\rho = 0.5$. Our qualitative results stand with higher or lower persistence.

hand, without these restrictions, the same growth shocks will be quickly accommodated, and will not lead to increased persistence or volatility. In this regard, it is immaterial whether the growth shocks originate on the demand or the supply side.

The literature, for obvious reasons, has focused on the extremely persistent and volatile post-1973 period. Given the geo-political realities, this period is unlikely to end in the foreseeable future. However, it is useful to recognize that throughout most of the history of oil, the ability to restrict access to supplies was actually sorely lacking, with the oil market showing remarkable flexibility and relative price stability as a result. This held true even in years when oil supply or demand were experiencing great upheavals, such as during World War II and the postwar re-building of Europe. Although rare, the history of oil shows that shifts in industry structure do occur, and the structural breaks in price behavior associated with these shifts are testimony to their importance.

References

- [1] Almoguera, Pedro A., and Herrera, Ana María (2007), "Testing for the Cartel in OPEC: Noncooperative Collusion or Just Noncooperative?" mimeo Wayne State University.
- [2] Alquist, Ron, and Kilian, Lutz (2008), "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Economics*, forthcoming.
- [3] Bai, Jushan and Pierre Perron (1998), "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica* 66, pp. 47-78.
- [4] Bai, Jushan and Pierre Perron (2003), "Computation and Analysis of Multiple Structural Change Models." *Journal of Applied Econometrics* 18, pp. 1-22.

- [5] Bairoch, Paul (1982), "International Industrialization Levels from 1750 to 1980." *Journal of European Economic History* 11, pp. 269-310.
- [6] Bentley, Jerome T. (1979), *The Effects of Standard Oil's Vertical Integration into Transportation on the Structure and Performance of the American Petroleum Industry, 1872-1884*. New York: Arno Press.
- [7] Barsky, Robert B. and Kilian, Lutz (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," *NBER Macroeconomics Annual* 2001, pp. 137–83.
- [8] Barsky, Robert B. and Kilian, Lutz (2004), "Oil and the Macroeconomy Since the 1970s." *Journal of Economic Perspectives* 18, pp. 115-134.
- [9] Busetti, Fabio and Taylor, A.M. Robert (2004), "Tests of Stationarity Against a Change in Persistence." *Journal of Econometrics* 123, pp. 33-66.
- [10] Caner, Mehmet and Kilian, Lutz (2001), "Size Distortions of Tests of the Null Hypothesis of Stationarity: Evidence and Implications for the PPP Debate," *Journal of International Money and Finance* 20, pp. 639-657.
- [11] Cavaliere, Giuseppe and Taylor, A.M. Robert (2008), "Testing for a Change in Persistence in the Presence of Non-stationary Volatility," *Journal of Econometrics* 147, pp. 84-98.
- [12] Deaton, Angus and Laroque, Guy (1992), "Behaviour of Commodity Prices." *Review of Economic Studies* 59(1), pp. 1-23.
- [13] Deaton, Angus and Laroque, Guy (1996), "Competitive Storage and Commodity Price Dynamics." *Journal of Political Economy* 104(5), pp. 896-923.
- [14] Dvir, Eyal and Rogoff, Kenneth S. (2009), "The Three Epochs of Oil," *NBER Working Paper No. 14927*.

- [15] Hamilton, James D. (2008a), "Oil and the Macroeconomy," in Durlauf, Steven N. and Blume, Lawrence E., *The New Palgrave Dictionary of Economics, 2nd Edition*. New York: Palgrave Macmillan.
- [16] Hamilton, James D. (2008b), "Understanding Crude Oil Prices." NBER Working Paper No. 14492.
- [17] Hansen, Bruce E. (2000), "Testing for Structural Change in Conditional Models." *Journal of Econometrics* 97, pp. 93-115.
- [18] Harvey, David I., Leybourne, Stephen J. and Taylor, A.M. Robert (2006), "Modified Tests for a Change in Persistence." *Journal of Econometrics* 134: 441-469.
- [19] Hotelling, Harold (1931), "The Economics of Exhaustible Resources." *Journal of Political Economy* 39, pp. 137-175.
- [20] International Energy Agency (2006, 2007), "Oil Information". Paris: OECD.
- [21] International Monetary Fund (2008), "World economic Outlook." Washington, DC.
- [22] Gustafson, Robert L. (1958), "Carryover Levels for Grains: A Method for Determining Amounts That Are Optimal under Specified Conditions." *USDA Technical Bulletin* 1178.
- [23] Judd, Kenneth L. (1998), *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- [24] Kilian, Lutz (2008a), "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature* 46, pp. 871-909.
- [25] Kilian, Lutz (2008b), "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics* 90, pp. 216-240.

- [26] Kilian, Lutz (2008c), "Not All Oil Price Shocks Are Alike: Disentangling demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, forthcoming.
- [27] Kim, Jae-Young (2000), "Detection of Change in Persistence of a Linear Time Series." *Journal of Econometrics* 95, pp. 97-116.
- [28] Kim, Jae-Young, Jorge Belaire-Franch, and Rosa Badillo Amador (2002), "Corrigendum", *Journal of Econometrics* 109, pp. 389-392.
- [29] Krautkraemer, Jeffrey A. (1998), "Nonrenewable Resource Scarcity," *Journal of Economic Literature* 36, pp. 2065-2107.
- [30] Kwiatkowski, Denis, Peter C.B. Phillips, Peter Schmidt, and Yongcheol Shin (1992), "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root." *Journal of Econometrics* 54, pp. 159-178.
- [31] Maddison, Angus (1998), *Chinese Economic Performance in the Long Run*. Paris: OECD.
- [32] Maddison, Angus (2003), *The World Economy: Historical Statistics*. Paris: OECD.
- [33] Miranda, Mario J. and Fackler, Paul L. (2002), *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press.
- [34] Ng, Serena, and Ruge-Murcia, Francisco J. (2000), "Explaining the Persistence of Commodity Prices," *Computational Economics* 16, pp. 149-171.
- [35] Pindyck, Robert S. (1999), "The Long-Run Evolution of Energy Prices." *The Energy Journal* 20(2), pp. 1-27.
- [36] Routledge, Bryan R., Seppi, Duane J., and Spatt, Chester S. (2000), "Equilibrium Forward Curves for Commodities." *Journal of Finance* 55(3), pp. 1297-1338.

- [37] Smith, James L. "Inscrutable OPEC? Behavioral Tests of the Cartel Hypothesis," *The Energy Journal* 26(1), pp. 51-82.
- [38] Smith, James L. "Organization of the Petroleum Exporting Countries (OPEC)," in Durlauf, Steven N. and Blume, Lawrence E., *The New Palgrave Dictionary of Economics, 2nd Edition*. New York: Palgrave Macmillan.
- [39] Schwartz, Eduardo, and Smith, James E. (2000), "Short-Term Variations and Long-Term Dynamics in Commodity Prices." *Management Science* 46(7), pp. 893-911.
- [40] Williams, Jeffrey C. and Wright, Brian D. (1991), *Storage and Commodity Markets*. Cambridge, England: Cambridge University Press.
- [41] Williamson, Harold F. and Daum, Arnold R. (1959), *The American Petroleum Industry, Vol. I*. Evanston, IL: Northwestern University Press.
- [42] Wright, Brian D. (2001), "Storage and Price Stabilization," in Gardner, Bruce L. and Rausser, Gordon C. (eds.) *Handbook of Agricultural Economics*, pp. 817-862.
- [43] Wright, Gavin (1990), "The Origins of American Industrial Success, 1879-1940." *American Economic Review* 80(4), pp. 651-668.
- [44] Yergin, Daniel (1991), *The Prize*. New York: Simon & Schuster.

Table 1: Sample Statistics of Oil Price Series

	Subsamples			Full Sample
	1861-1877	1878-1972	1973-2008	1861-2008
Price (2007 \$US)				
Mean	50.9	17.2	44.3	27.6
Std. Deviation	25.3	5.1	22.3	20.1
Annual price change (abs.)				
Mean	39.0%	14.2%	21.5%	18.7%
Std. Deviation	27.5%	13.7%	22.7%	19.5%

Table 2: Testing for Change in Persistence of Oil Price

Real oil prices are annual averages. The full sample 1861-2008 was split in order to run the tests, which do not accommodate multiple breaks. Tests are adjusted following Harvey et al. (2006) so that they are appropriately sized whether $H_0 : I(0)$ or $H_0 : I(1)$. The table presents three pairs of test statistics, with their respective bootstrap p-values in parentheses. A wild bootstrap with 1,000 repetitions was used to make inference robust to heteroskedasticity, following Cavaliere and Taylor (2008). Only statistically significant change-points are listed.

Sample	1861-1955	1881-2008
Change-point (direction)	1877 ($I(1)$ to $I(0)$)	1972 ($I(0)$ to $I(1)$)
95% CI	(1873, 1881)	(1967, 1977)

Mean Score Statistics:

MS	0.03 (0.962)	35.60** (0.021)
MS^R	125.52*** (0.002)	0.48 (0.266)

Mean-Exponential Statistics:

ME	0.01 (0.978)	146.36*** (<0.001)
ME^R	261.75*** (<0.001)	0.61* (0.085)

Maximum Statistics:

MX	0.21 (0.859)	320.61*** (<0.001)
MX^R	559.82*** (<0.001)	6.44* (0.067)

Table 3: Testing for Change in Volatility of Oil Price

Real oil prices are annual averages. Volatility is defined as the absolute value of the error terms from regressing the growth rate of real oil price on one lag. The table presents results from a test for multiple breakpoints in oil price volatility, following Bai and Perron (2003). See text for the sequential procedure used to select the number of breakpoints.

Estimated breakpoints, confidence intervals, and period mean volatilities:

	Breakpoint	95% CI	Mean Volatility (t-statistic)
Period I	1878	1873-1895	0.381 (9.4)
Period II	1934	1931-1946	0.195 (9.1)
Period III	1972	1949-1973	0.061 (2.3)
Period IV			0.221 (8.2)

The Bai-Perron test statistics:

Test Statistic	Value
$\text{supF}_T(1 \text{ break vs. no breaks})$	22.70***
$\text{supF}_T(2 \text{ breaks vs. no breaks})$	14.61***
$\text{supF}_T(3 \text{ breaks vs. no breaks})$	15.97***
Sequential Procedure:	
$\text{supF}_T(2 \text{ breaks vs. 1 break})$	6.39
$\text{supF}_T(3 \text{ breaks vs. 2 breaks})$	32.07***
$\text{supF}_T(4 \text{ breaks vs. 3 breaks})$	3.30

Figure 1: The Real Price of Oil 1861–2008






Figure 2: Real Oil Price Volatility, 1861–2008



**Figure 3: The Rational Expectations Equilibrium
Storage Rule**



Equilibrium Price




Figure 4: Effect of Increase in Relative Income on RE Equilibrium Storage Rule





Figure 5a: Response of Effective Availability to Positive Income Shock




Figure 5b: Price Response to Positive Income Shock




Figure 5c: Effective Storage Response to Positive Income Shock




Figure 5d: Effective Production Response to Positive Income Shock




Figure 6a: Response of Effective Availability to Positive Growth Shock




Figure 6b: Price Response to Positive Growth Shock




Figure 6c: Effective Storage Response to Positive Growth Shock




Figure 6d: Effective Production Response to Positive Growth Shock




Figure 7a: Response of Effective Availability to Positive Growth Shock




Figure 7b: Price Response to Positive Growth Shock




Figure 7c: Effective Storage Response to Positive Growth Shock




Figure 7d: Effective Production Response to Positive Growth Shock




Figure 8: Price Volatility Following a Positive Growth Shock

