Babylonian Risk Aversion*

Juliusz Radwański and Alex Stomper
Humboldt University, Berlin

July 2019

Abstract

The effects of risk on market prices depend on the risk aversion of market participants. We estimate this risk aversion based on crop prices recorded more than 2000 years ago in Seleucid Babylon. In these data, we find evidence that market participants were risk averse, revealing a degree of risk aversion comparable to today. The evidence suggests that generations in the far future will also price risk as we do in our markets.

Keywords: Relative risk aversion, crop markets, Seleucid Babylon.

JEL Classification Numbers: G12, N25, Q02.

1 Introduction

How stable is human risk aversion in the long run? The answer to this question matters for our policy-making with respect to risks affecting generations in the far future, e.g., risks of extreme weather events associated with climate change.[1] How will future generations assess such risks? Will they be more risk averse than we are? On which level of risk aversion should we base our policies?

*Correspondence to: stompera@hu-berlin.de. We thank Lawrence E. Blume, Max Bruche, Hermann Elendner, Mathieu Ossendrijver, Reinhard Pirngruber, Jacob S. Sagi, and Nikolaus Wolf for helpful comments.

[1] See Brock and Hansen 2018 for a discussion of various types of uncertainty associated with climate change.

This paper provides evidence that can be used to assess the stability of human risk aversion in the long run. Here, we estimate the risk aversion of crop traders in Seleucid Babylon. The estimates are surprisingly precise and are similar to estimates based on modern data, e.g., those reported in Myers 1989 which also concern crop market participants. This suggests that human risk aversion is rather stable over long periods of time, adding to evidence that it has genetic and neurophysiological foundations (Bell 2009; Christopoulos et al. 2009). Our evidence results from data recorded more than 2000 years ago. We know of no other estimates of human risk aversion based on these or older data.

The data behind our analysis are prices of ancient Babylon's main crops, i.e., barley and dates. The fluctuations of these prices determined the payoffs of crop storage in Babylon. We focus on short-run storage (10-60 days) and measure the payoffs in terms of returns. Economic theory (e.g., Merton 1973) predicts that the returns from storing dates and barley are – in expectation – proportional to the covariances of these returns with aggregate consumption risk. If storing a particular type of crop yields returns less correlated with aggregate consumption risk, then this type of crop storage is better suited as a form of insurance against this risk, and the expected returns can be smaller.

We compare returns to storing dates and barley. If these returns feature different correlations with aggregate consumption risk, then an increase in this risk should increase the difference between the expected values of the returns. The strength of this effect depends on the risk aversion of the representative market participant. We identify this risk aversion based on variation in risk associated with the Babylonian harvest cycle, and a standard model.

Our ancient data yield estimates of relative risk aversion (RRA) in the range 2-7 (90% confidence), consistent with estimates based on modern crop price data reported in Myers 1989, and also with the RRA-estimates recently reported in Calvet et al. 2019.[2] It thus appears that

[2] An intuitive definition of a person's RRA describes it as the number of percentage points by which the person's

human risk aversion is rather stable in the long run. This matters for the economic analysis of policies concerning the risk exposure of generations in the far future. The analyses are commonly based on models that specify the preferences of different generations in an identical manner, even if they are hundreds of years apart.[3] We provide empirical support for this practice with respect to RRA, i.e., a preference parameter that commonly appears in analyses of the social benefit of climate change mitigation, e.g., Giglio et al. 2018; Lemoine 2017. Our evidence suggests that the social cost of carbon can be measured based on data about suitable present-day markets (Giglio et al. 2018).

2 Background

This paper is based on price data from an ancient society which obtained most (70-75 %, see Aperghis 2004) of its caloric intake by consuming barley and dates. Prices of these crops affected the livelihoods of all people in this society, many of whom received crop rations as "salaries paid in kind" (Jursa 2008). The prices were recorded by Babylonian scholars in the "Astronomical Diaries", first analysed in Hunger and Sachs 1988, 1989, 1996. We draw on van der Spek's edition of the price data, available online (van der Spek 2005). The sample period is 385 to 61 BCE. On his website, van der Spek also provides extensive information regarding the Babylonian economy at the time.

Figure 1 plots the price observations behind our analysis. As discussed below, the analysis is based on all available pairs of prices recorded at most 60 days apart. We plot these observations in terms of shekels per *kurru* (180 litres), where a shekel corresponds to about 8.33 grams of

wealth could be reduced so that this person is (approximately) as worse-off as if subjected to the risk of a random percentage change in wealth with zero mean and a standard deviation of about 14%. See our Online Appendix A.

[3] For a survey, see Gollier 2013. There, the assumption appears on p. 29 in terms of a specification of the preferences of generation t's consumers in terms of the function u that also describes those of generation 0 in expression (2.2).

silver. The plot's shaded areas mark periods of military operations. We identify these periods based on information in a book by Pirngruber 2017, which also reports evidence that military operations were typically associated with higher prices.[4] Temin 2002 confirms and extends arguments of Slotsky 1997; Grainger 1999; van der Spek 2000 that the prices resulted from market transactions.

To the best of our knowledge, this paper is the first to measure the risk aversion of market participants in ancient Babylon. We make efficient use of scarce data to measure this risk aversion in terms of a single parameter, i.e., a coefficient of relative risk aversion of a representative market participant. A previous analysis by Foldvari and van Leeuwen 2014 uses a model with two parameters related to risk aversion, but stops short of estimating the parameters.

A literature in experimental economics suggests that estimates of risk aversion are context-dependent and that they tend to rise when the stakes increase.[5] By estimating the risk aversion of a representative market participant, we actually analyse the pricing of risk in a market, rather than the risk attitudes of individual market participants. We aim at results that can be compared to estimates based on modern price data since they result from a standard model and a market that existed throughout human history, i.e., the crop market. If the context of crop trading in antique markets differed from that of such trading today (e.g., with respect to the risk of political interference in crop storage), we may obtain estimates that differ from estimates based on modern crop prices. A key issue is certainly the extent of market participation. To interpret our results as measures of risk aversion of a representative agent in the Babylonian economy, we must trust that there were well-operating markets for silver, barley and dates, with a sufficiently large degree of direct market participation and/or sufficiently efficient risk-sharing between market participants and non-participants.[6]

[4] See Table 6.6. in Pirngruber 2017.

[5] For a survey, see Schildberg-Hörisch 2018.

[6] The existence of the system of crop rations suggests that our measures of risk aversion concern not only

With its evidence based on price data recorded 2000 years ago, our paper complements papers regarding shorter-run fluctuations in market participants' risk appetite, e.g., Malmendier and Nagel 2011. We test a hypothesis which appears in a literature on evolutionary "market selection" processes concerning the survival and price impact of market participants, e.g., Kogan et al. 2017 and Yan 2008, following a seminal contribution by Blume and Easley 1992. The hypothesis (discussed, e.g., in Lengwiler 2004, p. 90) states that, in the long run, market price formation should be governed by traders whose level of risk aversion is consistent with maximizing expected logarithmic wealth. We test this hypothesis based on the assumption that, at any point in time, the market participants know the distribution of future returns from crop storage. Our results are, however, consistent with values of relative risk aversion that appear in papers relaxing the assumption, e.g., Bossaerts 2004. With respect to harvest risk, the assumption may be less of a problem because this risk should have been well-understood by market participants in ancient Babylon.

3 Research Strategy

Our analysis is based on the idea that barley and date storage were alternative risk-management strategies in ancient Babylon. Our focus is on storage motivated by consumption needs, rather than by a desire to manipulate market prices through the creation of artificial shortages. This is consistent with evidence in Jursa 2010 that there was no hoarding of crops by big producers (i.e., the temples). Politically motivated price manipulation is an issue that will be addressed in robustness checks by changing our sample to exclude observations during years in which political intervention in crop storage/trading was more likely than in other years.

the market participants in Babylon's crop market, but also non-market participants who worked for the market participants and received the rations as as "salaries paid in kind" (Jursa 2008). This idea has been formalized in a literature on risk-sharing in labor markets. See Berk and Walden 2013.

Throughout our analysis below, we measure prices in units of barley, setting the price of barley to one. We thus focus on date storage as a strategy with a risky payoff depending on a future price in terms of litres of barley. This price would have to be paid for dates had they not been stored. It is the relative price of dates and barley, implied by their prices in units of silver. Its growth rate measures the profitability of storing dates, rather than barley.

Our estimates result from a condition (stated in expression (1) below) that is a criterion for the optimal extent of date storage. We use this condition to measure the risk aversion of the representative market participant in ancient Babylon, based on the notion that harvest risk was a key determinant of Babylonian crop storage strategies. In essence, our research strategy is built on an (a-priori) argument that the barley growing period (BGP) was a high-risk period in ancient Babylon. By using this argument, we avoid the need to identify variation in risk solely based on the data, which would reduce the power of our analysis.

4 Exploratory Analysis

We now show that the BGP was a time of high risk in ancient Babylonia. This is plausible because, during this period, the Babylonian economy received unanticipated information ("news") about its most important harvest's yield and timing. Bad/good news should have caused a drop/rise in aggregate consumption by raising/reducing the optimal extent of crop storage in the economy.

We also explore whether the date-growing period (DGP) was a high-risk period. Risk will be measured in terms of silver price fluctuations. This is based on the idea that, by changing aggregate consumption, news about harvests also changed the demand for silver as money used in consumption-induced transactions. We essentially assume that, during short intervals of time, the supply of silver was reasonably constant and that price changes can be attributed to demand shocks. In fact, most of our data comes from a particularly stable period in Babylon's history, i.e., the Seleucid period in which inflation was largely absent.[7]

In the data, we observe silver prices in units of barley, recorded by Babylonian scholars at intervals of (about) 10 days.[8] We first compute "raw" price growth rates as differences between logarithmic prices. Given that most records were lost, we can however only obtain a sufficiently big sample by including some growth rates based on prices observed more than 10 days apart. To compare the latter growth rates to the 10-day growth rates, we divide the raw growth rates by the number of 10-day periods between the underlying prices. Whenever this number exceeds one, we obtain averages of unobserved price growth rates during (more than one) 10-day periods. In using these averages to measure risk, we face the problem that some unobserved fluctuations in price growth become averaged-out. We address this problem in two ways. First, we only consider growth rates based on prices observed up to 60 days apart. Second, we re-scale the average growth rates so that they can be used to measure the average standard deviation of any unobserved 10-day growth rates that were averaged. See the Online Appendix B. The re-scaled growth rates appear in variance ratio tests, while the basic average growth rates are used to measure means and quantiles of price growth.

The 60-day limit yields samples of 73 growth rates based on prices of silver and dates observed at the same points in time. Descriptive statistics appear in Table 1, where we also test for differences between various samples.[9] In the remainder, we use the sample described

[7]See Figure 1. An inflationary shock did occur after the death of Alexander the Great, which triggered the war of the successors. This is discussed, for example, in Goetzmann 2016, pp. 69-70. The shock cannot affect our results because our estimates are based on a sample that excludes observations during periods of military operations in ancient Babylonia.

[8]The Babylonian scholars recorded quantities of barley that could be bought for one shekel of silver. van der Spek refers to the recorded values as "exchange values" (van der Spek 2000). We can directly use the barley exchange values since we measure prices in units of barley. In the archaeological literature, it is now common to compute prices in terms of shekels.

[9]In his edition of the price data recorded in the Babylonian Astronomical Diaries, van der Spek points out that some observations are based on doubtful records. Moreover, it is not always clear on which day of a month prices were recorded. We therefore define samples excluding observations based on doubtful records or observations with

in Panel A of Table 1. In terms of logarithmic prices, this sample appears in Figure 2. It features a covariance of silver and date price growth of 0.027 (measured using the "re-scaled" growth rates). As proxies for aggregate consumption, the silver prices are highly variable. While most of the variation is across years, we also observe substantial within-year variation. To some extent, this can be seen as a consequence of a lack of inter-regional trade in staple goods (Graslin-Thomé 2009), combined with low levels of inter-annual crop storage (Pirngruber 2017). With these features, the Babylonian economy should have been set up for large price fluctuations due to news about harvest yields, and it received such news more often than an economy with a single harvest period because its main harvests occurred about six months apart.

We next analyze the relative extent of silver price fluctuations during different parts of the Babylonian harvest cycle. Figure 3 presents box plots of the silver price growth rates based on all observations that fully fall inside/outside the DGP (Panel A) or the BGP (Panel B), where the DGP is defined as the three (Gregorian) months 8-10 while the BGP contains the three months 3-5.[10] The box plots in Panel A show that the DGP was similar to the rest of the year in terms of the interquartile range of silver price growth, depicted by shaded boxes. However, the same measure of risk reveals a substantial difference between the BGP and the rest of the year, as shown in Panel B. It appears that the BGP was a high-risk period with a particularly large interquartile range of silver price growth. A similar picture emerges in variance ratio tests (reported in the Online Appendix D), and also in terms of date price fluctuations. Figure 4 contrasts the BGP and the rest of the year by showing box plots of date price growth rates, again based on all observations that fully fall inside/outside the BGP. The BGP appears to be a high risk period, not only in terms of silver price fluctuations, but also in the market for dates.

imprecise within-month timing. For details, see the Online Appendices B and C.

[10] For further information, see van der Spek's commentary (van der Spek 2005). For example, he writes that barley was harvested mainly April to May. We add the month of March since it takes some time for barley to grow. Other specifications of the BGP and the DGP appear in robustness checks.

5 Main Analysis

We measure the risk aversion of the representative market participant (RMP) in ancient Babylon using the following condition, discussed below and in the Appendix:

$$E[(e^{\Delta p_{t,t+\tau}} - e^{\lambda\tau})e^{m_{t,t+\tau}}|I_t] = 0. \quad (1)$$

The condition requires a zero expected value of the function inside the expectations operator $E[\cdot]$, given the set of information available to market participants at any particular time t, denoted as I_t. The function is the product of two random variables. The first measures the profit (or loss) from storing dates between time t and a future time $t + \tau$. It is an "excess" return to date storage, i.e., the difference between a return due to date price growth, $e^{\Delta p_{t,t+\tau}}$, and a relative cost of storing dates, $e^{\lambda\tau}$. Given that we measure the date price in units of barley, the relative storage cost is the difference between the cost of storing dates and the cost of storing barley (both measured in terms of physical storage loss rates). The excess return to date storage is multiplied by a stochastic discount factor (SDF). This second random variable represents the way the RMP values the return to date storage, assigning higher values to profits in less desirable future scenarios.

Stated verbally, condition (1) requires that the excess return from storing additional dates has zero value. If so, the RMP would not want to increase or reduce the quantity of dates stored. This is a criterion for the extent of date storage being optimal, provided that it is non-zero.[11] This last requirement may not have been satisfied during periods of political instability in ancient Babylon. In fact, Figure 1 shows some exceptionally high crop prices (in units of silver) during periods of military operations after Alexander's death. During such periods, crop

[11] In the event of a stock-out, condition (1) would hold as an inequality. The effects of stock-outs on commodity prices are analyzed in Deaton and Laroque 1992, 1996 among others.

storage may have been affected by the risk of confiscation, and there may have been stock-outs. We therefore exclude periods of political instability from our empirical analysis, based on information in Pirngruber 2017 about military operations and instances of army presence in Babylonia, as well as wars abroad.[12]

To further explain our strategy for measuring the RMP's risk aversion, we consider the case of $\tau = 1$ and assume – counterfactually – that all random exponential values in condition (1) are lognormally distributed.[13] In this case, the condition can be stated as follows:

$$E_t[\Delta p_{t,t+1}] + \frac{1}{2}Var_t[\Delta p_{t,t+1}] - \lambda = -Cov_t[m_{t,t+1}, \Delta p_{t,t+1}], \quad (2)$$

where the expectation, variance, and covariance appear with subscripts indicating that we condition on the information set I_t. Under the assumption of conditional normality, the first two terms on the left-hand side add up to the expected arithmetic return to date storage. This expected return exceeds storage costs (λ) by the amount given by the right-hand side, i.e., a (possibly negative) risk-premium. The risk-premium is determined by the covariance of the return on date storage $\Delta p_{t,t+1}$ with the SDF as a variable which takes higher values in less desirable future scenarios. A positive value of the covariance therefore means that, in such scenarios, storing dates (rather than barley) tends to be especially profitable. In this way, the covariance measures the strength of market participants' insurance motive in date storage. If the covariance is negative, date storage is the opposite of insurance. In this case, market participants will only store dates if they can earn a positive risk premium.

In the Appendix, we specify the variable $m_{t,t+\tau}$ so that it increases when the silver price

[12] This information appears in Pirngruber 2017 in Tables 6.6. (military operations), 6.7. (army presence), and 6.3. (wars abroad).

[13] To test the assumption, we focus on the samples of growth rates based on prices observed no more than 10 days apart. All standard tests reject normality of the silver price growth rates, and we also find some evidence against normality of the date price growth rates. Moreover, the Doornik-Hansen test rejects the hypothesis of multivariate normality.

drops, indicating an increasing scarcity of barley (because we measure prices in units of barley). We use two textbook specifications of $m_{t,t+\tau}$, one of which is a restricted version of the other. Given this special case (stated in expression (13) in the Appendix), the above-stated condition can be written as follows:

$$E_t[\Delta p_{t,t+1}] + \frac{1}{2}Var_t[\Delta p_{t,t+1}] - \lambda = \gamma Cov_t[\Delta \pi_{t,t+1}, \Delta p_{t,t+1}], \tag{3}$$

where γ denotes the RMP's coefficient of relative risk aversion, and $\Delta\pi_{t+1}$ denotes the growth rate of the silver price. Our estimates actually reject the specification of $m_{t,t+\tau}$ behind the above-stated condition, but we still present it for expository purposes since it illustrates the effect of an increase in risk during the barley growing period (BGP) in a particularly simple way. Figure 5 contrasts the BGP (Panel A) and the rest of the year (Panel B) using plots of the expected growth rate of the date price, $E_t[\Delta p_{t,t+1}]$. Each plot features a dot showing estimates of this expected growth rate that result from condition (3) if we set γ and λ equal to our point estimates (reported below in row (1) of Panel A in Table 2), and estimate the remaining terms of the condition using maximum likelihood.[14] Given these estimates, the planes in the two plots show the expected date price growth implied by condition (3) for other values of γ and λ. The key take-away is that the value of γ determines the extent of the difference between the BGP and the rest of the year in terms of expected date price growth if the covariance $Cov_t[\Delta\pi_{t,t+1}, \Delta p_{t,t+1}]$ increases during the BGP (as a high-risk period). Our empirical analysis will yield estimates of γ and λ that are consistent with the variation in the average returns of date storage actually observed in our data.

Our null hypothesis is a unit value of the RMP's relative risk aversion, γ. This hypothesis follows from an evolutionary argument that, in the long run, market price formation should be

[14]The estimates are reported in the caption of Figure 5. The estimates regarding expected date price growth can be compared to the box plots in Figure 4.

governed by traders who maximize expected logarithmic wealth (which is optimal if $\gamma = 1$).[15] As mentioned above, it will be tested based on two specifications of $m_{t,t+\tau}$. The first is a specification which presupposes the null hypothesis. This is the specification which yields the equation in expression (3) if we - counterfactually - assume normally distributed silver and date price growth rates.

The second specification features an additional parameter, denoted as ω, i.e., the share of the RMP's aggregate consumption expenditure spent on dates. We can use this less parsimonious specification with the available data because we found a surprising amount of information regarding Babylonian nutrition, allowing us to calibrate the parameter ω as described the Appendix (below expression (15), which states the second specification of the SDF). With respect to condition (2), the specification can be stated as follows:

$$E_t[\Delta p_{t,t+1}] + \frac{1}{2} Var_t[\Delta p_{t,t+1}] - \lambda = \gamma Cov_t[\Delta \pi_{t,t+1}, \Delta p_{t,t+1}] + \omega(1-\gamma) Var_t[\Delta p_{t,t+1}]. \quad (4)$$

Below, we use this condition to compute the expected return of date storage.

The conditions in expressions (3) and (4) are based on the assumption that silver and date price growth rates are normally distributed. Our estimation process avoids this assumption by using Hansen's generalized method of moments (GMM, Hansen 1982) with moment conditions obtained by specifying the information set I_t in condition (1). The condition should hold for any information included in I_t. We focus on the information that the barley-growing period was a high-risk period, and we measure the extent to which this period overlaps with the period spanned by the prices behind each of our price growth rates. The resulting variable, denoted as z_t, appears in our moment conditions in place of I_t. The details of our estimation process are discussed in the Appendix.

[15] For further discussion, see Lengwiler 2004, p. 90, where the evolutionary argument behind our null hypothesis is attributed to Blume and Easley 1992.

6 Main Results

Table 2 presents our baseline estimates, including standard errors obtained by jackknifing, as described in Cameron and Trivedi 2005. The estimates result from price growth rates (differences between log prices of dates and silver) that satisfy two criteria. The first requires that the underlying prices span a period of up to 60 days. The second criterion excludes growth rates during years of military operations in Babylonia. During these years, it is likely that crop prices and crop storage were subject to non-market forces, as discussed above. To check the robustness of our results, we also exclude observations during years with an army presence in Babylon and during wars abroad. We do so using classifications of the years in our sample period that appear in Pirngruber 2017.

Panel A reports estimates resulting from our two-parameter specification, and Panel B reports estimates based on the three-parameter specification and the assumption that date consumption accounted for $\omega = 25\%$ of the RMP's aggregate (barley and date) consumption. Rows (0) of both panels report results obtained before excluding an influential observation in the Babylonian year 346 BCE. We detect this observation as described in the Appendix under "Estimation". It is not only influential, but also extreme, featuring the most negative price growth rates in the entire sample. In Pirngruber 2017, the year 346 BCE is included in the list of periods of an army presence in Babylon. During these periods, crop storage was probably affected by confiscation risk, suggesting that the price fluctuations during the year 346 BCE may have been amplified by a lack of stored crops occurring for an unusual reason (a revolt of the Sidonians eventually countered by troops sent from Babylonia).

Once the influential observation is excluded, we obtain rather precise estimates of RRA γ. These estimates, reported in rows (1) of Panel A and B, are similar to the RRA-estimates based on modern crop prices reported in Myers 1989. The latter estimates suggest values of RRA between 1 and 3. We can reject a unit value of RRA, as shown in the last column of Table 2.

Our estimates regarding the parameter λ are generally small. By basing our analysis on prices in units of barley, we implicitly specify this parameter as the difference between loss rates in storage of dates and of barley. We cannot reject that storing dates was equally costly as storing barley, but the point estimates suggest that storing dates was less costly.

Rows (2) and (3) of Panel A and B show robustness checks in which we restrict our sample by excluding all growth rates based on prices observed during any periods with an army presence in Babylonia or during wars abroad. Further robustness checks show that our results are also robust when varying key features of our analysis, such as the assumption (in Panel B) that $\omega = 25\%$. See the Online Appendix F.

All estimates discussed so far result from variation in the risk and average return from date storage associated with the barley-growing period (BGP) as a high-risk period. We tried to replicate the results in the second rows of Panels A and B in Table 2 based on a "placebo" high-risk period, i.e. the (Gregorian) months 12-2. This is a period of date storage (after the harvest month 10) that should not be a high-risk period, so that the replication should fail. This indeed happens. The results are reported in the last rows of Panels A and B. In both cases, we obtain nonsensical negative point estimates of RRA with large standard errors. Moreover, we had to resolve computational problems to obtain the estimates in the last row of Panel A.[16] These estimates are therefore based on a slightly extended sample which includes 4 growth rates based on price observations spanning periods between 60 and 90 days.

We end this section by using the condition in expression (4) to compare the BGP and the other months of the year in terms of the expected returns of date storage implied by our estimates. Figure 6 shows plots similar to those in Figure 5, but these plots result from condition (4) if $\omega = 25\%$ and λ and γ are set to the point estimates in row (1) of Panel B in Table 2. The

[16]The estimation terminated with the error: Hessian is not positive semidefinite.

dots in the two plots again show the expected date price growth rates according to our estimates. They suggest that the date price tended to increase during the BGP, but it tended to decrease during the other months of the year. Of course, the other months of the year include the date harvest. We therefore check whether our estimates are robust to excluding observations during the date harvest month. This robustness check is reported in Online Appendix F. It yields results similar to those reported in Table 2.

7 Conclusion

We use ancient price data to measure the relative risk aversion (RRA) of the representative market participant in Babylon's crop market. Using this standard measure of risk aversion, we obtain results that are comparable to estimates based on modern crop prices (Myers 1989) and to more recent RRA-estimates (Calvet et al. 2019). The evidence contributes to our understanding of the economy of ancient Babylon, but it also matters in a broader sense. It suggests that, in terms of the effects of risk on price formation, future markets will be similar to markets of today and to ancient markets. This is relevant for analyses of public policies affecting the risk exposure of future generations. By way of example, our evidence supports the notion that the social cost of carbon (SCC) can be measured using data about suitable present-day markets, e.g., real estate markets (Giglio et al. 2018). Our results also show that in ancient Babylon's crop market, traders responded to risk based on a level of RRA well above 1. This result is *per se* relevant to measuring the SCC: Lemoine 2017 shows that, if based on a RRA above 1, the SCC strongly depends on effects of greenhouse gas emissions on the risk exposure of future generations.

References

Aperghis, G. G. 2004. *The Seleukid royal economy. The finances and financial administration of the Seleukid Empire*. Cambridge University Press.

Arrow, K. J., H. B. Chenery, B. S. Minhas, and R. M. Solow. 1961. "Capital-labor substitution and economic efficiency". *The Review of Economics and Statistics* 43 (3): 225–250.

Bell, A. M. 2009. "Approaching the genomics of risk-taking behavior". *Advances in Genetics* 68:83–104.

Berk, J. B., and J. Walden. 2013. "Limited capital market participation and human capital risk". *The Review of Asset Pricing Studies* 3 (1): 1–37.

Blume, L., and D. Easley. 1992. "Evolution and market behavior". *Journal of Economic Theory* 58:9–40.

Bossaerts, P. 2004. "Filtering returns for unspecified biases in priors when testing asset pricing theory". *The Review of Economic Studies* 71 (1): 63–86.

Brock, W., and L. P. Hansen. 2018. "Wrestling with uncertainty in climate economic models". http://dx.doi.org/10.2139/ssrn.3008833.

Calvet, L. E., J. Y. Campbell, F. J. Gomes, and P. Sodini. 2019. "The cross-section of household preferences".

Cameron, A. C., and P. K. Trivedi. 2005. *Microeconometrics: methods and applications*. Cambridge University Press.

Christopoulos, G. I., P. N. Tobler, P. Bossaerts, R.J. Dolan, and W. Schultz. 2009. "Neural correlates of value, risk, and risk aversion contributing to decision making under risk". *The Journal of Neuroscience* 29:12574–12583.

Deaton, A., and G. Laroque. 1996. "Competitive storage and commodity price dynamics". *Journal of Political Economy* 104 (5): 896–923.

— . 1992. "On the behaviour of commodity prices". *The Review of Economic Studies* 59 (1): 1–23.

Fisher, I., and H. G. Brown. 1911. *The purchasing power of money, its determination and relation to credit, interest and crises*. The Macmillan Company.

Foldvari, P., and B. van Leeuwen. 2014. "Risk aversion and storage in autarkic societies: from Babylonian times until the era of globalization". In *A history of market performance*, 497–512. Routledge.

Giglio, S., M. Maggiori, K. Rao, J. Stroebel, and A. Weber. 2018. "Climate change and long-run discount rates: evidence from real estate". http://dx.doi.org/10.2139/ssrn.2639748.

Goetzmann, W. N. 2016. *Money changes everything: How finance made civilization possible*. Princeton University Press.

Gollier, C. 2013. *Pricing the planet's future*. Princeton University Press.

Grainger, J. D. 1999. "Prices in hellenistic babylon". *Journal of the Economic and Social History of the Orient* 43 (3): 303–325.

Graslin-Thomé, L. 2009. *Les échanges à longue distance en Mésopotamie au Ier millénaire: une approche économique*. De Boccard.

Hansen, L. P. 1982. "Large sample properties of generalized method of moments estimators". *Econometrica* 50 (4): 1029–1054.

Hunger, H., and A. Sachs. 1988, 1989, 1996. *Astronomical Diaries and related texts from Babylonia, vol. I-III*. Verlag der Österreichischen Akademie der Wissenschaften.

Jursa, M. 2010. *Aspects of the economic history of Babylonia in the first millenium BC. Economic geography, economic mentalities, agriculture, the use of money and the problem of economic growth*. Ugarit Verlag.

— . 2008. "The remuneration of institutional laborers in an urban context in babylonia in the first millenium BC". In *L'archive des fortifications de Persépolis: état des questions et perspectives de recherches*, 387–427. De Boccard.

Kogan, L., S. Ross, J. Wang, and M. Westerfield. 2017. "Market selection". *Journal of Economic Theory* 168:209–236.

Lemoine, D. 2017. "The climate risk premium: how uncertainty affects the social cost of carbon". http://dx.doi.org/10.2139/ssrn.2560031.

Lengwiler, Y. 2004. *Microfoundations of financial economics: an introduction to general equilibrium asset pricing*. Princeton University Press.

Malmendier, U., and S. Nagel. 2011. "Depression babies: do macroeconomic experiences affect risk taking?" *Quarterly Journal of Economics* 126 (1): 373–416.

Merton, R. C. 1973. "An intertemporal capital asset pricing model". *Econometrica* 41 (5): 867–887.

Munk, C. 2013. *Financial asset pricing theory*. Oxford University Press.

Myers, R. J. 1989. "Econometric testing for risk-averse behavior in agriculture". *Applied Economics* 21:542–552.

Pirngruber, R. 2017. *The economy of late Achaemenid and Seleucid Babylonia*. Cambridge University Press.

Poyck, A. P. G. 1962. *Farm studies in Iraq: An agro-economic study of the agriculture in the Hilla-Diwaniya area in Iraq*. H. Veenman and Zonen.

Schildberg-Hörisch, H. 2018. "Are risk preferences stable?" *Journal of Economic Perspectives* 32 (2): 135–54.

Slotsky, A. L. 1997. *The bourse of Babylon: market quotations in the astronomical diaries of Babylonia*. CDL Press Bethesda.

Temin, P. 2002. “Price behavior in ancient Babylon”. *Explorations in Economic History* 39:46–60.

van der Spek, R. J. 2005. “Commodity prices in Babylon 385–61 BC”. *International Institute of Social History, accessed at http://www.iisg. nl/hpw/babylon.php*.

— . 2000. “The effect of war on the prices of barley and agricultural land in Hellenistic Babylonia”. In *Économie antique. La guerre dans les économies antiques*, ed. by J. Andreau, P. Briant, and R. Descat. Saint-Bertrand-de-Comminges, Musée archéologique départemental.

Yan, H. 2008. “Natural selection in financial markets: does it work?” *Management Science* 54 (11): 1935–1950.

Appendix

In this Appendix, we discuss the theoretical foundations of our analysis measuring the risk aversion of the representative market participant (RMP) in ancient Babylon. We first present the derivation of expression (1). After that, we present the details of our specifications regarding the stochastic discount factor (SDF) $m_{t,t+\tau}$ which appears in this expression. Given that we use textbook specifications, we base our presentation on Munk 2013, where further information can be found.

Expression (1)

We start with a standard asset pricing condition for economies with two types of consumption (goods), denoted as c_t and q_t. This condition appears above Theorem 6.8 in Munk 2013, p. 215, and it is reproduced below:

$$P_{i,t} = E_t \left[e^{-\delta} \frac{u_c(c_{t+1}, q_{t+1})}{u_c(c_t, q_t)} (P_{i,t+1} + D_{i,t+1}) \right].$$

In this equation, $E_t[\cdot]$ denotes the conditional expectation operator $E[\cdot|I_t]$ that also appears in expression (1), t is a point in time, $P_{i,t}$ denotes the price of an asset i at time t, $D_{i,t+1}$ denotes a payoff that investors obtain at time $t + 1$ upon investing into the asset, and δ is a parameter measuring the RMP's aversion to postponing consumption.

We can drop the subscript i. In our setting, the price P_t is the price of dates at time t, measured in units of barley. The payoff from storing one unit of dates (a given physical quantity, e.g., a litre) for a single period is $(1 - d)$ units of dates, where d denotes a proportional storage loss rate. Multiplying this quantity of dates by the price of dates in units of barley (our "numéraire") converts it into a payoff $(1 - d)P_{t+1}$. We substitute this payoff for the sum

$P_{i,t+1} + D_{i,t+1}$ in the above-stated condition, and obtain the following equation:

$$P_t = E_t \left[e^{-\delta} \frac{u_c(c_{t+1}, q_{t+1})}{u_c(c_t, q_t)} (1-d) P_{t+1} \right].$$

We use a multi-period version of this condition, so that we can analyze date storage over τ periods. To derive this version of the condition, we start with listing the single-period conditions for all points in time between t and $t + \tau - 1$. All of these are versions of the above-stated equation, with different time subscripts:

$$\begin{aligned} P_t =& E_t \left[\xi_{t,t+1} (1-d) P_{t+1} \right], \\ P_{t+1} =& E_{t+1} \left[\xi_{t+1,t+2} (1-d) P_{t+2} \right], \\ &\vdots \\ P_{t+\tau-1} =& E_{t+\tau-1} \left[\xi_{t+\tau-1,t+\tau} (1-d) P_{t+\tau} \right], \end{aligned}$$

where we have used the definition

$$\xi_{t+k,t+k+1} = e^{-\delta} \frac{u_c(c_{t+k+1}, q_{t+k+1})}{u_c(c_{t+k}, q_{t+k})}.$$

The above-stated equations can be combined through a process of recursive substitution. The substitution process (using the second equation to eliminate P_{t+1} from the first equation, etc.) results in an expression involving iterated expectations $E_t[E_{t+1}[\dots]]$ which can be replaced by a simple expectation $E_t[\dots]$ by the law of iterated expectations. We thus obtain

$$P_t = E_t \left[M_{t,t+\tau} (1-d)^\tau P_{t+\tau} \right], \tag{5}$$

where $M_{t,t+\tau}$ is defined as the product $\xi_{t,t+1} \xi_{t+1,t+2} \dots \xi_{t+\tau-1,t+\tau}$.

A similar condition must hold with respect to storing barley. In this case, the prices are normalized to one (because the value of barley is measured in its own units), so that the condition is simply

$$1 = E_t[M_{t,t+\tau}(1-b)^\tau], \tag{6}$$

where the parameter b is the proportional storage loss rate in barley storage.

We now combine the equations in expressions (5) and (6). Dividing both sides of the first equation by P_t, subtracting the second equation, and dividing both sides of the resulting equation by $(1-d)^\tau$, yields the following condition:

$$E_t\left[M_{t,t+\tau}\left(\frac{P_{t+\tau}}{P_t} - \left(\frac{1-b}{1-d}\right)^\tau\right)\right] = 0.$$

Given that we use a time scale in which one period is an interval of 10 days, the proportional storage loss rates d and b are small. We can therefore use the approximation $((1-b)/(1-d))^\tau \approx e^{-\lambda\tau}$, with λ defined as the difference $d-b$. This allows us to rewrite the last equation so that it matches expression (1) (using the definitions $m_{t,t+\tau} = \ln M_{t,t+\tau}$, and $\Delta p_{t,t+\tau} = \ln P_{t+\tau} - \ln P_t$.)

Our Specifications

Overview

We use the following specification:

$$m_{t,t+\tau} = -\delta\tau - \gamma\Delta c_{t,t+\tau} - (1-\gamma)\omega(\Delta c_{t,t+\tau} - \Delta q_{t,t+\tau}), \tag{7}$$

where $\Delta c_{t,t+\tau} = \ln C_{t+\tau} - \ln C_t$, and $\Delta q_{t,t+\tau} = \ln Q_{t+\tau} - \ln Q_t$ are defined as the growth rates of the RMP's consumption concerning dates and barley, respectively,

Expression (7) states the Cobb-Douglas version of a textbook specification of the SDF which is based on CES consumption preferences, proposed in Arrow et al. 1961. It features three parameters, i.e., the RMP's coefficient of relative risk aversion (RRA), γ, the share of the RMP's aggregate (barley and date) consumption expenditures spent on dates, ω, and a parameter δ measuring the RMP's aversion to postponing consumption (impatience). This last parameter also appeared in the last section.

Given the structure of condition (1), it holds for any value of the parameter δ in our specification of the SDF. As a consequence, this parameter does not affect our estimates regarding RRA γ. The estimates will, however, depend on ω, except in the case in which $\gamma = 1$. This case is our null hypothesis. To test the null hypothesis, we use two specifications. The first is based on a simplified version of expression (7):

$$m_{t,t+\tau} = -\delta\tau - \gamma\Delta c_{t,t+\tau}. \tag{8}$$

This specification should be reasonable if the last term of expression (7) can be ignored. In this sense, the estimates resulting from the above-stated specification are conditional on the null hypothesis that $\gamma = 1$. By using this specification, we can test whether $\gamma = 1$ without estimating or specifying the value of ω. We will, however, also compute estimates using expression (7) and assumptions about ω. This is possible because the assumptions can be based on a surprising amount of information regarding Babylonian nutrition, described below (following expression 16)). Estimating both γ and ω requires more data than are available.

By substituting the specifications (7) or (8) into condition (1), we obtain equations in terms of growth rates of the RMP's barley and date consumption, $\Delta c_{t,t+\tau}$ and $\Delta q_{t,t+\tau}$, respectively. We measure these growth rates using the growth rates of the prices of dates and silver (in units of barley), denoted as P_t and Π_t, respectively. This approach is based on the following equations, discussed below:

$$\begin{aligned}&\text{(i)}\;\; P_t = \frac{\omega}{1-\omega}\frac{C_t}{Q_t}, \text{so that } \Delta p_{t,t+\tau} = \Delta c_{t,t+\tau} - \Delta q_{t,t+\tau},\\&\text{(ii)}\;\; \Pi_t \propto C_t, \text{so that } \Delta \pi_{t,t+\tau} = \Delta c_{t,t+\tau}.\end{aligned} \tag{9}$$

The first equation in line (i) follows from the RMP's optimization of barley and date consumption, denoted as C_t and Q_t, respectively. In the optimum, the RMP must neither gain nor lose from a marginal change in date consumption associated with a marginal change in barley consumption so that the RMP's budget constraint remains satisfied. Given this criterion, expression (i) follows from the fact that the date price P_t measures the cost of increasing the RMP's date consumption, and from the Cobb-Douglas specification of the DM's consumption preferences that also yields equation (7).

Line (ii) states that the silver price, denoted as Π_t, is proportional to barley consumption, so that the growth rate of this price $\Delta\pi_{t,t+\tau} = \ln \Pi_{t+\tau} - \ln \Pi_t$ equals the barley consumption growth $\Delta c_{t,t+\tau}$. This specification is based on the argument that the silver price measures the demand for silver as money used in consumption-induced transactions. Below, we formalize this argument based on a common notion in the field of monetary economics, i.e., the equation of exchange, proposed in Fisher and Brown 1911. This equation determines the price level in an economy by equating money supply to money demand, defined as the ratio of the volume of transactions at market prices and the velocity of money circulation. In our setting, the equation determines the price of silver. See expression (12) below.

We next describe how we derive expressions (7) and (9), based on results from Munk 2013. Readers not interested in these details can jump the following section.

Derivations

Expression (7) This is a special case of expression (6.28) in Munk 2013, reproduced below:

$$\zeta_t = e^{-\delta t}\left(\frac{c_t}{c_0}\right)^{-\gamma}\left(\frac{1+\frac{b}{a}\left(\frac{q_t}{c_t}\right)^{\frac{\psi-1}{\psi}}}{1+\frac{b}{a}\left(\frac{q_0}{c_0}\right)^{\frac{\psi-1}{\psi}}}\right)^{-\frac{\psi\gamma-1}{\psi-1}} .$$

We start by re-stating the above-stated expression in terms of our notation. In the above-stated expression:

- Time t denotes the point in time which we denote as $t+\tau$ in expression (7), i.e., the end point of a storage period, while time 0 denotes this period's start time, which we denote as time t in expression (7);
- ζ_t corresponds to the variable $M_{t,t+\tau}$ in our notation;
- Lower-case variables are used in place of the upper-case variables that we use (while reserving lower case for logarithmic values of upper-case variables). For example, c_t is used to denote the consumption that we denote as $C_{t+\tau}$, c_0 corresponds to our C_t, q_0 corresponds to our Q_t, etc.

Translated to our notation, the above-stated expression appears as follows:

$$M_{t,t+\tau} = e^{-\delta\tau}\left(\frac{C_{t+\tau}}{C_t}\right)^{\theta-1}\left\{\frac{\left[(1-\omega)(C_{t+\tau})^{\theta}+\omega(Q_{t+\tau})^{\theta}\right]^{\frac{1}{\theta}}}{\left[(1-\omega)(C_t)^{\theta}+\omega(Q_t)^{\theta}\right]^{\frac{1}{\theta}}}\right\}^{1-\theta-\gamma} . \qquad (10)$$

where we have defined $\omega = b/(a+b)$, and $\theta = 1-1/\psi$.

Our expression (7) is a version of expression (10) which allows us to estimate the risk aversion parameter γ with the available data because it features only three parameters, rather than four. We obtain the three-parameter specification by assuming that the RMP demands barley and dates so that they are neither gross complements nor substitutes. This neutral Cobb-Douglas case is a limit case of expression (10) for $\theta \to 0$, which corresponds to $\psi \to 1$ in the notation of Munk 2013. To derive the Cobb-Douglas specification, we need to compute the limit of the ratio inside the curly braces in expression (10). We do so by using the following limit:

$$
\begin{gathered}
\lim_{\theta \to 0} \ln \left[(1-\omega)X^{\theta} + \omega Y^{\theta}\right]^{\frac{1}{\theta}} = \\
= \lim_{\theta \to 0} \frac{\partial}{\partial \theta} \ln \left[(1-\omega)X^{\theta} + \omega Y^{\theta}\right] = \\
= (1-\omega)\ln X + \omega \ln Y = \\
= \ln X^{1-\omega} Y^{\omega},
\end{gathered}
$$

where the second line uses the L'Hôspital's rule. Removing the logs yields

$$
\lim_{\theta \to 0} \left[(1-\omega)X^{\theta} + \omega Y^{\theta}\right]^{\frac{1}{\theta}} = X^{1-\omega} Y^{\omega}.
$$

This result can be used to replace the expressions in the numerator and the denominator of the ratio stated inside the curly braces in expression (10). After simplifying, we obtain

$$
M_{t,t+\tau} = e^{-\delta\tau} \left(\frac{C_{t+\tau}}{C_t}\right)^{-\gamma} \left(\frac{C_{t+\tau}}{C_t}\frac{Q_t}{Q_{t+\tau}}\right)^{-(1-\gamma)\omega}. \tag{11}
$$

Taking logs yields expression (7).

Expression (9) Line (i) of this expression follows from formula (6.9) in Munk 2013, reproduced below with time subscripts added:

$$P_t = \frac{b}{a}\left(\frac{q_t}{c_t}\right)^{-\frac{1}{\psi}}.$$

The notation used in this expression maps to our notation as discussed above expression (10). We again consider the Cobb-Douglas specification, i.e. the limit of the above-stated expression when $\psi \to 1$. In terms of our notation, the limit is given by $P_t = \omega/(1-\omega)(C_t/Q_t)$, where we have defined $\omega = b/(a+b)$.

We next turn to line (ii) of expression (9). This line is based on the argument that the silver price measures the demand for silver as money used in consumption-induced transactions. We formalize this argument in terms of an ”equation of exchange”, based on Fisher and Brown 1911:

$$(1-\eta)(C_t + Q_t P_t) = vA\Pi_t. \tag{12}$$

The left-hand side specifies the demand for silver associated with the consumption of barley and dates, where η is a parameter which allows for silver not being required in some fraction of the transactions underlying barley and date consumption.[17] The right-hand side is the amount of silver in circulation, valued in barley, where v and A are parameters measuring the velocity of silver circulation, and the quantity of silver available, respectively. Both sides of the equation have barley units, since P_t is the date price in units of barley, while Π_t is the silver price in units of barley. If v and A changed over time, then these changes would also affect the silver price.

[17] It seems that barter was common in Babylon's crop market. In fact, van der Spek writes on his website (from which we obtained our data) that the system of price quotations in the Babylonian crop market was partly motivated by the need to facilitate barter. The quotations were exchange values of silver, i.e., quantities that can be bought for a shekel of silver. van der Spek writes: “Thinking in exchange values facilitates trade by barter. When you know that you can buy 180 litres of barley and 120 litres of dates, you can also exchange these two items.” In addition, crop consumption in Babylon was partly based on workers receiving crop rations. See the Section “Wages” on van der Spek's website.

To abstract from this possibility, we focus on silver price fluctuations over short periods of time, essentially assuming that the velocity-adjusted silver supply vA remained constant in the short run. Given this assumption, $\Delta\pi_{t,t+\tau} = \Delta c_{t,t+\tau}$ since the left-hand side of expression (12) is proportional to C_t because $P_t = \omega/(1-\omega)(C_t/Q_t)$.

Estimation Process

The estimates in our paper are based on Hansen's generalized method of moments (Hansen 1982). Here, we explain how we obtain the underlying moment conditions. The conditions follow from expression (1), upon substituting the expressions stated below for m_{t+1} and specifying the information set I_t. We do the latter using a variable, denoted as z_t, which measures the overlap of the barley-growing period with the interval spanned by the prices behind each of our price growth rates. For example, suppose that we want to measure the overlap between the BGP and a price growth rate based on price information recorded at the beginning of month 2, and at the end of month 3 of some year. If the BGP is defined as the period consisting of three months between the beginning of month 3, and the end of month 5, then the overlap is $1/2$ since the underlying prices span a period of 2 months, one of which is part of the BGP.

Moment Conditions

We use two different moment conditions. The first results from the simple specification in expression (8), stated above. Given that $\Delta c_{t,t+\tau} = \Delta\pi_{t,t+\tau}$ (line (ii) of expression (9)),

$$m_{t,t+\tau} = -\delta\tau - \gamma\Delta\pi_{t,t+\tau}. \tag{13}$$

If this expression is substituted into expression (1), we obtain a condition from which we can cancel $\exp(-\delta\tau)$. This yields the following condition (with I_t replaced by z_t):

$$E\left[e^{-\gamma\Delta\pi_{t,t+\tau}}(e^{\Delta p_{t,t+\tau}} - e^{\lambda\tau})|z_t\right] = 0. \tag{14}$$

This moment condition generates the results reported in Panel A of Table 2 and Panels A of Tables S7-S10 in the Online Appendix F.

The second moment condition is based on the specification in expression (7), stated above. We can re-state this specification in terms of the price growth rates $\Delta\pi_{t,t+\tau}$ and $\Delta p_{t,t+\tau}$ based on the system of equations in expression (9): $\Delta c_{t,t+\tau} = \Delta\pi_{t,t+\tau}$ and $\Delta c_{t,t+\tau} - \Delta q_{t,t+\tau} = \Delta p_{t,t+\tau}$. This yields:

$$m_{t,t+\tau} = -\delta\tau - \gamma\Delta\pi_{t,t+\tau} - (1-\gamma)\omega\Delta p_{t,t+\tau}. \tag{15}$$

If this expression is substituted into expression (1), we again obtain a condition from which we can cancel $\exp(-\delta\tau)$. This yields the following condition (with I_t again replaced by z_t):

$$E\left[e^{-\gamma\Delta\pi_{t,t+\tau}-(1-\gamma)\omega\Delta p_{t,t+\tau}}(e^{\Delta p_{t,t+\tau}} - e^{\lambda\tau})|z_t\right] = 0. \tag{16}$$

This moment condition generates the results reported in Panel B of Table 2, Panels A and B of Table S6 in the Online Appendix F, and Panels B of Tables S7-S10. To obtain the results, we fix the value of ω, avoiding the need to estimate an additional parameter. We do so using two sources of information. The first (Aperghis 2004) compares barley and dates, stating (p. 254) that "taking calorie content and price into consideration, it is not unreasonable to consider dates the rough equivalent of barley, or slightly less." This statement suggests that we can base a crude estimate of ω on the relative contribution of barley and dates to the caloric intake of ancient Babylonians. To do so, we draw on evidence in Poyck 1962 that dates accounted for about

25.6% of the calorie intake of farmer families living in the area in 1960. Given this evidence, we use $\omega = 25\%$ as our baseline value, effectively assuming that, in ancient Babylonia, dates were an equally important source of calories as in 1960. In robustness checks (reported in Table S6 in the Online Appendix F), we will vary the value of ω by $\pm 5\%$. A much higher value of ω is unreasonable because of the low protein content of dates relative to barley, and the greater need for protein of ancient Babylonians if their diet included less meat than that of farmers living in the area in 1960.

Estimation

Estimation is performed using the "gmm onestep" algorithm in Stata. All estimates result from starting values $\gamma = 1$ and $\lambda = 0$. We also tried other reasonable starting values, but we could not detect an effect on the estimates. All standard errors are estimated by means of jackknife resampling, as described in Cameron and Trivedi 2005, p. 375. To do so, we use the "vce(jack)" option of the algorithm. We also tried bootstrapping, but this proved impossible given the small size of our sample.

Jackknife resampling is also used to detect influential observations. The starting point is the set of estimates obtained by excluding observations from the estimation sample, one at a time. We first compute this set of estimates based on the "jackknife" algorithm in Stata. For each $\gamma-$estimate in the set, we then compute the difference between this estimate and the mean of all γ-estimates in the set, divided by their standard deviation. The results are z-values that show the degree to which observations are influential. It turns out that there is one observation with a z-value above 3 (given either set of estimates resulting from our two moment conditions). This is an observation in the year 346 BCE that features the most negative silver and date price growth rates in the entire sample, so that it is not only influential but also extreme. As discussed in the "Main Results" section of our paper, we have reasons to believe that the influential and extreme nature of the observation results from date storage in 346 BCE being affected by non-market forces (risk of confiscation). We therefore exclude the influential observation from our analysis, but we do report the results that we obtain before excluding this observation. These results appear in the rows (0) of Panels A and B in Table 2. Comparing the results to those in the rows (1) shows that excluding the influential observation substantially reduces our standard errors, but it yields point estimates that are less than one standard error away from the point estimates in the rows (0) (based on these rows' standard errors). The influential observation is also excluded in the analyses that yield the estimates in the rows (2)-(4) of Panels A and B in Table 2.

Figures and Tables

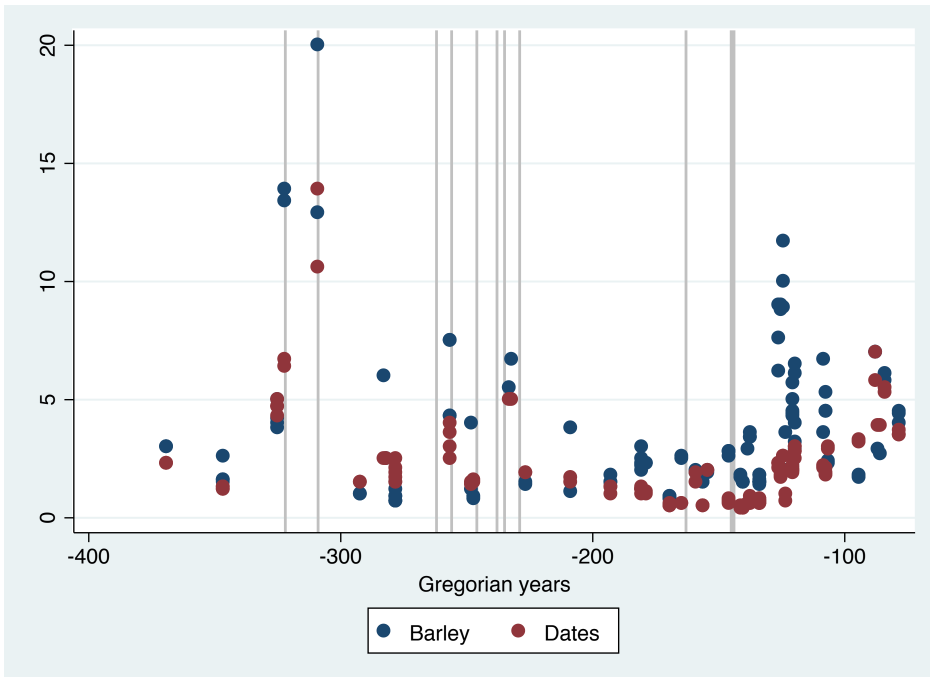

Figure 1.

Prices of barley and dates in Babylon, 385 to 61 BCE. The prices are measured in shekels per *kurru* (180 litres). The plot shows all available prices observed at most 60 days apart, i.e., the prices behind the short-run price growth rates used in our empirical analysis.

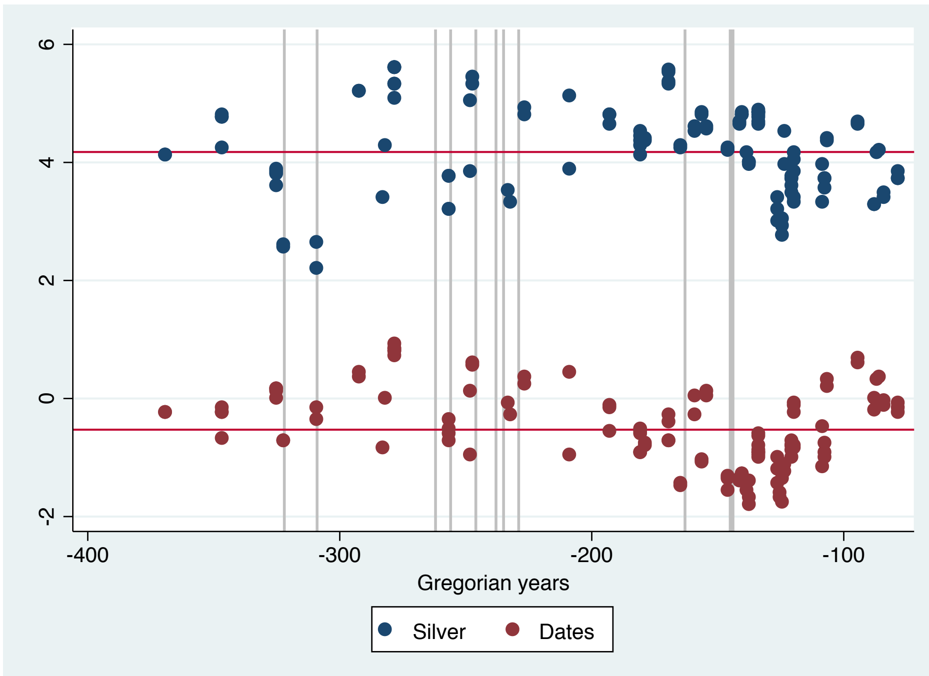

Figure 2.

Logarithmic prices of silver and dates in Babylon, 385 to 61 BCE. The prices are measured in *kurru* (180 litres) of barley. The silver price is the price of 1 shekel while the date price is the price of 1 *kurru* of dates. The plot shows logarithmic values of all available prices observed at most 60 days apart, i.e., the prices behind the short-run price growth rates used in our empirical analysis.

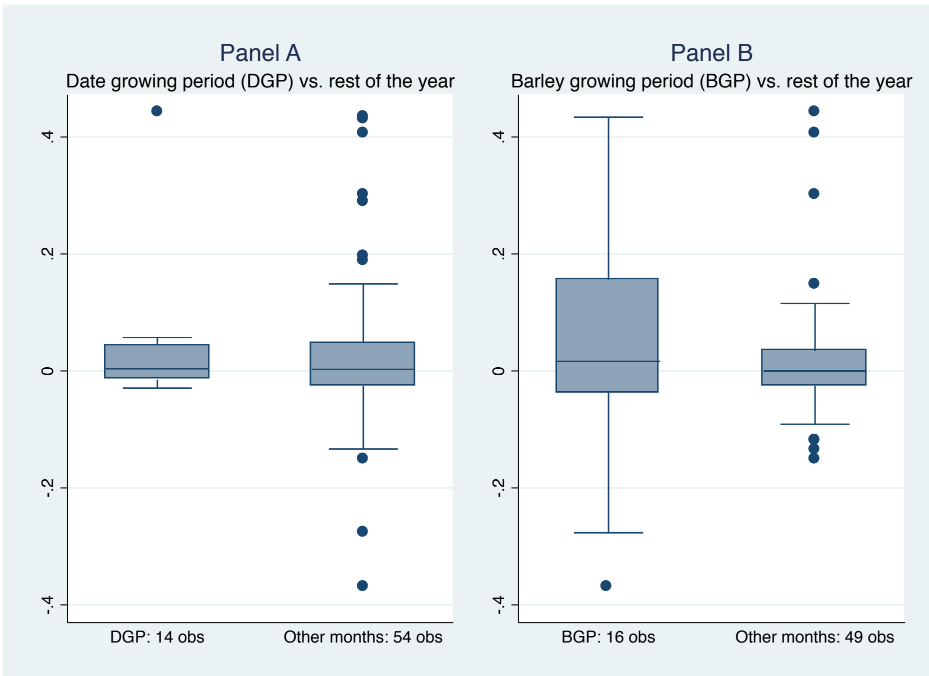

Figure 3.

Average 10-day growth rates of the silver price in Babylon, 385 to 61 BCE. The growth rates are differences between log prices (in units of barley), divided by the number of 10-day periods between the underlying price observations. All growth rates result from prices spanning periods of up to 60 days. In Panel A, we compare observations that are either fully inside or fully outside the date-growing period (DGP, defined as Gregorian months 8-10). In Panel B, we show a similar comparison concerning the barley-growing period (BGP, defined as months 3-5).

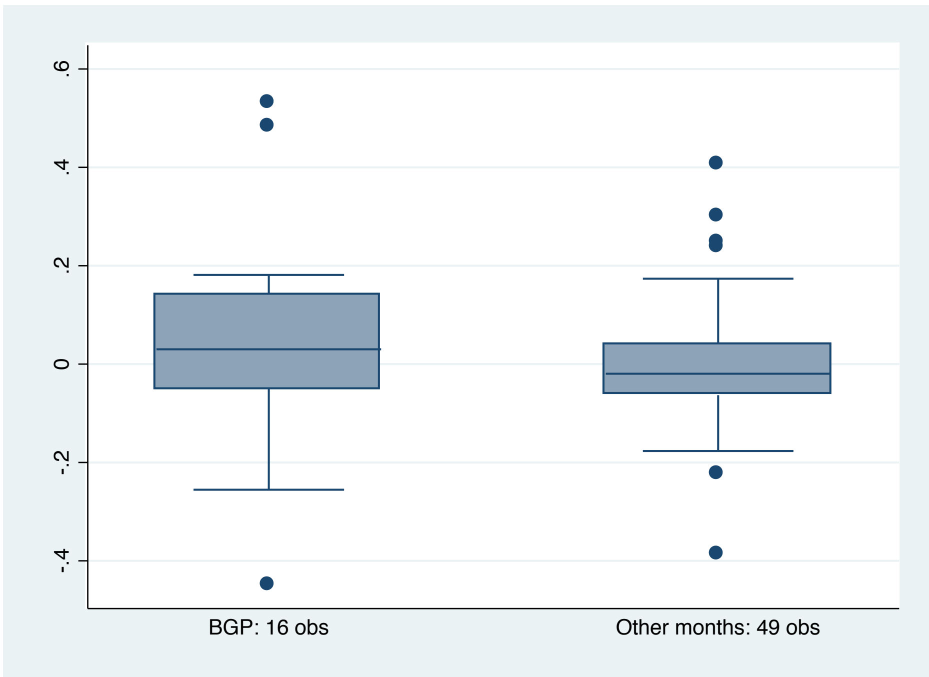

Figure 4.

Average 10-day growth rates of the date price in Babylon, 385 to 61 BCE. The growth rates are differences between log prices (in units of barley), divided by the number of 10-day periods between the underlying price observations. All growth rates result from prices spanning periods of up to 60 days. The box plots compare observations that are either fully inside or fully outside the barley-growing period (BGP, defined as the months 3-5).

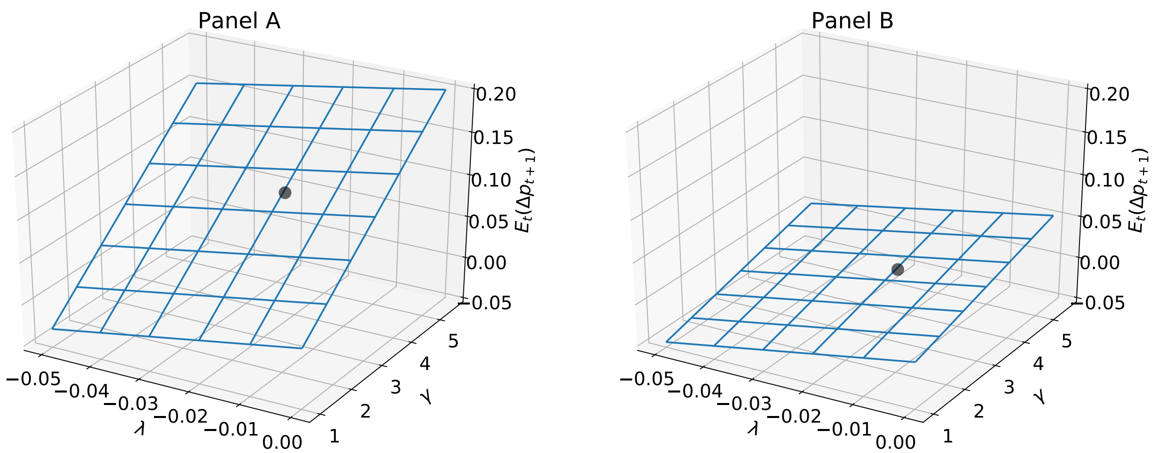

Figure 5.

Expected 10-day date price growth during the barley growing period (BGP, defined as months 3-5) and the other months of the year. The parameter λ is the difference between the cost of storing dates and the cost of storing barley (both measured in terms of physical storage loss rates). The parameter γ is a coefficient of relative risk aversion. The plots are are based on expression (3) and maximum likelihood estimates of the terms in this expression, given the values of γ and λ stated in row (1) of Panel A in Table 1. The estimates of the covariance of silver and date price growth are 0.0410 (BGP) and 0.0119 (other months). The estimates of the variance of date price growth are 0.0477 (BGP) and 0.0215 (other months).

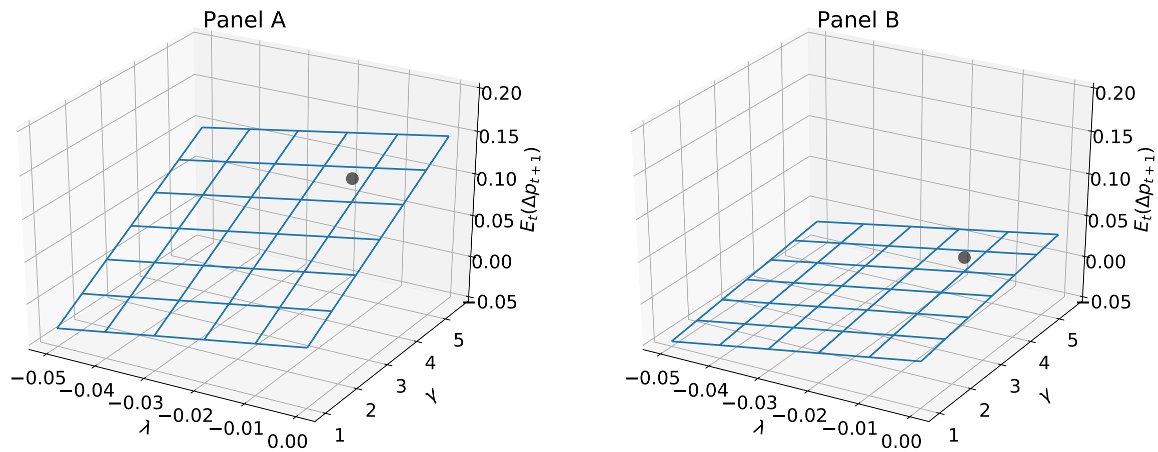

Figure 6.

Expected 10-day date price growth during the barley growing period (BGP, defined as months 3-5) and the other months of the year. The parameter λ is the difference between the cost of storing dates and the cost of storing barley (both measured in terms of physical storage loss rates). The parameter γ is a coefficient of relative risk aversion. The plots are are based on expression (4) and maximum likelihood estimates of the terms in this expression, given the values of γ and λ stated in row (1) of Panel B in Table 1, and the assumption that date consumption accounted for $\omega = 25\%$ of aggregate (barley and date) consumption in ancient Babylon. The estimates of the covariance of silver and date price growth are 0.0409 (BGP) and 0.0121 (other months). The estimates of the variance of date price growth are 0.0478 (BGP) and 0.0222 (other months).

Descriptive Statistics					
Panel A: Baseline Samples	**1st Quartile**	**3rd Quartile**	**Mean**	**Std. Dev.**	**N**
Silver Price Growth	-0.0252	0.0445	0.0262	0.1689	73
Date Price Growth	-0.0576	0.0515	0.0114	0.1882	73
Panel B: Samples excl. Observations with Imprecise Timing					
Silver Price Growth	-0.0257	0.0530	0.0241	0.1562	52
Date Price Growth	-0.0513	0.0652	0.0070	0.1689	52
Panel C: Samples excl. Doubtful Observations					
Silver Price Growth	-0.0252	0.0488	0.0226	0.2005	41
Date Price Growth	-0.0605	0.0661	0.0092	0.2213	41
Panel D: Core Samples (Intersections of Samples in Panels A-C)					
Silver Price Growth	-0.0371	0.0387	-0.0003	0.1631	28
Date Price Growth	-0.0576	0.0661	-0.0184	0.1637	26
Panel E: Complements of Core Samples (in Panel A's Samples)					
Silver Price Growth	-0.0148	0.0445	0.0428**	0.1706	45
Date Price Growth	-0.0629	0.0446	0.0279	0.1996	47
Panel F: Comparing Statistics in Panels D and E (p-values for $H_0 : D = E$)					
Silver Price Growth			0.213	0.818	
Date Price Growth			0.218	0.289	

Table 1.

Descriptive statistics concerning price growth rates in Babylon, 385 to 61 BCE. We focus on growth rates spanning at most 60 days. The columns "1st Quartile", "3rd Quartile", and "Mean" are based on average 10-day price growth rates defined as differences of log prices (in units of barley), divided by the number of 10-day periods between the underlying price observations. The column "Std. Dev." is based on re-scaled growth rates, defined as differences of log prices (in units of barley), divided by the square root of the number of 10-day periods between the underlying price observations. Column "N" reports the numbers of growth rates in each sample. Panel A describes our baseline samples, while Panels B and C concern samples obtained by excluding price observations affected by imprecise within-month timing (Panel B) or doubtful silver exchange values (Panel C). Panel D describes the growth rates that appear in all of the samples behind Panels A-C, and Panel E describes this sample's complement within our baseline sample (Panel A). Panel F reports p-values of two-tailed tests comparing the samples described in Panels D and E. The p-values in the column "Mean" result from t-tests (allowing for unequal variances) while those in the column "Std. Dev" result from variance ratio tests.

Panel A: GMM Estimates Conditional on the Null Hypothesis $H_0 : \gamma = 1$						
	Relative Storage Cost λ		Relative Risk Aversion γ			$H_0 : \gamma = 1$
	Estimate	St. Error	Estimate	St. Error	Obs.	p-val.
(0)	-0.0100	0.0168	1.7989	2.0121	67	0.693
(1)	-0.0196	0.0128	3.5898***	0.9220	66	0.007
(2)	-0.0214	0.0131	3.6279***	0.9159	63	0.006
(3)	-0.0110	0.0148	3.2359***	0.9828	61	0.027
(4)	0.0608	0.0459	-0.5176	0.9106	70	0.101

Panel B: GMM Estimates Conditional on the Share of Dates in Consumption $\omega = 0.25$						
	Relative Storage Cost λ		Relative Risk Aversion γ			$H_0 : \gamma = 1$
	Estimate	St. Error	Estimate	St. Error	Obs.	p-val.
(0)	-0.0080	0.0135	2.0460	2.7141	67	0.701
(1)	-0.0129	0.0127	4.5011***	1.3761	66	0.013
(2)	-0.0143	0.0132	4.5575***	1.3741	63	0.012
(3)	-0.0050	0.0143	3.9749***	1.4246	61	0.041
(4)	0.0610	0.0631	-1.3229	1.8818	66	0.222

Table 2.

GMM estimates based on growth rates (differences in log prices) of silver and dates in Babylon, 385 to 61 BCE. The underlying price observations span periods of up to 60 days which do not overlap with periods of military operations in Babylonia. We estimate a relative cost of storing dates and barley, λ, and a coefficient of relative risk aversion γ, exploiting variation in risk associated with the barley-growing period (BGP, months 3-5) as a high-risk period. The moment conditions appear in the Appendix. In Panel A, we use a simple condition that presupposes the null hypothesis $\gamma = 1$. In Panel B, we use a condition with an additional parameter, set to $\omega = 0.25$ as discussed in the Appendix. In all rows except rows (0), we exclude one influential observation during a period of army presence, i.e. the year 346 BCE. In rows (2), we exclude all periods of army presence. In rows (3), we exclude all periods of wars abroad. Rows (4) report results based on a “placebo” high-risk period, i.e. the months 12-2. Standard errors result from jackknifing. The symbols *, **, and *** mark estimates significant at levels of 10%, 5%, and 1%, respectively.