

DISCUSSION PAPER SERIES

DP15688

Individual Discount Rates: A Meta- Analysis of Experimental Evidence

Jindrich Matousek, Tomas Havranek and Zuzana
Irsova

FINANCIAL ECONOMICS

CEPR

Individual Discount Rates: A Meta-Analysis of Experimental Evidence

Jindrich Matousek, Tomas Havranek and Zuzana Irsova

Discussion Paper DP15688
Published 20 January 2021
Submitted 20 January 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Financial Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Jindrich Matousek, Tomas Havranek and Zuzana Irsova

Individual Discount Rates: A Meta-Analysis of Experimental Evidence

Abstract

A key parameter estimated by lab and field experiments in economics is the individual discount rate--and the results vary widely. We examine the extent to which this variance can be attributed to observable differences in methods, subject pools, and potential publication bias. To address the model uncertainty inherent to such an exercise, we employ Bayesian and frequentist model averaging. We obtain evidence consistent with publication bias against unintuitive results. The corrected mean annual discount rate is 0.33. Our findings also suggest that discount rates are independent across domains: people tend to be less patient when health is at stake compared to money. Negative framing is associated with more patience. Finally, the results of lab and field experiments differ systematically, and it also matters whether the experiment relies on students or uses broader samples of the population.

JEL Classification: D01, C83, C90

Keywords: Discount rate, Experiment, Publication bias, meta-analysis, Bayesian model averaging, frequentist model averaging

Jindrich Matousek - jindrich.matousek@fsv.cuni.cz
Charles University, Prague

Tomas Havranek - tomas.havranek@ies-prague.org
Charles University, Prague and CEPR

Zuzana Irsova - zuzana.irsova@ies-prague.org
Charles University, Prague

Acknowledgements

Matousek acknowledges support from the Czech Science Foundation (grants \#18-02513S and \#21-09231S) and Charles University (project Primus/17/HUM/16). Havranek and Irsova acknowledge support from the Czech Science Foundation (grant \#19-26812X).

Individual Discount Rates: A Meta-Analysis of Experimental Evidence*

Jindrich Matousek, Tomas Havranek, and Zuzana Irsova

Charles University, Prague

January 20, 2021

Abstract

A key parameter estimated by lab and field experiments in economics is the individual discount rate—and the results vary widely. We examine the extent to which this variance can be attributed to observable differences in methods, subject pools, and potential publication bias. To address the model uncertainty inherent to such an exercise, we employ Bayesian and frequentist model averaging. We obtain evidence consistent with publication bias against unintuitive results. The corrected mean annual discount rate is 0.33. Our findings also suggest that discount rates are independent across domains: people tend to be less patient when health is at stake compared to money. Negative framing is associated with more patience. Finally, the results of lab and field experiments differ systematically, and it also matters whether the experiment relies on students or uses broader samples of the population.

Keywords: Discount rate, experiment, publication bias, meta-analysis, Bayesian model averaging, frequentist model averaging

JEL Codes: D01, C83, C90

*Corresponding author: Zuzana Irsova, zuzana.irsova@ies-prague.org. An online appendix with data and code is available at meta-analysis.cz/discrate. Matousek acknowledges support from the Czech Science Foundation (grants #18-02513S and #21-09231S) and Charles University (project Primus/17/HUM/16). Havranek and Irsova acknowledge support from the Czech Science Foundation (grant #19-26812X).

1 Introduction

Intertemporal trade-offs are key to a host of decision problems at both the private and public levels. For some of these decisions, it is appropriate to employ the market discount rate, which is detectable from financial time series. For others, however, we must try to recover the underlying discount rates of individuals—rates that also reflect the underlying transaction costs of borrowing money that households face (Kovacs & Larson, 2008). Policies addressing climate change, particularly those underpinned by the literature on the social cost of carbon, constitute a typical example of choices for which individual discounting of future costs and benefits plays a crucial role (Tol, 1999; Goulder & Stavins, 2002; Fujii & Karp, 2008; Anthoff *et al.*, 2009).

Individual discount rates can be either observed from existing data (such as in Lawrance, 1991; Dreyfus & Viscusi, 1995; Warner & Pleeter, 2001) or measured experimentally (Benzion *et al.*, 1989; Chapman & Elstein, 1995; Coller & Williams, 1999; Harrison *et al.*, 2010, among others). We focus on the latter: experiments. Controlled experiments provide a natural framework for exploring time discounting in both laboratory and field conditions by enabling researchers to vary the parameters in order to infer the subject’s preferences. However, despite decades of work and dozens of experiments devoted to eliciting time preferences, no consensus on how to best measure discounting has emerged (Andreoni *et al.*, 2015). It is safe to say that the discount rate differs across individuals and its estimates vary a great deal throughout the literature, sometimes by orders of magnitude (Coller & Williams, 1999; Frederick *et al.*, 2002).

In this paper we take stock of the evidence and aim to trace the differences in the reported discount rates to the design of experiments while accounting for model uncertainty. We also control for the effects of potential selective reporting, a phenomenon found to be widespread in economics and other fields (Doucouliagos & Stanley, 2013; Ioannidis *et al.*, 2017). Focusing on aspects related to study design, methodology, and subject pool characteristics, we collect a set of 22 explanatory variables and employ Bayesian model averaging (BMA; Raftery *et al.*, 1997) and frequentist model averaging (FMA; Hansen, 2007) to examine which ones matter the most for the differences among the reported estimates. Model averaging techniques estimate many regressions with various combinations of the 22 variables and then weight the models according to data fit, parsimony, and collinearity.

The closest work to our own is the meticulous meta-analysis by Imai *et al.* (2020a), who employ a similar methodology but focus on the present-bias parameter estimated using the convex time budget protocol. They find that the literature implies the present-bias parameter to lie between 0.95 and 0.97 on average and describe the

sources of heterogeneity: for example, experiments that use monetary rewards tend to find little evidence of present bias. Other related recent studies include Brown *et al.* (2020), who meta-analyze the estimates of loss aversion, Imai *et al.* (2020b), who estimate the degree of publication bias in laboratory experiments in economics, and a series of important works evaluating the replicability of experiments in economics and other social sciences (Camerer *et al.*, 2016, 2018; Altmejd *et al.*, 2019).

Our results are consistent with the notion that selective reporting (which causes publication bias) represents an important factor in the literature. When selective reporting is present, insignificant and negative estimates are discriminated against. A zero or negative discount rate, of course, makes little sense in most contexts. Nevertheless, given sufficient noise in the experimental setup, we should sometimes observe insignificant estimates and sometimes very large positive estimates. If non-positive estimates (which are unintuitive) are discarded but large positive estimates (for which it is difficult to determine whether they are intuitive or not) are kept, harmful publication bias arises. This outcome is paradoxical because selective reporting can be beneficial at the micro level: for an individual study, it is most likely a wise choice not to build the story around negative or insignificant estimates of the discount rate. However, at the macro level, the discarding rule is asymmetrical since large estimates are typically not omitted. Our findings indicate that such publication bias is associated with exaggerating the mean reported annual discount rate from 0.33 to 0.80.

Aside from publication bias, which manifests as a correlation of the discount rate estimates with their standard errors, the differences in results seem to be caused primarily by the experimental design of discounting tasks. We find evidence in line with domain independence in intertemporal choice (Loewenstein *et al.*, 2003; Ubfal, 2016): it matters what the experimental subjects should be patient or impatient about. Subjects are more patient with regard to money than health or more exotic contexts (such as vacations, certificates, and kisses from movie stars). The results support the hypothesis that liquidity constraints play a key role in intertemporal choice experiments (Dean & Sautmann, 2020), since health and kisses from movie stars are more difficult than money to transfer over time (Bleichrodt *et al.*, 2016). We also find that negative framing is associated with more patience, which corroborates the notion that anticipation of dread is important in intertemporal decisions (Harris, 2012).

Our results offer three broad implications for economics experiments in general. First, it matters whether the experiment is conducted in the lab or in the field. Lab experiments yield systematically larger discount rates, indicating greater impatience. Second, the composition of the sample of experimental subjects (the subject pool) has a systematic impact on the results. Experiments working exclusively with students

show less evidence for patience than experiments using mixed population samples. Taken together, these two results might question the external validity of some experiments. Nevertheless, we also have good news for the experimental economics community: Third, it does not matter systematically for the results whether experiments use real or hypothetical rewards.

Three caveats of our results are in order. First, we are unlikely to cover all experiments ever conducted on the discount rate. Nevertheless, a meta-analysis does not have to collect the entire universe of available studies; it is important only to avoid selecting studies based on their results. Second, fewer than two-thirds of the collected estimates are reported together with a measure of uncertainty from which we can directly compute standard errors. We address this problem partially by re-sampling standard errors at the study level for observations with missing data. (Limiting our attention to the studies that explicitly report precision would not change our main results.) Third, although we control for the differences in many features of study design, experiments involve unique methodological as well as procedural details that are difficult to codify but that can cause differences in the results of individual studies. Some of these unobserved features might be correlated not only with the reported discount rate but also with the reported standard error, which might make our results concerning publication bias spurious. We partially address this problem by using study fixed effects, caliper tests, p-uniform*, and by employing the number of observations in primary studies as an instrument for the standard error.

The remainder of the paper is structured as follows. Section 2 reviews the basic concepts of discounted utility models and discusses the methods of discount rate elicitation. Section 3 describes our approach to data collection and presents an overview of our dataset. Section 4 examines the extent of publication bias using meta-regression and other meta-analysis techniques. Section 5 investigates the sources of heterogeneity in the estimated discount rates using Bayesian model averaging. Section 6 concludes the paper. Supplementary data, codes, statistics, and diagnostics for the BMA and robustness checks to all analyses presented in the main body are available in Appendix A, Appendix B, and online at meta-analysis.cz/discrete.

2 Estimating the Discount Rate

In this section we do not attempt to provide a comprehensive review of the methodology used to measure discounting but briefly describe the basic concepts that are necessary for the understanding of our meta-analysis. For a more detailed treatment, we refer the reader to the authoritative works by Frederick *et al.* (2002), Andersen *et al.*

(2014), Cheung (2016), and Cohen *et al.* (2020).

The theory of intertemporal choice and discounting dates back to Irving Fisher's *Theory of Interest* (Fisher, 1930) and Paul Samuelson's *Note on Measurement of Utility*, in which he postulated the discounted utility model (Samuelson, 1937). His model was widely accepted together with its central idea of concentrating various decisions about intertemporal choice into a single parameter—the discount rate. Several modifications to the original discount function have been introduced to capture various features, such as hyperbolic (Ainslie, 1975; Mazur, 1984) or quasi-hyperbolic (Phelps & Pollak, 1968; Laibson, 1997) discounting functions.

The discounted utility model captures the time preferences of an individual—more specifically, an individual's preference for immediate utility over delayed utility, represented by her intertemporal utility function $U^t(c_t, \dots, c_T)$, which can be described by the functional form presented in Equation 1:

$$U^t(c_t, \dots, c_T) = \sum_{k=0}^{T-t} D(k) \cdot u(c_{t+k}) , \quad (1)$$

where $D(k)$ is the discount function and $u(c_{t+k})$ is an instantaneous utility function that can be interpreted as an individual's well-being in period $t + k$. The discount function $D(k)$ represents the relative weight that the individual places in period t on her well-being in period $t + k$ and encompasses parameter δ , which represents the individual's discount rate. This discount function can have different functional forms.

The standard exponential model, a well-known functional form used in the majority of practical applications, follows:

$$D^E(k) = \frac{1}{(1 + \delta)^k} , \quad k \geq 0 \quad (2)$$

where the discount rate d is $d^E(k) = \delta$. The key feature of this model is that the discount rate $d^E(k)$ is constant over time, i.e., the rate at which an individual discounts future well-being between today and tomorrow is identical to the rate at which she discounts well-being between one month from today and one month from tomorrow. In contrast, a widely documented situation in which an individual has a declining rate of time preference is described as hyperbolic discounting, which generally means that the implicit discount rate over longer time horizons is lower than the implicit discount rate over shorter time horizons. A typical case from the family of hyperbolic discounting functions proposed by Mazur (1984) is described in Equation 3:

$$D^H(k) = \frac{1}{1 + \delta k} , \quad (3)$$

where the hyperbolic discount rate $d^H(k) = (1 + \delta k)^{\frac{1}{k}} - 1$ (Andersen *et al.*, 2014).¹ Phelps & Pollak (1968) further introduced a quasi-hyperbolic specification of the discount function for use in a social planner problem:

$$D^{QH}(k) = \begin{cases} 1, & \text{if } k = 0 \\ \frac{\beta}{(1+\delta)^k}, & \text{if } k > 0 \end{cases} \quad (4)$$

where $\beta \leq 1$ and the quasi-hyperbolic discount rate $d^{QH}(k) = \left(\frac{\beta}{(1+\delta)^k}\right)^{-\frac{1}{k}} - 1$.² A characteristic feature of the quasi-hyperbolic specification is the discontinuity at time $t = 0$. This specification was applied by Laibson (1997) to model individual agent behavior.

Several experimental methods are available to elicit time preferences in both laboratory and field settings, such as lotteries, choice lists, and bidding; however, there is no consensus on how to best measure discounting (Andreoni *et al.*, 2015). The basic method for eliciting individual discount rates is conceptually simple—asking subjects questions about whether they prefer an amount of money today (option A) or the same amount + \$X tomorrow (option B). By changing X, a researcher can infer bounds for the subject’s individual discount rate.³ Experiments therefore involve a series of questions aligned in lists, such as in the classical choice list design of Collier & Williams (1999) or Harrison *et al.* (2002). Modifications to this basic method are further used to elicit preferences more precisely, such as variations in the delay between options A and B, the domain in which preferences are revealed (money, health, etc.), and the magnitude or the nature of the reward (hypothetical or real).

Several types of elicitation methods are routinely used in the experimental literature (Frederick *et al.*, 2002): i) choice, ii) matching, iii) rating, and iv) pricing. The most common type of elicitation is the choice method, where subjects are presented alternative options and are asked to simply choose between them. This method provides discount rate intervals pre-generated by the experimenter rather than precise estimates of the discount rate for specific individuals. The matching method, in contrast, provides an exact inference of the individual’s discount rate since she reveals her true indifference point by filling the blank field to equate two intertemporal options. In rating tasks, subjects evaluate individual options by rating their attractiveness on

¹In a hyperbolic specification, the discount rate is the value of $d^H(k)$ that solves $D^H(k) = 1/(1 + d^H)^k$, i.e., the equation $1/(1 + \delta k) = 1/(1 + d^H)^k$.

²Again, in the quasi-hyperbolic specification, the discount rate is the value of $d^{QH}(k)$ that solves $D^{QH}(k) = 1/(1 + d^{QH})^k$, i.e., the equation $\beta/(1 + \delta)^k = 1/(1 + d^{QH})^k$.

³The point of the first switch to option B gives a measure of the upper bound of her discount rate. We assume linear utility here for simplicity and discuss relaxing of this assumption later.

a predefined scale, while in pricing tasks, subjects specify their willingness to pay for individual options in which they either obtain or avoid a particular outcome. In contrast to choice and matching tasks, rating and pricing tasks allow the researcher to manipulate the time variable between subjects since immediate and delayed options are evaluated separately.

Each method described briefly above has its strengths and limitations. When subjects are asked to evaluate multiple options at once in a standard choice list, the earlier choices inevitably influence the choices made later. This procedural limitation—the anchoring effect—can be partially addressed by employing titration procedures and exposing subjects to a sequence of different opposing anchors (Frederick *et al.*, 2002). The timing of an outcome was found to have a much lower effect when evaluating a single option compared to a situation when two options occurring in different times are evaluated against each other at once (Loewenstein, 1987). The timing of two evaluating options is further argued to cause the more general problem of an additional risk or transaction costs imposed on the future option. The recent literature, represented by Harrison *et al.* (2005), Andersen *et al.* (2014), and others, deal with this risk by employing a front-end delay, thereby shifting the immediate option to the nearer future and imposing transaction costs on the instant payoff.

Harrison *et al.* (2005) argue that standard choice tasks often executed through multiple price lists (MPL) have three possible disadvantages: i) they elicit only interval responses; ii) they allow subjects to switch back and forth while moving down the list; and iii) they can be subject to framing effects. Harrison *et al.* (2005) therefore introduces an *iterative Multiple Price List* (iMPL) that allows the subjects to iteratively specify their choices through refined options within an interval chosen in the last option.

The inference of discount rates from the experimental task depends on the utility function presented in the discounted utility model (Equation 1). This function, however, is unobserved and therefore usually assumed to be linear, generating biased estimates for individuals with non-linear utility functions (Cheung, 2016). Recent papers by Andersen *et al.* (2008, 2014) use the *joint elicitation strategy* to measure time preferences by controlling for non-linear utility. Using the equivalence of utility for risk and time, these authors use a series of binary choices to infer the discount function conditional on the utility function elicited through Holt & Laury (2002)'s risk preference task. Further modifications of the design to measure time preferences by controlling for non-linear utility include, among others, the work of Laury *et al.* (2012), who interact risk with time using a lottery to be paid out with probability p_t in time t and with probability p_{t+k} in time $t + k$, where $p_t \leq p_{t+k}$ and p_{t+k} vary through the choice list. Further experiments measuring time preferences while controlling for non-linear util-

ity are conducted by Takeuchi (2011), who employs separate choices under risk and over time using matched pairs of payoffs; Andreoni & Sprenger (2012), Andreoni & Sprenger (2012b), and Andreoni *et al.* (2015), who examine risk and time preferences through individual elicitation methods—convex time budgets and double multiple price list tasks—and Attema *et al.* (2016), who introduce a *direct method* to measure discounting that is not dependent on the knowledge or measurement of utility.

An alternative method for inferring discount rates was devised by Chabris *et al.* (2008b), who not only derive intertemporal preferences from standard choice tasks but also adopt an approach of using response times from these choices, i.e., how long it actually takes the subjects to choose between option A and option B. The authors assume that “*subjects should take longest to decide when the two options are most similar in their discounted values*” and therefore argue that the inference from response times should, in principle, work (Chabris *et al.*, 2008, p. 7). The results of Chabris *et al.* (2008) suggest that choice-based and response-time-based estimates are nearly identical in their setting.

3 The Dataset

The first step of a meta-analysis is the collection of primary studies. To this end, we search Google Scholar for the literature on discounting and then examine the references of the retrieved studies to search for other usable studies (this method is called “snowballing” in the meta-analysis context). We use Google Scholar because it provides powerful fulltext search. Specifically, we employ the following query: `discount method experiment ‘discount rate’ OR ‘discount factor.’` The query is designed to yield the well-known experimental studies on discounting among the first hits, while being sufficiently inclusive. We go through the first 300 studies returned by the search and examine the abstract of each paper. If the abstract suggests at least a remote possibility that the paper contains estimates of the discount rate, we download the paper and inspect it; this way we inspect 178 studies. Next, we collect the references of these studies and download the 30 papers that are most often quoted in the literature but are not returned by our baseline Google Scholar search.

We apply three inclusion criteria. Each study included in our dataset must be an experiment, either lab or field, and must report an estimate of the discount rate (or the discount factor in a way that allows re-computation to the discount rate). Next, we exclude estimates of the discount rate derived from very short delays (several hours)—these are extreme cases for which it is often difficult to find use in practice. Finally, we include only studies published in peer-reviewed journals. The major reason for the

last inclusion criterion is feasibility, but we also hope that peer review sets a bar for quality. Moreover, journal articles generally contain fewer typos and other mistakes in the presentation of results compared to unpublished manuscripts

Table 1: Studies used in the meta-analysis

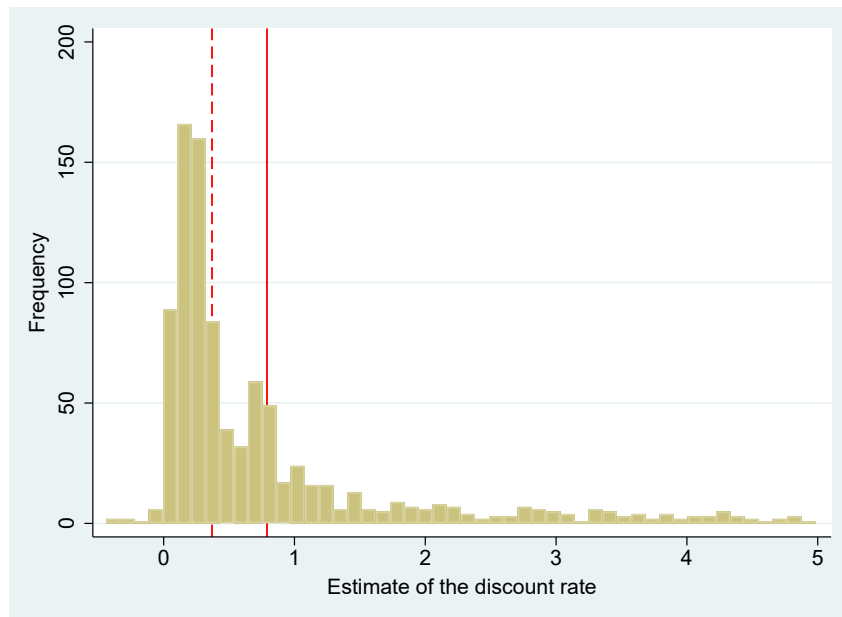
Abdellaoui <i>et al.</i> (2010)	Castillo <i>et al.</i> (2011)	Ifcher & Zarghamee (2011)
Andersen <i>et al.</i> (2006)	Chabris <i>et al.</i> (2008)	Kirby & Marakovic (1995)
Andersen <i>et al.</i> (2008)	Chabris <i>et al.</i> (2009)	Kirby & Marakovic (1996)
Andersen <i>et al.</i> (2010)	Chapman & Elstein (1995)	Kirby <i>et al.</i> (1999)
Andersen <i>et al.</i> (2013)	Chapman & Winquist (1998)	Loewenstein (1987)
Andersen <i>et al.</i> (2014)	Chapman (1996)	McClure <i>et al.</i> (2007)
Andreoni & Sprenger (2012)	Chapman <i>et al.</i> (1999)	Meier & Sprenger (2010)
Andreoni <i>et al.</i> (2015)	Chesson & Viscusi (2000)	Meier & Sprenger (2013)
Attema <i>et al.</i> (2016)	Coller & Williams (1999)	Meier & Sprenger (2015)
Bauer & Chytilova (2010)	Deck & Jahedi (2015a)	Newell & Siikamaki (2015)
Bauer & Chytilova (2013)	Deck & Jahedi (2015b)	Olivola & Wang (2016)
Bauer <i>et al.</i> (2012)	Dolan & Gudex (1995)	Read & Read (2004)
Benzion <i>et al.</i> (1989)	Duquette <i>et al.</i> (2012)	Sutter <i>et al.</i> (2013)
Booij & van Praag (2009)	Field <i>et al.</i> (2013)	Tanaka <i>et al.</i> (2010)
Brown <i>et al.</i> (2009)	Finke & Huston (2013)	Thaler (1981)
Burks <i>et al.</i> (2012)	Hardisty <i>et al.</i> (2013)	Voors <i>et al.</i> (2012)
Cairns & van der Pol (1997)	Harrison <i>et al.</i> (2002)	Warner & Pleeter (2001)
Carlsson <i>et al.</i> (2012)	Harrison <i>et al.</i> (2010)	Zauberman <i>et al.</i> (2009)
Cassar <i>et al.</i> (2017)	Hausman (1979)	

We terminate the search for studies on January 15, 2020. Our final dataset covers 56 studies comprising 927 estimates of the discount rate. Of these, 715 were reported explicitly as discount rates, and the remaining 212 estimates were reported as discount factors that we recomputed to rates according to the corresponding discounting formulas. All discount rates are annualized. The oldest study in our sample was published in 1979,⁴ and our meta-analysis thus spans four decades of research in the area. An overview of primary studies included in the meta-analysis is presented in Table 1; the full dataset (together with estimation codes for R and Stata) is available in an online appendix at meta-analysis.cz/discrate. We follow the reporting guidelines for meta-analysis compiled by Havranek *et al.* (2020).

Apart from the key variables for our analysis—the estimated discount rate and its standard error—we codify additional explanatory variables to control for the sources of variation in our data sample. We control for the length of the time horizon presented to the subjects, i.e., the delay of the experimental task. Moreover, we include a dummy variable describing whether the reported estimate relates to hyperbolic or exponential discounting. We further control for whether the study employs front-end delay; if it

⁴The oldest paper we use is Hausman (1979), which is not an experiment in the strict sense but is still based on real choices. The paper estimates discount rates from trade-offs between upfront capital costs and future savings of operating costs, looking at purchasing decisions of air conditioners.

Figure 1: Histogram of discount rate estimates




Notes: The figure depicts a histogram of annualized discount rate estimates reported by individual studies. Extreme values are omitted from the graph but included in all regressions. The solid line denotes the sample mean; the dashed line denotes the sample median.

is performed in the lab or in the field; if payoffs used in the study are hypothetical or real, i.e., paid out at the end of the experiment; what the stakes of the experiment are in terms of the maximum payoff related to median personal expenditure; which elicitation method (choice, matching, and rating) and domain (money, health, and others) is used to identify the estimate; and whether the framing of the task is positive (gaining), negative (losing) or neutral. We also control for the characteristics of the subject pool: whether it contains students or a more general sample of the population; the gender of the subjects it includes (exclusively males, females, or both); and the continent from which the subject pool was drawn. Additionally, we control for study age and the number of Google Scholar citations weighted by the number of years since the first version of the study appeared in Google Scholar. We describe these variables in more detail in Section 5, which also includes the corresponding Bayesian model averaging analysis.

The estimated discount rates in our dataset have a mean of 0.80 and a standard deviation of 0.97. A histogram of the estimates is presented in Figure 1: the distribution is apparently skewed, with a median value of 0.37. Negative values of the discount rate estimates are rare, though present, and often the matter of negative framing (for example, choosing to pay a fine or experience an illness now rather than later). The distribution thus offers several outliers on both sides. We address the potential influ-

Figure 2: Within- and between-study variation of discount rate estimates



Notes: The figure shows a box plot of annualized discount rate estimates reported in individual studies. Extreme values are omitted from the graph but included in all regressions. N = the number of estimates reported in the study.

ence of these outliers on our analysis by winsorizing at the 5% level (the main results are robust to changes in the winsorization level; without winsorizing, the minimum reported discount rate is -0.4 , the maximum is 13.7).

To be able to employ modern meta-analysis methods, we need measures of precision for individual estimates. Nevertheless, the standard errors of the discount rate estimates are reported only for 539 of the 927 estimates in our dataset. Researchers in the field sometimes mention that the discount rates they report are large and robust to various changes in the specifications, which constitutes the implicit apology for not reporting precision. As a robustness check (available on request and in the working paper version of this article), we exclude these studies from the dataset and focus only on those for which standard errors can be obtained directly. However, doing so reduces the power of our estimations and does not affect our main results. Therefore, in the baseline case, we also use studies that do not report precision explicitly. To approximate precision at least at the study level, we apply the bootstrap resampling technique. We then combine the explicitly reported standard errors with the standard errors obtained by bootstrapping at the study level. The substantial within- and between-study heterogeneity of discount rate estimates, the rationale for a meta-regression analysis, is apparent from Figure 2.

4 Publication Bias

The selective reporting of some estimates (typically those that are intuitive and statistically significant) has been identified as a serious threat to the credibility of empirical economics (Ioannidis *et al.*, 2017).⁵ When estimation noise is large, and therefore standard errors are large, researchers have incentives to preferentially report large point estimates that become statistically significant. McCloskey & Ziliak (2019) liken selective reporting to the Lombard effect, in which speakers increase their vocal effort in the presence of noise. Selective reporting (which is conventionally called publication bias but is not confined to published papers) thus manifests as a correlation between point estimates and their standard errors.

The general prior among economists and psychologists is that the discount rate is positive. People are impatient; they value the present more than the future. In contrast, a negative estimate of the discount rate means that an individual is willing to

⁵Other recent papers documenting the extent of publication bias in various fields in economics include Blanco-Perez & Brodeur (2020); Brodeur *et al.* (2016, 2020a); Campos *et al.* (2019); Doucouliagos & Paldam (2011); Duan *et al.* (2020); Geyer-Klingenberg *et al.* (2019); Havranek (2010); Havranek & Irsova (2010); Havranek *et al.* (2015b); Havranek & Kokes (2015); Nelson & Moran (2020); Tokunaga & Iwasaki (2017); Ugur *et al.* (2018, 2020); Valickova *et al.* (2015); Xue *et al.* (2020); Zigraviova & Havranek (2016).

accept an offer in the future with a lower value than what is available now, indicating an extraordinary preference for such a state of the world. Negative (and positive but insignificant) estimates are rare in our sample but do occur, which suggests that any potential publication bias in the literature is occasional and not universal. We do not claim that the average discount rate should be zero or even negative. However, the crux of the publication bias problem is the following: with sufficient imprecision and liberal elicitation techniques, we always obtain insignificant or negative estimates from time to time. For the same reason we also obtain large positive estimates. If negative and zero findings are often discarded (they are obviously implausible), while large positive estimates are often retained (it is less obvious whether they are far from the true value), the literature as a whole presents distorted results. The typical reported estimate is biased upwards.

The idea of publication bias is illustrated by Figure 3, the so-called funnel plot (Egger *et al.*, 1997). The horizontal axis depicts the magnitude of the estimate, while the vertical axis depicts the estimate's precision. With no publication bias, the most precise estimates should be close to the underlying average effect. With decreasing precision, we obtain increasing dispersion, which creates the shape of an inverted funnel. However, in the absence of publication bias, there is no reason for asymmetry in the funnel. If, in contrast, imprecise negative estimates are discarded but imprecise large positive estimates are reported, we obtain asymmetry—which is precisely what we see from the figure. The funnel plot can thus serve as a visual check of publication bias (Stanley & Doucouliagos, 2010; Rusnak *et al.*, 2013).

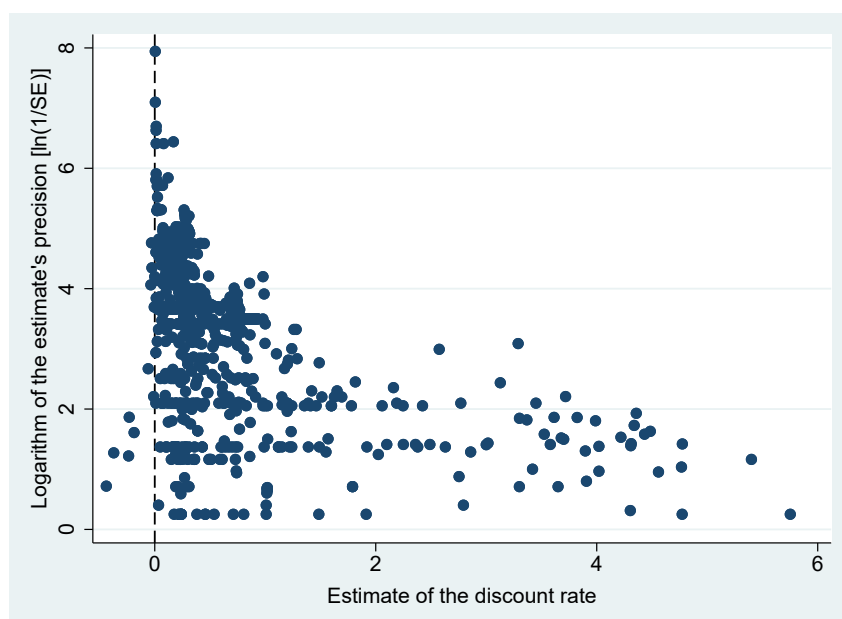
Next, we examine the correlation between the discount rate estimates and their standard errors quantitatively to test for the presence of publication bias (the so-called funnel asymmetry test, Egger *et al.*, 1997):

$$\hat{\delta}_{ij} = \delta_1 + \gamma_1 \cdot SE(\hat{\delta}_{ij}) + u_{ij}. \quad (5)$$

Here, the $\hat{\delta}_{ij}$ is the i -th estimate of the discount rate from the j -th study, $SE(\hat{\delta}_{ij})$ is the corresponding standard error, γ_1 measures publication bias, and δ_1 is the mean discount rate corrected for the bias; u_{ij} is a disturbance term. The first part of Table 2 shows the results of the funnel asymmetry test; we always cluster standard errors at the study level. The first column in the table shows a simple OLS regression; the second column presents a weighted least squares specification (with precision as the weight) which addresses the apparent heteroskedasticity of Equation 5.

The results presented in Panel A of Table 2 are consistent with the finding of publication bias: the correlation between estimates and standard errors is statistically significant at least at the 10% level in both specifications and the corrected mean is smaller than the simple uncorrected mean (0.26–0.52 vs. 0.80). But, as Stanley & Doucouliagos

Figure 3: Funnel plot suggests publication bias



Notes: The figure depicts the funnel plot of annualized discount rate estimates. Extreme values are omitted from the graph but included in all regressions.

(2014) show, while the linear funnel asymmetry test is a valid tool for testing the presence of publication bias, it is not a good estimator of the underlying corrected mean. The reason is that selective reporting is a more complex function of the standard error, and Monte Carlo simulations have shown that a linear approximation does not suffice (Stanley, 2008). For this reason, in Panel B of Table 2 we employ more advanced non-linear techniques.

The first non-linear technique presented in Table 2 is the Weighted Average of Adequately Powered estimates (WAAP) due to Ioannidis *et al.* (2017). The technique computes the statistical power of each estimate and uses only those whose power exceeds 80%. From these “adequately powered” estimates Ioannidis *et al.* (2017) compute a weighted average with weights proportional to the precision of the estimate. From this technique we obtain a mean discount rate of 0.33, which lies between the two estimates we obtained in Panel A (but as we have noted, estimates of the underlying effect derived from linear models in Panel A are not reliable). The second non-linear approach we use is the stem-based technique by Furukawa (2020). The “stem” in the title of the methods refers to the stem of the funnel plot; the technique focuses on the most precise estimates. It follows Stanley *et al.* (2010), who suggest that “discarding 90% of the [most imprecise] published findings greatly reduces publication selection bias and is often more efficient than conventional summary statistics.” (Stanley *et al.*, 2010, p. 70).

Table 2: Funnel asymmetry tests indicate publication bias

PANEL A: Linear models				
	OLS		Precision	
Standard error (<i>publication bias</i>)	0.535 ^{***} (0.0299)		1.031 ^{**} (0.449)	
Constant (<i>effect beyond bias</i>)	0.518 ^{***} (0.114)		0.259 ^{***} (0.0373)	
Observations	927		927	
PANEL B: Non-linear models				
	WAAP of Ioannidis <i>et al.</i> (2017)	Stem-based method of Furukawa (2020)	Selection model of Andrews & Kasy (2019)	Endogenous kink of Bom & Rachinger (2019)
Effect beyond bias	0.331 ^{***} (0.0131)	0.282 ^{***} (0.00915)	0.252 ^{***} (0.0140)	0.145 ^{***} (0.00321)
Observations	927	927	927	927

Notes: The table reports the results of regression $\delta_{ij} = \delta_1 + \gamma_1 \cdot SE(\delta_{ij}) + u_{ij}$, where δ_{ij} denotes the i -th annualized discount rate estimated in the j -th study, and $SE(\delta_{ij})$ denotes its standard error. Panel A shows estimation by OLS and weighted least squares where estimates are weighted by precision, the inverse of their standard error. Panel B shows the recently developed non-linear estimation techniques; WAAP stands for the Weighted Average of the Adequately Powered estimates. Standard errors, clustered at the study level, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Instead of discarding an arbitrary portion of estimates, which is generally suboptimal, Furukawa (2020) optimizes the trade-off between efficiency (which decreases when estimates are discarded) and bias (which increases when less precise estimates are included). The cut-off percentage is thus determined endogenously in the model, and in our case it yields an estimate of 0.28 for the mean discount rate.

The third non-linear technique is the selection model developed by Andrews & Kasy (2019). The selection model assumes that the probability of publication changes abruptly after reaching pre-defined thresholds for the t-statistic (in our case: 0, 1.65, 1.96, 2.33). The technique then computes how much estimates from each bracket are over- or under-represented in the literature, and re-weights them accordingly. The selection model gives us an estimate of 0.25 for the mean discount rate. Finally, the fourth non-linear specification we employ is the Endogenous Kink technique introduced recently by Bom & Rachinger (2019). The logic of the estimator is similar to both the linear funnel asymmetry test and the stem-based technique by Furukawa (2020): it also assumes that highly precise estimates are unbiased, but fits the publication bias function using two linear segments. The first segment is horizontal (no bias, therefore no relation between estimates and standard errors for the most precise estimates) and the second segment has a slope equal to the correlation between estimates and standard error for less precise estimates. Bom & Rachinger (2019) show how the “kink”

(that is, the point where both segments join) can be identified. The technique yields an estimate of 0.15 for the mean discount rate.

In sum, Table 2 gives us significant estimates for publication bias (Panel A) and estimates of the corrected mean discount rate in the range 0.15–0.33 (Panel B). We prefer to focus on the most conservative estimate from Panel B, 0.33. These results indicate that publication bias exaggerates the mean reported discount rate more than twofold, from 0.33 to 0.80 (the simple uncorrected mean). But again we have to note that our results hinge on the assumption that in the absence of publication bias there is no correlation between estimates and standard errors; even the selection model by Andrews & Kasy (2019) uses this assumption for identification. There are two reasons why the assumption might not hold in the case of the discounting literature, and we thank two anonymous referees of this Journal for articulating the reasons. First, researchers are likely to design the experiment in a way that is tuned to detect discount rates near zero and does not uniformly cover the entire interval of possible rates. Consequently, smaller discount rates are likely to be measured with greater precision, and thus the correlation between estimates and standard errors can arise even in the absence of publication bias. Second, negative estimates of the discount rate can be missing from the literature simply because elicitation techniques used by the researchers do not allow for negative values: for instance, if experimental subjects are always offered a larger sum of money in the future compared with the immediate option.⁶

While we see no bulletproof way how to measure the quantitative importance of these two caveats for our results, a useful exercise is to conduct a caliper test inspired by Gerber & Malhotra (2008) and Brodeur *et al.* (2020b). Caliper tests are typically employed to identify a systematic break related to publication bias at a particular psychologically important threshold (such as 0 for the point estimate or 1.96 for the t-statistic). For example, Brodeur *et al.* (2020b) show how, for many quasi-experimental techniques commonly used in economics, estimates that are just significant at the 5% level (that is, have t-statistics slightly larger than 1.96) are more likely to get published than estimates that are just insignificant. The essence of the caliper test is thus to compare the number of estimates just below and just above a particular threshold: given a sufficiently narrow caliper, there should be no difference. In this paper we use a different tactic and employ calipers of varying width to constrain our baseline linear

⁶If the correlation between estimates and standard errors is driven by this second caveat, certainly it cannot be called publication bias. (The issue is also discussed by Nobel *et al.*, 2020, p. 11.) But it can still represent another type of research bias that should be corrected in meta-analysis: suppose an extreme case in which the true discount rate is negative. If an experiment does not consider the possibility of negative discounting, it will inevitably produce estimates biased upwards. A similar bias will arise on average in a more plausible setting in which the true discount rate is positive but small, because most elicitation designs will allow large positive outliers, but not negative ones.

Table 3: Caliper tests for different ranges of discount rate estimates

Caliper test for $\delta \in \langle -0.5, 0.5 \rangle$	OLS	Precision
Standard error (<i>publication bias</i>)	0.0919** (0.0367)	0.473** (0.190)
Constant	0.214*** (0.0139)	0.184*** (0.0188)
Observations	538	538
Caliper test for $\delta \in \langle -1, 1 \rangle$	OLS	Precision
Standard error (<i>publication bias</i>)	0.205*** (0.0398)	0.949** (0.409)
Constant	0.325*** (0.0444)	0.232*** (0.0313)
Observations	717	717
Caliper test for $\delta \in \langle 0.25, 0.75 \rangle$	OLS	Precision
Standard error (<i>publication bias</i>)	0.0835** (0.0395)	0.536* (0.288)
Constant	0.429*** (0.0351)	0.371*** (0.0428)
Observations	313	313
Caliper test for $\delta \in \langle 0.5, 1.5 \rangle$	OLS	Precision
Standard error (<i>publication bias</i>)	0.125*** (0.0126)	0.199** (0.0786)
Constant	0.801*** (0.0295)	0.764*** (0.0341)
Observations	244	244

Notes: The table reports the results of regression $\delta_{ij} = \delta_1 + \gamma_1 \cdot SE(\delta_{ij}) + u_{ij}$, where δ_{ij} denotes the i -th annualized discount rate estimated in the j -th study, and $SE(\delta_{ij})$ denotes its standard error. The regressions only include estimates within the bounds indicated by the caliper. The table shows estimation by OLS and precision weighting. Standard errors, clustered at the study level, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

regression (of estimates on their standard errors) in an attempt to address the important caveats mentioned earlier.

We use two groups of calipers. First, we focus on small estimates, both positive and negative. If the correlation between estimates and standard errors persists when large positive outliers are excluded, the finding of publication bias is probably not fully driven by the clustering of tradeoffs near zero typical in the discounting literature. Second, we focus on positive estimates approximately around the mean and median of the reported discount rates. If the correlation between estimates and standard errors persists when only safely positive estimates are considered, the finding of publication bias is not fully driven by the impossibility of negative discount rates in many

experimental designs. The results of caliper tests of funnel asymmetry are shown in Table 3. Note that here we cannot interpret the means corrected for publication bias (the constant in the regression), because the calipers are arbitrary slices of the data. We can interpret the slopes in this regression, and they all suggest a positive correlation between estimates and standard errors. It is important to point out, however, that we still have to assume that the standard error is exogenous within individual calipers. If there is a mechanical relationship between the estimates and standard errors within calipers in the absence of publication bias, caliper tests fail to address the two caveats.

Another way to approach this problem is to use techniques that do not need the assumption of zero correlation between estimates and standard errors in the absence of publication bias—or, in the case of one technique, at least not between studies. Table 4 shows the corresponding results. In the first column we apply p-uniform*, a brand new technique to test publication bias and estimate the corrected mean. The technique was developed by van Aert & van Assen (2020) for psychology, but it can be applied to an experimental economics setting as well. (In fact, it is probably better suited to experimental economics than the traditional publication bias tests that are designed to aggregate regressions.) At the basis of p-uniform* lies the statistical principle that p-values should be uniformly distributed *at the mean underlying effect size*: i.e., when testing the hypothesis that the estimated coefficient equals the underlying value of the effect (not necessarily zero). The reported t-statistics and p-values, of course, in almost all cases correspond to tests that relate the estimated coefficient to zero. It follows that if the reported p-values are uniformly distributed, the literature is consistent with a zero underlying effect. The idea of p-uniform* is to find a coefficient at which the distribution of p-values is uniform; this is done by recomputing p-values for various potential values of the underlying effect and then comparing the resulting distribution to the uniform one. Similarly the technique's test for publication bias evaluates whether p-values are uniformly distributed at the simple mean reported in the literature. Technical details and more discussions are available in van Aert & van Assen (2020). The results in Table 4 show evidence of publication bias significant at the 1% level. The mean corrected discount rate is small (0.18) but imprecisely estimated.

In the second column of Table 4 we use the inverse of the square root of the number of observations as an instrument for the standard error following Havranek (2015) and Astakhov *et al.* (2019): some method choices in the primary studies can influence both the discount rate and the standard error, which would make our OLS results spurious. (There can also exist a more direct mechanical relationship between estimates and standard errors, as we discussed in the context of the caliper test.) The number of observations is a natural instrument, because it correlates with the standard error

Table 4: Relaxing the exogeneity assumption

	p-uniform*	Instrument	Fixed effects
Publication bias	YES*** (0.007)	0.316* (0.183)	0.875*** (0.0154)
Effect beyond bias	0.176 (0.663)	0.633*** (0.158)	0.341*** (0.00806)
Observations	927	927	927

Notes: In the first column the table reports the results of the p-uniform* test for publication bias developed by van Aert & van Assen (2020); p-values are reported in parentheses. For the remaining two specifications, which show regressions along the lines of the first panel of Table 2, standard errors are reported in parentheses and are clustered at the study level. The second column reports an instrumental variable specification (where the instrument for the standard error is the inverse of the square root of the number of observations in a study), and the third column reports a study-level fixed effects specification. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

by definition. Nevertheless, while not the product of the estimation technique (in contrast to the standard error), in the studies estimating the discount rate the number of observations can be still correlated with the choice of the technique. Therefore the instrumental variable technique cannot be expected to fully address the exogeneity problem. The results in Table 4 indicate publication bias significant at the 10% level and an underlying mean discount rate of 0.63. Finally, in the last column of the table we explore whether publication bias appears within studies. This specification still needs the exogeneity condition to hold within individual studies, but relaxes it between studies as the latter source of variation in discount rate estimates is not used. Once again we obtain evidence of publication bias, now significant at the 1% level, and underlying mean effect smaller than the uncorrected simple mean (0.34 vs. 0.8). Overall we prefer this fixed effects estimation because it is simple, elegant, and its results are consistent with the most conservative non-linear technique presented earlier.

The Appendix harbors four sets of further robustness checks. First, in Table A1 we cluster standard errors at the level of authors instead of studies. Several researchers have co-authored many of the studies in our dataset, and consequently the results of these studies do not have to be independent of each other. We have identified 31 clusters for which no co-authors overlap. The results are almost identical to the baseline case, with the exception of the IV specification, in which we lose statistical significance. Second, in Table A3 we exclude estimates for which the discounting model is not explicitly specified. Once again the results are similar, but we obtain smaller estimates of the mean discount rate corrected for publication bias.

Third, in Table A4 we run funnel asymmetry tests with the discount rate in the absolute value. Aside from the standard error, on the right-hand side we include the interaction of the standard error and a dummy variable that equals one for negative

values. In consequence, this specification reveals different mechanisms of selective reporting for positive and negative estimates. For positive estimates, our findings are consistent with publication probability increasing with an increasing t-statistic. For negative estimates, our findings are consistent with the opposite: insignificant negative estimates tend to be easier to publish, probably because they are more feasible. Fourth, in Table A2 we investigate how publication bias differs between medians and means of individual-specific discounting. To this end, we include an interaction of the standard error and a dummy variable that equals one for median estimates. Medians comprise 15% of the data set, and the results of the table show mixed findings. According to most techniques, there is little difference in the extent of publication bias between means and medians. Our preferred fixed effects specification, however, indicates that median estimates are substantially less biased than mean estimates.

5 Heterogeneity

The substantial differences in the estimates of the discount rate reported in the experimental literature have already been stressed by several previous studies (Frederick *et al.*, 2002; Percoco & Nijkamp, 2009; Andersen *et al.*, 2014; Cheung, 2016). As Frederick *et al.* (2002, p. 352) puts it: “While the discounted utility model assumes that people are characterized by a single discount rate, this literature reveals spectacular variation across (and even within) studies.” Figure 2 shows strong differences in the results at the study level. In this section we try to explain the differences by regressing the estimated discount rates on their standard errors together with 21 additional explanatory variables that reflect observable variation in the context in which researchers obtain the estimates. We start from the linear model of publication bias, which is the reason why we retain the standard error variable in the regression. Therefore the second goal of this section is to find out whether our previous findings concerning publication bias prove robust to controlling for heterogeneity.

The first option for estimating such an extended model is simply running a regression with all the collected variables. The problem is that not all the variables are equally important; some are probably redundant, and including all variables would substantially diminish the precision of our point estimates for the effects of the important variables. However, we do not know *ex ante* which variables are redundant. A common approach would be to eliminate potential redundant variables in a step-wise fashion (sequential t-tests); but in doing so, we can never be sure that we have arrived at the best underlying model. Furthermore, the theory can help us stress some particular variables, but we still do not want to completely ignore the remaining ones. In

other words, we face extensive model uncertainty, which is a typical feature of meta-regression analysis. The formal response to model uncertainty in the Bayesian setting is Bayesian model averaging (Raftery *et al.*, 1997), our first method of choice.

Bayesian model averaging (BMA) tackles the problem of uncertainty by estimating models with all possible combinations of explanatory variables in the dataset⁷ and constructing a weighted average over the estimated coefficients across all these models. The weights used for averaging stem from posterior model probabilities derived from Bayes' theorem and are analogous to information criteria in frequentist econometrics. Posterior model probabilities (PMPs) measure how well the particular model fits the data, conditional on model size. BMA produces posterior inclusion probability (PIP) for each variable, which is the sum of the posterior model probabilities for the models in which the variable is included. Recent applications of Bayesian model averaging in meta-analysis include, for example, Irsova & Havranek (2013); Babecky & Havranek (2014); Havranek & Irsova (2017); Cazachevici *et al.* (2020). More details on BMA, including a formal derivation, can be found in Raftery *et al.* (1997) or Eicher *et al.* (2011).

The application of BMA, however, is not straightforward since estimating the millions of possible model combinations is infeasible. A solution is to approximate the whole model space by applying the Markov chain Monte Carlo algorithm that walks only through the models with high posterior model probabilities (Madigan *et al.*, 1995). For approximation we use the BMS package for R developed by Zeugner & Feldkircher (2015). Bayesian model averaging is sensitive to the estimation framework, particularly to the use of *priors* representing the researcher's prior beliefs on the probability of each model (the model prior: how much confidence we place in the prior that, for example, all models have the same probability) and regression coefficients (Zellner's g-prior: how much confidence we place in the prior that, for example, all regression coefficients are zero). In the baseline specification, we follow the agnostic (that is, easily overridden by data) priors supported by Eicher *et al.* (2011), who show that these intuitive priors yield the best predictive performance: the unit information prior (UIP) for Zellner's g-prior, which assigns the prior that coefficients are zero the same weight as one observation of data, and the uniform model prior, which gives each model the same prior probability, irrespective of the number of variables included in the model. On top of the uniform model prior we use the dilution prior suggested by George (2010). In this prior the relative weight of each model is further multiplied by the determinant of the correlation matrix of the variables included in

⁷If the matrix of explanatory variables X contains K potential variables, this means estimating 2^K variable combinations, i.e., 2^K models. This estimation results in $2^{22} = 4,194,304$ models in our case.

the model. The dilution prior is designed to address collinearity: models with high collinearity will have small determinants of the correlation matrix, and therefore little weight in our implementation of BMA.⁸

5.1 Variables

The explanatory variables we have collected are listed in Table 5; we include the description of each variable, its mean, standard deviation, and the mean weighted by the inverse of the number of estimates reported per study, which effectively equalizes the impact each study has on the statistics. For ease of exposition, we divide the explanatory variables into 4 categories: estimation characteristics, experimental characteristics, subject pool characteristics, and publication characteristics.

Estimation Characteristics

The variation among the reported discount rate estimates can stem from the theoretical assumptions of the intertemporal choice model used in the experimental task presented to subjects, that is, mainly from the type of the discounting model and the time horizon that subjects face in their decision. The studies included in our dataset use the hyperbolic discounting model most frequently (373 observations; 40% of the data), followed by the exponential discounting model (133; 14%). Special cases of discounting models such as exponential mixture share, quasi-hyperbolic discounting, or mixed general model occur rarely in our dataset. Due to a lack of information reported in primary studies, we cannot identify the precise type of the discounting model in some of the cases and use this “unidentified” group as a reference category. The time horizon of the decisions presented to the subjects spans from one week to 50 years, while the mean value is 4.07 years. We also take into account whether the study uses front-end delay. With front-end delay the immediate option is shifted to the future, thereby imposing transaction costs on the instant payoff. Last but not least, we control for the general estimation setup—that is, whether the study employs a controlled laboratory experiment or a field experiment.

⁸A robustness check using the BRIC g-prior suggested by Fernandez *et al.* (2001) and the beta-binomial model prior according to Ley & Steel (2009) can be found in Appendix B; our main results would not change if we opted for this alternative set of priors. A detailed discussion of the priors used in the robustness checks is beyond the scope of the paper; for more details, see Zeugner & Feldkircher (2015). For example, the beta-binomial model prior gives the same weight to each model size (a certain number of variables included in the model), not the same weight to each model. The reason is that moderate model sizes are over-represented: there are many models that have 2^{10} variables, but only one model that has 2^{22} variables.

Table 5: Description and summary statistics of regression variables

Variable	Description	Mean	SD	WM
Discount rate	The reported estimate of the discount rate.	0.798	0.973	0.710
Standard error	The standard error of the discount rate estimate.	0.522	1.149	0.214
<i>Estimation characteristics</i>				
Hyperbolic discounting	= 1 if the discounting type is hyperbolic.	0.402	0.491	0.368
Exponential discounting	= 1 if the discounting type is exponential.	0.143	0.351	0.199
Delay	The logarithm of the time horizon of the task.	-0.255	2.222	-0.782
Front-end delay	= 1 if the immediate option is shifted to the future, thereby imposing transaction costs on the instant payoff.	0.338	0.473	0.364
Lab experiment	= 1 if a controlled laboratory experiment is used instead of a field experiment.	0.650	0.477	0.549
<i>Experimental characteristics</i>				
Real reward	= 1 if the reward subjects received is real instead of hypothetical.	0.629	0.483	0.754
Matching task	= 1 if matching is used for elicitation.	0.243	0.429	0.149
Health domain	= 1 if the experiment concerns health questions.	0.055	0.228	0.055
Other domain	= 1 if the experiment concerns questions other than health or money (such as vacation or a kiss from a movie star).	0.082	0.274	0.100
Negative framing	= 1 if the framing of the experimental task is presented as negative, i.e., "losing."	0.086	0.281	0.072
Neutral framing	= 1 if the framing of the experimental task is presented as neutral.	0.023	0.149	0.031
Stakes	The ratio of the logarithm of the highest payoff possible in the experiment to the logarithm of the median monthly expenditure in the country where the experiment was conducted.	0.817	0.373	0.753
<i>Subject pool characteristics</i>				
Sample size	The logarithm of the sample size used for the experiment.	4.889	1.617	5.035
Students	= 1 if the subject pool consists of students only.	0.528	0.500	0.445
Males only	= 1 if the subject pool contains males only.	0.029	0.168	0.027
Females only	= 1 if the subject pool contains females only.	0.030	0.171	0.054
North America	= 1 if the experiment is conducted in North America.	0.588	0.492	0.589
Asia	= 1 if the experiment is conducted in Asia.	0.058	0.234	0.107
Africa	= 1 if the experiment is conducted in Africa.	0.030	0.171	0.036
<i>Publication characteristics</i>				
Citations	The logarithm of the number of citations the study received in Google Scholar normalized by the number of years since the first draft of the study appeared in Google Scholar.	2.691	1.278	2.776
Publication year	The standardized publication year of the study.	0.000	1.001	0.283

Notes: SD = standard deviation, WM = mean weighted by the inverse of the number of estimates reported per study. The variable *Stakes* is only available for 777 observations; statistics for all other variables are calculated using the full sample of 927 observations. Data on median expenditure are obtained from World Bank (2020).

Experimental Characteristics

The results of any experiment can be affected by procedural subtleties. The second set of explanatory variables therefore comprises experimental and behavioral characteristics of the task presented to the subject pool. Psychological research suggests that there should be no systematic difference observed between real and hypothetical payoffs in discounting experiments (Johnson & Bickel, 2002; Kuhnberger *et al.*, 2002; Locey *et al.*, 2011). The recent literature, however, provides more ambivalent results stating that hypothetical conditions yield patterns of discounting that mirror those for real effort tasks, but these may change with repeated exposure to the decisions. The nature of the payoffs provided with the repetition of those tasks therefore needs to be taken into account when designing discounting studies (Malesza, 2019). We therefore control for this payoff effect by extracting the information on the nature of the reward from primary studies; 53% of the discount rates are computed for hypothetical payoffs. For a subsample of estimates, we are able to collect data on the size of the maximum payoff available in the experiment. We relate the maximum payoff size to World Bank data on household median monthly expenditure in the country, and the resulting variable is labeled “Stakes.” Note that this variable is not included in the baseline model, because doing so would imply disregarding all the observations for which the variable is not available.

Following the reasoning of Frederick *et al.* (2002) and others, we control for the variation in the estimates caused by the elicitation method used in the experiment. We include a dummy variable for matching tasks, taking choice tasks as the reference category present in 66% of cases. An important behavioral aspect of the corresponding task is represented by the domain over which the intertemporal decision is made. The majority of observations utilize monetary payoffs (87%); we therefore use them as the natural reference category in this regard. We codify the remaining domains by using dummy variables, distinguishing between the health domain and other domains—typically, more exotic ones (e.g. vacation, certificate, or a kiss from a movie star).

The design of any experiment is seldom immune to the issues of framing effects that refer to the finding that subjects often respond differently to different descriptions of the same problem (Tversky & Kahneman, 1981). The majority of discounting tasks are presented (framed) as positive decisions, e.g., choices between an amount of money today and a greater amount tomorrow (89.1%). There are, however, also negative framings of the tasks present in our dataset (8.6%). For example, Chapman & Winqvist (1998) and Hardisty *et al.* (2013) use monetary losses in their experiments. Other studies with negative framing operate with the health domain (Dolan & Gudex, 1995; Read & Read, 2004). Neutral framing applies for only 2.3% of the observations.

Subject Pool Characteristics

We describe the subject pool characteristics of an individual study by several variables. First, we control for the size of the subject pool by coding the number of subjects used for deriving the estimate; the mean is 271. Second, we control for the composition of the subject pool by incorporating dummy variables reflecting whether the pool consists exclusively of male or female subjects. The majority of studies, however, use non-exclusive subject pools consisting of both males and females (94.1%).


A general concern of any experimental study is its external validity, i.e., the extent to which its results can be generalized to other situations. Economic experiments are often criticized for using university students (typically economics majors) as experimental subjects—a pool of people with specific characteristics not always generalizable to the whole population (Marwell & Ames, 1981; Carter & Irons, 1991; Frank *et al.*, 1993). The behavior of decision makers recruited from natural markets has been examined in a variety of contexts, and it has typically not differed from that exhibited by more standard (and far less costly) student subject pools (Davis & Holt, 1993, p. 17).⁹ We control for the potential effect of a subject pool composed exclusively of student subjects. In addition, as recommended by an anonymous referee, we include an interaction of the student and lab experiment dummy variables. These two variables are correlated, because lab experiments often rely on students, and students, who are commonly familiar with lab experiments, may potentially behave differently in lab and field settings. Finally, the heterogeneity in the reported discount rates may stem from different cultural characteristics of populations. The primary studies do not give us much information to build on systematically, but at least we can control for continents out of which the subject pool was recruited. The majority of studies recruit subjects from European countries (32.4% obs.) and North America (58.8%). We also experimented with including dummy variables for each individual region, but doing so creates collinearity problems.

Publication Characteristics

We do not exclude any journal articles based on their supposedly poor quality, but we try to control for it—even poor-quality studies can bring useful information, especially if their results differ from those of high-quality studies. Some of the aspects related to quality are captured by the data and method characteristics described above. However, other quality aspects are surely more difficult to observe. Therefore we use two

⁹More recent evidence on differences between student and non-student samples is provided by Depositario *et al.* (2009).

Figure 4: Correlation matrix



Notes: The figure presents Pearson correlation coefficients for the variables reported in Table 5. Correlations for *Stakes* are computed using the 777 observations for which the variable is available. For all the other variables the figure shows correlations calculated at the full sample of 927 observations.

rough proxies: the age of the study and the number of citations. These are no perfect controls for quality, but other things being equal, newer and highly cited studies tend to be more reliable. For computing the age of the study we do not use the year of journal publication; due to different publication lags in different economics and psychology journals, such a measure would be of limited use. Therefore, we use the date of the first appearance of a draft of the paper in Google Scholar. For citations, we also rely on Google Scholar and compute the number of per-year citations that the primary study has obtained since the first draft appeared.

Figure 4 shows the correlations between the variables we consider. Several patterns emerge that are informative for understanding the types of experiments observed in the data. For example, lab experiments tend to use matching tasks with hypothetical rewards and rely on students. Recent and highly cited studies typically employ real

rewards. Recent studies are also less likely to use negative and neutral framing compared to older studies. Payoffs in experiments tend to be smaller when students are used.

5.2 Results

The results of the BMA estimation are visualized in Figure 5. The variables are displayed on the vertical axis and sorted by posterior inclusion probability. PIP can be thought of as a Bayesian analogy of statistical significance—we therefore see the most “significant” variables at the top of the figure. The horizontal axis denotes individual regression models sorted according to the posterior model probability, from left to right. The PMP represents how well the model fits the data relative to its size; the width of the columns is proportional to the PMP. The colors of individual cells denote the sign of the corresponding regression coefficients. Blue (darker in grayscale) depicts a positive sign, while red (lighter in grayscale) depicts a negative sign. Blank cells denote the exclusion of the variable from the given model.


The numerical results of BMA are reported in the left-hand panel of Table 6, which shows the posterior mean and standard deviation for each variable together with the posterior inclusion probability. Not counting the intercept, which is included by default in all models, eleven variables have PIPs above 50%: the standard error, the dummy for lab experiments, the dummy for health domain, the dummy for other (exotic) domains, the dummy for negative framing, sample size, the dummy for students in the subject pool, the interaction between student and lab experiment dummies, the dummy for subjects drawn from Asia, the dummy for Africa, and publication year. In the remainder of this subsection we will go through these results in more detail.

The first important result of the BMA analysis concerns publication bias. Standard errors are robustly correlated with the point estimates of the discount rate even when we control for 21 additional aspects of studies and estimates. The result corroborates our previous findings that the correlation is not spurious and does not result from an omission of factors that influence both the standard error and the point estimate. Moreover, both the posterior mean in BMA and the point estimate in the frequentist check suggest that the correlation is strong.

Results for Estimation Characteristics

An often-discussed factor potentially affecting the heterogeneity in discount rate estimates is the length of the delay over which the decision is made. This factor is inherently embedded as the parameter k in the discounted utility model presented in

Figure 5: Model inclusion in Bayesian model averaging



Notes: The response variable is the estimate of the discount rate reported in a primary study. The columns denote individual models; variables are sorted by posterior inclusion probability in descending order. The horizontal axis denotes cumulative posterior model probabilities. The estimation is based on the unit information prior recommended by (Eicher *et al.*, 2011) and the dilution prior suggested by George (2010), which takes into account collinearity. Blue color (darker in grayscale) depicts variables with a positive estimated sign. Red color (lighter in grayscale) depicts variables with a negative estimated sign. Variables with no color are not included in the given model. The numerical results of the BMA exercise are reported in Table 6.

Table 6: Explaining the heterogeneity in discount rate estimates

Variable:	Bayesian model averaging			Frequentist check (OLS)		
	Post. mean	Post. SD	PIP	Mean	SE	p-value
Constant	-0.244	NA	1.000	-0.253	0.163	0.126
Standard error	0.549	0.021	1.000	0.542	0.035	0.000
<i>Estimation characteristics</i>						
Hyperbolic discounting	0.039	0.062	0.352			
Exponential discounting	0.006	0.030	0.076			
Delay	0.000	0.002	0.041			
Front-end delay	0.014	0.041	0.143			
Lab experiment	0.155	0.101	0.776	0.222	0.091	0.018
<i>Experimental characteristics</i>						
Real reward	-0.005	0.027	0.077			
Matching task	0.017	0.046	0.161			
Health domain	0.345	0.088	0.993	0.356	0.076	0.000
Other domain	0.441	0.070	1.000	0.442	0.153	0.006
Negative framing	-0.148	0.106	0.734	-0.205	0.102	0.049
Neutral framing	0.003	0.031	0.046			
<i>Subject pool characteristics</i>						
Sample size	0.075	0.014	1.000	0.076	0.029	0.012
Students	0.877	0.111	1.000	0.901	0.223	0.000
Students * Lab experiment	-0.753	0.144	1.000	-0.813	0.239	0.001
Males only	0.013	0.052	0.090			
Females only	-0.001	0.023	0.041			
North America	0.012	0.041	0.127			
Asia	0.385	0.103	0.990	0.428	0.117	0.001
Africa	3.170	0.118	1.000	3.174	0.066	0.000
<i>Publication characteristics</i>						
Citations	-0.003	0.011	0.095			
Publication year	0.121	0.026	1.000	0.114	0.051	0.030
Observations	927			927		
Studies	56			56		

Notes: Response variable = annualized estimates of the discount rate. In the first specification from the left we employ Bayesian model averaging (BMA) using the unit information prior recommended by (Eicher *et al.*, 2011) and the dilution prior suggested by George (2010), which takes into account collinearity. The second specification, frequentist check (OLS), includes variables recognized by the BMA as having a posterior inclusion probability above 50%. Standard errors in the frequentist check are clustered at the study level. SD = standard deviation, PIP = Posterior inclusion probability, SE = standard error. All variables are described in Table 5.

Equation 1. According to the exponentially discounted utility theory, the values of all future outcomes should be discounted at a constant rate (Frederick *et al.*, 2002). Our results do not disagree: we find little systematic relationship between reported estimates of the discount rate and the length of the delay. This finding contrasts the results of, among others, Mazur (1984), who presents evidence for hyperbolic discounting, or, more recently Tsukayama & Duckworth (2010), who find that subjects discount rewards more steeply when they find the discounting domain particularly tempting. On the other hand, our results are in line with Andersen *et al.* (2014). A related effect is the importance of the dummy for exponential discounting, of which the constant discount rate is a key property. Our analysis suggests that tasks with exponential setups, i.e., with a constant discount rate between decisions with different delays, do not systematically differ from other studies in terms of the reported discount rates. Moreover, the estimates in our sample do not seem to be significantly different when hyperbolic discounting is applied.

Two additional results related to estimation characteristics are important. The first result is the low posterior inclusion probability and therefore the absence of the variable *Front-end delay* in most BMA models, which again contrasts many previous findings in the literature that front-end delay tends to decrease estimated discount rates (for example, Coller & Williams, 1999), but is consistent with the results of Andersen *et al.* (2014). A second important result is the difference between field and laboratory experiments. This finding suggests that a controlled laboratory environment produces more evidence for impatience than field study environment.

Results for Experimental Characteristics

Several studies find that individual discount rates are not very correlated across different domains such as money and health—this diversity is called domain independence. Cairns (1992), for example, estimates discount rates that are different for future health as compared to future wealth states; Chapman & Elstein (1995) demonstrate in two experiments that decision makers use different discount rates for health-related decisions and money-related decisions, with less patience for the health domain. See Loewenstein *et al.* (2003) for more examples of domain independence.

Our results suggest that people tend to be more impatient when the experiment concerns health than when it concerns money. It is difficult to transfer health states over time, so questions about health are, to some extent, similar to questions about money when liquidity constraints are binding (see Bleichrodt *et al.*, 2016). When liquidity constraints are present and binding, people cannot increase current consumption at the expense of consumption in the future. A high discount rate follows. In

addition, we also find that people tend to be more impatient when making their decisions in more exotic domains than money: holiday preferences, gift certificates, kisses from movie stars. Our results thus strongly corroborate domain independence.

Describing the estimation characteristics in Section 5, we referred to the literature suggesting there should be no difference whether real or hypothetical payoffs are used in discounting experiments. Our results confirm that it indeed does not matter whether the decision is made with fictive payoffs only. Real rewards do not systematically affect the estimates of the discount rate. Researchers can thus use hypothetical questions that have advantages in the elicitation of time preferences since hypothetical setting allows us to ask questions involving long time horizons and large payoffs (Wang *et al.*, 2016).

We find no substantial effect for some other experimental characteristics. Different experimental tasks do not bring substantially different results: matching does not seem to differ significantly from choice tasks, which suggests that the inference of an individual's discount rate by the matching method does not systematically outperform the interval elicitation provided by choice tasks. In contrast, the estimated discount rates are affected by framing, and negative framing is associated with smaller estimates. The result is consistent with Harris (2012) and Hardisty *et al.* (2013), among others, who stress the role of dread in intertemporal choices: it is itself aversive to wait for an aversive outcome, and for many subjects it is preferable to get it over with. Finally, we find that the stakes of the experiment (maximum possible payoff relative to personal expenditure) are associated with smaller reported discount rates. (Note that the BMA specification featuring this variable is included in Table B3 in the Appendix; the variable is not available for all observations, and thus is not included in the baseline BMA estimation.) The result is consistent with a large literature (for example, Thaler, 1981; Benzion *et al.*, 1989; Warner & Pleeter, 2001; Meyer, 2015), and a possible explanation is that non-monetary transaction costs of borrowing or saving that increase the discount rate may be relatively larger for smaller payments.

Results for Subject Pool Characteristics

The long-term debate over the external validity of the experiments performed on student samples is reflected in our analysis by the variable *Students*. Our results suggest that students make more impatient choices in discounting tasks than the general population, which is consistent with Harrison *et al.* (2002) and can be explained by the fact that students tend to be more liquidity-constrained. In contrast, the interaction between student and lab experiment dummies shows a negative coefficient: students that participate in laboratory experiments tend to display relatively little impatience.

This finding can be caused by several factors, out of which the standard argument would point to the self-selection of students into subject pools in laboratory experiments. The vast majority of lab experiments are conducted with university students majoring in economics, who have been shown, for example, to be more selfish than the general population (Marwell & Ames, 1981). Two types of hypotheses explain why this may be the case: 1) the selection hypothesis, according to which individuals concerned with economic incentives opt for economic studies, and 2) the learning hypothesis, which states that individuals studying economics learn behavioral patterns out of the theories and models they pursue (Carter & Irons, 1991). It might be true that not only more “selfish” individuals self-select into study fields such as economics but also that more patient students self-select into the roles of experimental subjects.

Our results provide some evidence that discount rates elicited from subject pools in Asia and Africa significantly differ from those obtained in other parts of the world. The Asian and (especially) African population is, according to our analysis, more impatient than the population of other continents. This result is in line with the results of the large cross-country study on time preferences by Wang *et al.* (2016, p. 17), who observe that “*Africa has the lowest percentage of participants choosing to wait (33%).*” The benchmark demographic area—Europe—seems to follow similar patterns of discounting as North America and display lower discount rates. Again, a possible explanation is related to liquidity constraints, which might be larger in Asia and Africa than in the West. Nevertheless, a disclaimer is in order: for Africa we only have two studies in our sample. Next, we also obtain evidence of an impact of the sample size on the discount rate estimates: large experiments seem to produce larger discount rates, though the effect is economically weak. Finally, neither exclusively male nor female subject pools report significantly different results of discount rates in our sample compared to the baseline (mixed) subject pools.

Results for Publication Characteristics

Out of the publication characteristics that we consider, the number of citations does not matter for the estimated discount rates, while publication year is positively associated with the estimates: other things being equal and on average, newer studies show more evidence for impatience. The age of the study can be considered a rough proxy for (unobserved) quality aspects that are not captured by the variables discussed earlier. There are certainly quality aspects that we do not control for, and an obvious solution is the addition of study-level fixed effects. We opt for the fixed-effects estimator in the previous section that focuses on publication bias, but here, it is not feasible: for many variables in which we are interested the within-study variation is very small.

5.3 Robustness Checks

In Appendix B we perform several different sensitivity checks in order to confirm whether our baseline BMA results presented earlier in this section are robust. First, we combine the reduction in model uncertainty resulting from BMA estimation with traditional frequentist estimation: in other words, we use a Bayesian technique for the selection of variables and a frequentist technique for estimation. The best model identified by the BMA exercise includes eleven explanatory variables (plus the intercept). These variables also have a posterior inclusion probability above 0.5 and therefore should, according to the classification by Kass & Raftery (1995), have a non-negligible impact on our response variable. We re-estimate this best BMA model using the standard OLS technique, clustering standard errors at the study level. The results of this estimation are provided in the right-hand panel of Table 6 and are very similar to the baseline BMA results.

Second, we perform a robustness check using an alternative set of BMA priors, employing the BRIC g-prior suggested by Fernandez *et al.* (2001) together with the beta-binomial model prior, which gives each model size (in contrast to each model) equal prior probability (Ley & Steel, 2009). We label this estimation according to the g-prior parameter as “BRIC.” The results of this robustness check are reported in Table B2 in the appendix and are again similar to those of the baseline specification. In the right-hand panel of the same table we report the results of a fully frequentist technique, FMA. It employs Mallows’ weights, which have been shown by Hansen (2007) to be optimal for frequentist model averaging, and the orthogonalization of model space suggested by Amini & Parmeter (2012). FMA has recently been applied in meta-analysis, for example, by Bajzik *et al.* (2020); Havranek *et al.* (2017, 2018a,b,c). Also this robustness check corroborates the results we have discussed previously.

Third, in Table B3 we present three BMA specifications that use a subset of discount rate estimates, a different set of variables, or both. The first specification from the left excludes the standard error. While the exclusion might introduce an omitted-variable bias (the standard error, our proxy for the extent of publication bias, is a key variable in all our previous models), it reduces the danger of endogenous controls. Of the eleven variables with posterior inclusion probability above 50% in the benchmark model, two (health domain and other domain) slip below the 50% threshold, though in the case of health only slightly (to 44%). Nevertheless, there are 5 new variables that achieve a posterior inclusion probability above 50%, including *Real reward*. Our results thus suggest that if we ignored publication bias in the heterogeneity analysis, we would (erroneously, in our opinion given the remaining evidence) conclude that the use of hypothetical rewards biases the results of experiments. The second specifi-

cation from the left includes a variable reflecting the size of stakes in the experiment, information that is available only for a subset of the discount rate estimates. The estimated effect of the variable is negative, which is consistent with the magnitude effect (Meyer, 2015). The third specification excludes discount rate estimates for which the discounting model is not explicitly specified in the paper. Here we lose high posterior inclusion probability for the variable reflecting student samples, but we note that the variable proves to be important in all other specifications.

Finally, in Table B4 we consider two specifications that feature i) an interaction term between *Money domain* and *Non-linearity correction* and ii) a sub-sample of estimates for which the measurement error in the variable *Delay* is reduced. The interaction term is meant to capture the difference between discount rates estimated with and without correcting for non-linearity in utility functions (non-linearity is discussed in Section 2). Nevertheless, the interaction attains a very low posterior inclusion probability. Hence we fail to obtain evidence which would suggest that this variable is important for systematically explaining the heterogeneity in the reported discount rates. Regarding the right-hand part of Table B4, we use a sub-sample of estimates for which delay is precisely defined. For 61% estimates of the discount rate in our sample, the corresponding delay is clearly reported in the papers. The remaining estimates are derived from a series of questions with varying horizons, where for “delay” we use the maximum horizon to which a subject is exposed in a given experimental task. Similarly to the baseline BMA result, we fail to obtain the anticipated significant negative coefficient. The insignificance result would likewise hold if we used the mean or median instead of the maximum to approximate the delay variable for discount rate estimates obtained from questions with varying horizons.

6 Concluding Remarks

We provide a quantitative synthesis of the literature that uses experiments to identify individual discount rates. We examine 927 estimates of the discount rate reported in 56 primary studies. By employing meta-regression and other methods, we detect selective reporting against null and negative results. The mean reported discount rate is 0.80. Using conservative techniques, we find that the mean drops to about 0.33 after we correct for publication bias—that is, people are more patient on average than what is indicated by a naive summary of the conclusions of the experiments. This result is in line with Imai *et al.* (2020a), who report evidence of modest selective reporting in the literature estimating the present bias parameter. In contrast, Imai *et al.* (2020b) find little evidence of publication bias in laboratory economics experiments.

The estimates of the discount rate vary a great deal. We explain this heterogeneity by using Bayesian model averaging, a method accounting for model uncertainty inherent in meta-analysis. We corroborate the presence of selective reporting in the literature by showing that the standard error is an important factor in the heterogeneity of discount rate estimates. We corroborate the domain independence hypothesis stressed by the previous literature (Cairns, 1992; Chapman & Elstein, 1995; Loewenstein *et al.*, 2003) since discount rates for different questions (for example, health on one hand and money on the other) differ systematically. Other important results include the systematic difference between lab and field experiments and the importance of framing and the composition of the subject pool.

The results of our study can be used in various settings. The discount rate has implications for decisions regarding savings, education, smoking, exercise, and other contexts of day-to-day behavior (e.g., Chabris *et al.*, 2008; Meier & Sprenger, 2010). Accurate measures of discounting parameters can provide helpful guidance in welfare analyses on the potential impacts of policies and provide useful diagnostics for effective policy targeting (Andreoni *et al.*, 2015); moreover, they can be applicable to modeling political campaigns, advertisement, and R&D investment (Deck & Jahedi, 2015b). Other examples of applications are discussed by Deck & Jahedi (2015a), who examine discounting in strategic settings, such as auctions or experimental contests, in which it is often critical to accurately predict the behavior of counterparts.

Climate change policies, in which the individual pure rate of time preference or the social discount rate is needed to evaluate the long-term effects, can serve as an example of a welfare analysis application of our results. The pure rate of time preference together with the growth rate of per capita consumption and the elasticity of marginal utility of consumption create the basis for the calculation of the Ramsey discount rate consisting of *time* and *growth* discounting elements (Fearnside, 2002; Anthoff *et al.*, 2009; Foley *et al.*, 2013). Our discount rate synthesis together with the results of Havranek *et al.* (2015a), who provide a meta-analysis of the elasticity of marginal utility of consumption, can be employed to calculate the pure rate of time preference from the Ramsey discount rate.

Our results also have broad implications for future experimental research on discounting. The potential for publication bias is correlated with the occurrence of large positive outliers, which means that estimates of the median discount rate are more robust to the bias than estimates of the average discount rate. Indeed, we find some direct evidence in our data set that median estimates may suffer less from publication bias compared to mean estimates. Papers that estimate individual-specific discounting often report median statistics for this reason (see, for example, Kuhn *et al.*, 2017). Lab

experiments seem to yield, *ceteris paribus*, larger estimates of the discount rate compared to field experiments. Because both lab and field experiments have their pros and cons (Al-Ubaydli & List, 2015), we need more studies along the lines of Andersen *et al.* (2010) that would evaluate the results of both in a comparable environment. We obtain robust evidence that the estimated discount rates are not systematically affected by the fact whether rewards in the experiment are real or hypothetical. In contrast, discount rates vary a lot across domains: subjects display substantially less patience for goods where intertemporal markets are limited compared to money—health, vacations, kisses from movie stars. In conjunction with the finding that discount rates tend to be larger for groups that are likely to be liquidity-constrained (e.g., students), these results suggest that the experimental subjects’ decisions are not fully divorced from outside conditions. If this is the case, current experimental measures may not allow us to properly identify preference parameters, though they are useful for understanding the intertemporal behavior of subjects under various external constraints (Dean & Sautmann, 2020). The literature thus awaits novel techniques that will ensure narrow bracketing and enable an even cleaner identification of the underlying discount rates.

References

- ABDELLAOUI, M., A. ATTEMA, & H. BLEICHRODT (2010): “Intertemporal Trade-offs for Gains and Losses: An Experimental Measurement of Discounted Utility.” *The Economic Journal* **120(545)**: pp. 845–866.
- VAN AERT, R. C. & M. VAN ASSEN (2020): “Correcting for publication bias in a meta-analysis with the p-uniform* method.” *Working paper*, Tilburg University & Utrecht University.
- AINSLIE, G. (1975): “Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control.” *Psychological Bulletin* **82(4)**: pp. 463–496.
- AL-UBAYDLI, O. & J. A. LIST (2015): “Do Natural Field Experiments Afford Researchers More or Less Control Than Laboratory Experiments?” *American Economic Review* **105(5)**: pp. 462–466.
- ALTMEJD, A., A. DREBER, E. FORSELL, J. HUBER, T. IMAI, M. JOHANNESSON, M. KIRCHLER, G. NAVE, & C. CAMERER (2019): “Predicting the replicability of social science lab experiments.” *PLOS ONE* **14(12)**: pp. 1–18.
- AMINI, S. M. & C. F. PARMETER (2012): “Comparison of model averaging techniques: Assessing growth determinants.” *Journal of Applied Econometrics* **27(5)**: pp. 870–876.
- ANDERSEN, S., G. HARRISON, M. LAU, & E. RUTSTROM (2006): “Elicitation Using Multiple Price List Formats.” *Experimental Economics* **9(4)**: pp. 383–405.
- ANDERSEN, S., G. HARRISON, M. LAU, & E. RUTSTROM (2008): “Eliciting Risk and Time Preferences.” *Econometrica* **76(3)**: pp. 583–618.
- ANDERSEN, S., G. HARRISON, M. LAU, & E. RUTSTROM (2010): “Preference Heterogeneity in Experiments: Comparing the Field and Laboratory.” *Journal of Economic Behavior & Organization* **73(2)**: pp. 209–224.
- ANDERSEN, S., G. HARRISON, M. LAU, & E. RUTSTROM (2013): “Discounting Behaviour and the Magnitude Effect: Evidence from a Field Experiment in Denmark.” *Economica* **80(320)**: pp. 670–697.
- ANDERSEN, S., G. HARRISON, M. LAU, & E. RUTSTROM (2014): “Discounting Behavior: A Reconsideration.” *European Economic Review* **71(C)**: pp. 15–33.
- ANDREONI, J., M. A. KUHN, & C. SPRENGER (2015): “Measuring Time Preferences: A Comparison of Experimental Methods.” *Journal of*

- Economic Behavior & Organization* **116(C)**: pp. 451–464.
- ANDREONI, J. & C. SPRENGER (2012): “Estimating Time Preferences from Convex Budgets.” *American Economic Review* **102(7)**: pp. 3333–3356.
- ANDREONI, J. & C. SPRENGER (2012b): “Risk Preferences Are Not Time Preferences.” *American Economic Review* **102(7)**: pp. 3357–3376.
- ANDREWS, I. & M. KASY (2019): “Identification of and Correction for Publication Bias.” *American Economic Review* **109(8)**: pp. 2766–94.
- ANTHOFF, D., R. S. J. TOL, & G. W. YOHE (2009): “Risk Aversion, Time Preference, and the Social Cost of Carbon.” *Environmental Research Letters* **4(2)**: pp. 240–242.
- ASTAKHOV, A., T. HAVRANEK, & J. NOVAK (2019): “Firm Size And Stock Returns: A Quantitative Survey.” *Journal of Economic Surveys* **33(5)**: pp. 1463–1492.
- ATTEMA, A. E., H. BLEICHRODT, Y. GAO, Z. HUANG, & P. P. WAKKER (2016): “Measuring Discounting without Measuring Utility.” *American Economic Review* **106(6)**: pp. 1476–1494.
- BABECKY, J. & T. HAVRANEK (2014): “Structural reforms and growth in transition.” *The Economics of Transition* **22(1)**: pp. 13–42.
- BAJZIK, J., T. HAVRANEK, Z. IRSOVA, & J. SCHWARZ (2020): “Estimating the Armington elasticity: The importance of study design and publication bias.” *Journal of International Economics* **127(C)**: p. 103383.
- BAUER, M. & J. CHYTILOVA (2010): “The Impact of Education on Subjective Discount Rate in Ugandan Villages.” *Economic Development and Cultural Change* **58(4)**: pp. 643–669.
- BAUER, M. & J. CHYTILOVA (2013): “Women, Children and Patience: Experimental Evidence from Indian Villages.” *Review of Development Economics* **17(4)**: pp. 662–675.
- BAUER, M., J. CHYTILOVA, & J. MORDUCH (2012): “Behavioral Foundations of Microcredit: Experimental and Survey Evidence from Rural India.” *American Economic Review* **102(2)**: pp. 1118–1139.
- BENZION, U., A. RAPOPORT, & J. YAGIL (1989): “Discount Rates Inferred from Decisions: An Experimental Study.” *Management Science* **35(3)**: pp. 270–284.
- BLANCO-PEREZ, C. & A. BRODEUR (2020): “Publication Bias and Editorial Statement on Negative Findings.” *Economic Journal* **130(629)**: pp. 1226–1247.
- BLEICHRODT, H., Y. GAO, & K. I. M. ROHDE (2016): “A measurement of decreasing impatience for health and money.” *Journal of Risk and Uncertainty* **52(3)**: pp. 213–231.
- BOM, P. R. D. & H. RACHINGER (2019): “A Kinked Meta-Regression Model for Publication Bias Correction.” *Research Synthesis Methods* **10(4)**: pp. 497–514.
- BOOIJ, A. S. & B. M. VAN PRAAG (2009): “A Simultaneous Approach to the Estimation of Risk Aversion and the Subjective Time Discount Rate.” *Journal of Economic Behavior & Organization* **70(1-2)**: pp. 374–388.
- BRODEUR, A., N. COOK, & A. HEYES (2020a): “A Proposed Specification Check for p-Hacking.” *AEA Papers and Proceedings* **110**: pp. 66–69.
- BRODEUR, A., N. COOK, & A. HEYES (2020b): “Methods Matter: P-Hacking and Causal Inference in Economics.” *American Economic Review* **110(11)**: pp. 3634–60.
- BRODEUR, A., M. LE, M. SANGNIER, & Y. ZYLBERBERG (2016): “Star Wars: The Empirics Strike Back.” *American Economic Journal: Applied Economics* **8(1)**: pp. 1–32.
- BROWN, A., T. IMAI, F. VIEIDER, & C. F. CAMERER (2020): “Meta-Analysis of Empirical Estimates of Loss-Aversion.” *mimeo*, LMU Munich.
- BROWN, A. L., Z. E. CHUA, & C. F. CAMERER (2009): “Learning and Visceral Temptation in Dynamic Saving Experiments.” *The Quarterly Journal of Economics* **124(1)**: pp. 197–231.
- BURKS, S., J. CARPENTER, L. GOTTE, & A. RUSTICHINI (2012): “Which Measures of Time Preference Best Predict Outcomes: Evidence from a Large-Scale Field Experiment.” *Journal of Economic Behavior & Organization* **84(1)**: pp. 308–320.
- CAIRNS, J. A. (1992): “Health, Wealth and Time Preference.” *Project Appraisal* **7(1)**: pp. 31–40.
- CAIRNS, J. A. & M. VAN DER POL (1997): “Constant and Decreasing Timing Aversion for Saving Lives.” *Social Science & Medicine* **45(11)**: pp. 1653–1659.
- CAMERER, C. F., A. DREBER, T. H. HO, J. HUBER, M. JOHANNESSEN, M. KIRCHLER, J. ALMENBERG, A. ALTMEJD, T. CHAN, E. FORSELL, E. HEIKENSTEN, F. HOLZMEISTER, T. IMAI, S. ISAKSSON, G. NAVE, T. PFEIFFER, M. RAZEN, & H. WU (2016): “Evaluating Replicability of Laboratory Experiments in Economics.” *Science* **351(6280)**: pp. 1433–1436.
- CAMERER, C. F., A. DREBER, F. HOLZMEIS-

- TER, T. H. HO, J. HUBER, M. JOHANNESSEN, M. KIRCHER, G. N. G, B. A. NOSEK, T. PFEIFFER, A. ALTMEJD, N. BUTTRICK, T. CHAN, Y. CHEN, E. FORSELL, A. GAMPA, E. HEIKENSTEN, L. HUMMER, T. IMAL, S. ISAKSSON, D. MANFREDI, J. ROSE, E. J. WAGENMAKERS, & H. WU (2018): "Evaluating the Replicability of Social Science Experiments in Nature and Science between 2010 and 2015." *Nature Human Behaviour* 2: pp. 637–644.
- CAMPOS, N. F., J. FIDRMUC, & I. KORHONEN (2019): "Business cycle synchronisation and currency unions: A review of the econometric evidence using meta-analysis." *International Review of Financial Analysis* 61(C): pp. 274–283.
- CARLSSON, F., H. HE, P. MARTINSSON, P. QIN, & M. SUTTER (2012): "Household Decision Making in Rural China: Using Experiments to Estimate the Influences of Spouses." *Journal of Economic Behavior & Organization* 84(2): pp. 525–536.
- CARTER, J. R. & M. D. IRONS (1991): "Are Economists Different, and If So, Why?" *Journal of Economic Perspectives* 5(2): pp. 171–177.
- CASSAR, A., A. HEALY, & C. VON KESSLER (2017): "Trust, Risk, and Time Preferences after a Natural Disaster: Experimental Evidence from Thailand." *World Development* 94(C): pp. 90–105.
- CASTILLO, M., P. J. FERRARO, J. L. JORDAN, & R. PETRIE (2011): "The Today and Tomorrow of Kids: Time Preferences and Educational Outcomes of Children." *Journal of Public Economics* 95(11): pp. 1377–1385.
- CAZACHEVICI, A., T. HAVRANEK, & R. HORVATH (2020): "Remittances and Economic Growth: A Meta-Analysis." *World Development* 134: p. 105021.
- CHABRIS, C. F., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, & D. TAUBINSKY (2008): "Individual Laboratory-Measured Discount Rates Predict Field Behavior." *Journal of Risk and Uncertainty* 37(2–3): pp. 237–269.
- CHABRIS, C. F., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, & D. TAUBINSKY (2008b): "Measuring Intertemporal Preferences Using Response Times." *NBER Working Paper 2008/14353*, National Bureau of Economic Research, Cambridge: MA.
- CHABRIS, C. F., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, & D. TAUBINSKY (2009): "The Allocation of Time in Decision-Making." *Journal of the European Economic Association* 7(2–3): pp. 628–637.
- CHAPMAN, G. B. (1996): "Temporal Discounting and Utility for Health and Money." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22(3): pp. 771–791.
- CHAPMAN, G. B. & A. S. ELSTEIN (1995): "Valuing the Future: Temporal Discounting of Health and Money." *Medical Decision Making* 15(4): pp. 373–386.
- CHAPMAN, G. B., R. NELSON, & D. B. HIER (1999): "Familiarity and Time Preferences: Decision Making about Treatments for Migraine Headaches and Crohn's Disease." *Journal of Experimental Psychology: Applied* 5(1): pp. 17–34.
- CHAPMAN, G. B. & J. R. WINQUIST (1998): "The Magnitude Effect: Temporal Discount Rates and Restaurant Tips." *Psychonomic Bulletin & Review* 5(1): pp. 119–123.
- CHESSON, H. & W. K. VISCUSI (2000): "The Heterogeneity of Time-Risk Tradeoffs." *Journal of Behavioral Decision Making* 13(2): pp. 251–258.
- CHEUNG, S. L. (2016): "Recent Developments in the Experimental Elicitation of Time Preference." *Journal of Behavioral and Experimental Finance* 11(C): pp. 1–8.
- COHEN, J., K. M. ERICSON, D. LAIBSON, & J. M. WHITE (2020): "Measuring Time Preferences." *Journal of Economic Literature* 58(2): pp. 299–347.
- COLLER, M. & M. B. WILLIAMS (1999): "Eliciting Individual Discount Rates." *Experimental Economics* 2(2): pp. 107–127.
- DAVIS, D. D. & C. A. HOLT (1993): *Experimental Economics*. Princeton: Princeton University Press, Princeton: NJ.
- DEAN, M. & A. SAUTMANN (2020): "Credit Constraints and the Measurement of Time Preferences." *Review of Economics and Statistics* (forthcoming).
- DECK, C. & S. JAHEDI (2015a): "An Experimental Investigation of Time Discounting in Strategic Settings." *Journal of Behavioral and Experimental Economics* 54(C): pp. 95–104.
- DECK, C. & S. JAHEDI (2015b): "Time Discounting in Strategic Contests." *Journal of Economics & Management Strategy* 24(1): pp. 151–164.
- DEPOSITARIO, D. P. T., R. M. NAYGA, X. WU, & T. P. LAUDE (2009): "Should Students Be Used as Subjects in Experimental Auctions?" *Economics Letters* 102(2): pp. 122–124.
- DOLAN, P. & C. GUDEX (1995): "Time Preference, Duration and Health State Valuations." *Health Economics* 4(4): pp. 289–299.
- DOUCOULIAGOS, C. & T. D. STANLEY (2013): "Are

- All Economic Facts Greatly Exaggerated? Theory Competition and Selectivity." *Journal of Economic Surveys* **27(2)**: pp. 316–339.
- DOUCOULIAGOS, H. & M. PALDAM (2011): "The ineffectiveness of development aid on growth: An update." *European Journal of Political Economy* **27(2)**: pp. 399–404.
- DREYFUS, M. K. & W. K. VISCUSI (1995): "Rates of Time Preference and Consumer Valuations of Automobile Safety and Fuel Efficiency." *The Journal of Law and Economics* **38(1)**: pp. 79–105.
- DUAN, J., K. K. DAS, L. MERILUOTO, & W. R. REED (2020): "Estimating the effect of spillovers on exports: a meta-analysis." *Review of World Economics* **156(2)**: pp. 219–249.
- DUQUETTE, E., N. HIGGINS, & J. HOROWITZ (2012): "Farmer Discount Rates: Experimental Evidence." *American Journal of Agricultural Economics* **94(2)**: pp. 451–456.
- EGGER, M., G. DAVEY SMITH, M. SCHNEIDER, & C. MINDER (1997): "Bias in Meta-Analysis Detected by a Simple, Graphical Test." *British Medical Journal* **315(7109)**: pp. 629–34.
- EICHER, T. S., C. PAPAGEORGIOU, & A. E. RAFTERY (2011): "Default Priors and Predictive Performance in Bayesian Model Averaging, with Application to Growth Determinants." *Journal of Applied Econometrics* **26(1)**: pp. 30–55.
- FEARNSIDE, P. M. (2002): "Time Preference in Global Warming Calculations: A Proposal for a Unified Index." *Ecological Economics* **41(1)**: pp. 21–31.
- FERNANDEZ, C., E. LEY, & M. F. STEEL (2001): "Benchmark Priors for Bayesian Model Averaging." *Journal of Econometrics* **100(2)**: pp. 381–427.
- FIELD, E., R. PANDE, J. PAPP, & N. RIGOL (2013): "Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India." *American Economic Review* **103(6)**: pp. 2196–2226.
- FINKE, M. S. & S. J. HUSTON (2013): "Time preference and the importance of saving for retirement." *Journal of Economic Behavior & Organization* **89(C)**: pp. 23–34.
- FISHER, I. (1930): *The Theory of Interest*. New York: Macmillan.
- FOLEY, D. K., A. REZAI, & L. TAYLOR (2013): "The Social Cost of Carbon Emissions: Seven Propositions." *Economics Letters* **121(1)**: pp. 90–97.
- FRANK, R. H., T. GILOVICH, & D. T. REGAN (1993): "Does Studying Economics Inhibit Cooperation?" *Journal of Economic Perspectives* **7(2)**: pp. 159–171.
- FREDERICK, S., G. LOEWENSTEIN, & T. O'DONOGHUE (2002): "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* **40(2)**: pp. 351–401.
- FUJII, T. & L. KARP (2008): "Numerical Analysis of Non-Constant Pure Rate of Time Preference: A model of Climate Policy." *Journal of Environmental Economics and Management* **56(1)**: pp. 83–101.
- FURUKAWA, C. (2020): "Publication Bias under Aggregation Frictions: From Communication Model to New Correction Method." *Working paper*, MIT.
- GEORGE, E. I. (2010): "Dilution priors: Compensating for model space redundancy." In "IMS Collections Borrowing Strength: Theory Powering Applications - A Festschrift for Lawrence D. Brown," volume 6, pp. 158–165. Institute of Mathematical Statistics.
- GERBER, A. & N. MALHOTRA (2008): "Do Statistical Reporting Standards Affect What Is Published? Publication Bias in Two Leading Political Science Journals." *Quarterly Journal of Political Science* **3(3)**: pp. 313–326.
- GEYER-KLINGEBERG, J., M. HANG, & A. W. RATHGEBER (2019): "What drives financial hedging? A meta-regression analysis of corporate hedging determinants." *International Review of Financial Analysis* **61(C)**: pp. 203–221.
- GOULDER, L. H. & R. N. STAVINS (2002): "Discounting: An Eye on the Future." *Nature* **419(6908)**: pp. 673–674.
- HANSEN, B. (2007): "Least Squares Model Averaging." *Econometrica* **75(4)**: pp. 1175–1189.
- HARDISTY, D. J., K. F. THOMPSON, D. H. KRANTZ, & E. U. WEBER (2013): "How to Measure Time Preferences: An Experimental Comparison of Three Methods." *Judgment and Decision Making* **8(3)**: pp. 236–249.
- HARRIS, C. (2012): "Feelings of Dread and Intertemporal Choice." *Journal of Behavioral Decision Making* **25(1)**: pp. 13–28.
- HARRISON, G. W., M. I. LAU, E. E. RUTSTROM, & M. B. SULLIVAN (2005): "Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues." In "Field Experiments in Economics (Research in Experimental Economics, Vol. 10)," pp. 125–218. Emerald Group Publishing Limited, Bingley: UK.
- HARRISON, G. W., M. I. LAU, & E. E. RUTSTROM (2010): "Individual Discount Rates and Smoking: Evidence from a Field Experiment in Den-

- mark." *Journal of Health Economics* **29(5)**: pp. 708–717.
- HARRISON, G. W., M. I. LAU, & M. B. WILLIAMS (2002): "Estimating Individual Discount Rates in Denmark: A Field Experiment." *American Economic Review* **92(5)**: pp. 1606–1617.
- HAUSMAN, J. (1979): "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell Journal of Economics* **10(1)**: pp. 33–54.
- HAVRANEK, T. (2010): "Rose effect and the euro: is the magic gone?" *Review of World Economics* **146(2)**: pp. 241–261.
- HAVRANEK, T. (2015): "Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting." *Journal of the European Economic Association* **13(6)**: pp. 1180–1204.
- HAVRANEK, T., D. HERMAN, , & Z. IRSOVA (2018a): "Does Daylight Saving Save Electricity? A Meta-Analysis." *The Energy Journal* **39(2)**: pp. 35–61.
- HAVRANEK, T., R. HORVATH, Z. IRSOVA, & M. RUSNAK (2015a): "Cross-Country Heterogeneity in Intertemporal Substitution." *Journal of International Economics* **96(1)**: pp. 100–118.
- HAVRANEK, T. & Z. IRSOVA (2010): "Meta-Analysis of Intra-Industry FDI Spillovers: Updated Evidence." *Czech Journal of Economics and Finance* **60(2)**: pp. 151–174.
- HAVRANEK, T. & Z. IRSOVA (2017): "Do Borders Really Slash Trade? A Meta-Analysis." *IMF Economic Review* **65(2)**: pp. 365–396.
- HAVRANEK, T., Z. IRSOVA, K. JANDA, & D. ZILBERMAN (2015b): "Selective reporting and the social cost of carbon." *Energy Economics* **51(C)**: pp. 394–406.
- HAVRANEK, T., Z. IRSOVA, & T. VLACH (2018b): "Measuring the Income Elasticity of Water Demand: The Importance of Publication and Endogeneity Biases." *Land Economics* **94(2)**: pp. 259–283.
- HAVRANEK, T., Z. IRSOVA, & O. ZEYNALOVA (2018c): "Tuition Fees and University Enrollment: A Meta-Regression Analysis." *Oxford Bulletin of Economics and Statistics* **80(6)**: pp. 1145–1184.
- HAVRANEK, T. & O. KOKES (2015): "Income elasticity of gasoline demand: A meta-analysis." *Energy Economics* **47(C)**: pp. 77–86.
- HAVRANEK, T., M. RUSNAK, & A. SOKOLOVA (2017): "Habit Formation in Consumption: A Meta-Analysis." *European Economic Review* **95(C)**: pp. 142–167.
- HAVRANEK, T., T. D. STANLEY, H. DOUCOLIAGOS, P. BOM, J. GEYER-KLINGEBERG, I. IWASAKI, W. R. REED, K. ROST, & R. C. M. VAN AERT (2020): "Reporting Guidelines for Meta-Analysis in Economics." *Journal of Economic Surveys* **34(3)**: pp. 469–475.
- HOLT, C. A. & S. K. LAURY (2002): "Risk Aversion and Incentive Effects." *American Economic Review* **92(5)**: pp. 1644–1655.
- IFCHER, J. & H. ZARGHAMEE (2011): "Happiness and Time Preference: The Effect of Positive Affect in a Random-Assignment Experiment." *American Economic Review* **101(7)**: pp. 3109–3129.
- IMAI, T., T. A. RUTTER, & C. F. CAMERER (2020a): "Meta-Analysis of Present-Bias Estimation using Convex Time Budgets." *Economic Journal* (forthcoming).
- IMAI, T., K. ZEMLIANOVA, N. KOTECHA, & C. F. CAMERER (2020b): "How Common are False Positives in Laboratory Economics Experiments? Evidence from the P-Curve Method." *mimeo*, LMU Munich.
- IOANNIDIS, J. P., T. D. STANLEY, & H. DOUCOLIAGOS (2017): "The Power of Bias in Economics Research." *Economic Journal* **127(605)**: pp. F236–F265.
- IRSOVA, Z. & T. HAVRANEK (2013): "Determinants of Horizontal Spillovers from FDI: Evidence from a Large Meta-Analysis." *World Development* **42(C)**: pp. 1–15.
- JOHNSON, M. W. & W. K. BICKEL (2002): "Within-Subject Comparison of Real and Hypothetical Money Rewards in Delay Discounting." *Journal of the Experimental Analysis of Behavior* **77(2)**: pp. 129–146.
- KASS, R. E. & A. E. RAFTERY (1995): "Bayes Factors." *Journal of the American Statistical Association* **90(430)**: pp. 773–795.
- KIRBY, K., N. PETRY, & W. BICKEL (1999): "Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls." *Journal of Experimental Psychology: General* **128(1)**: pp. 78–87.
- KIRBY, K. N. & N. N. MARAKOVIC (1995): "Modeling Myopic Decisions: Evidence for Hyperbolic Delay-Discounting within Subjects and Amounts." *Organizational Behavior and Human Decision Processes* **64(1)**: pp. 22–30.
- KIRBY, K. N. & N. N. MARAKOVIC (1996): "Delay-

- Discounting Probabilistic Rewards: Rates Decrease as Amounts Increase." *Psychonomic Bulletin and Review* **3(1)**: pp. 100–104.
- KOVACS, K. F. & D. M. LARSON (2008): "Identifying Individual Discount Rates and Valuing Public Open Space with Stated-Preference Models." *Land Economics* **84(2)**: pp. 209–224.
- KUHN, M. A., P. KUHN, & M. C. VILLEVAL (2017): "Decision-environment effects on intertemporal financial choices: How relevant are resource-depletion models?" *Journal of Economic Behavior & Organization* **137(C)**: pp. 72–89.
- KUHNBERGER, A., M. SCHULTE-MECKLENBECK, & J. PERNER (2002): "Framing Decisions: Hypothetical and Real." *Organizational Behavior and Human Decision Processes* **89(2)**: pp. 1162–1175.
- LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting." *The Quarterly Journal of Economics* **112(2)**: pp. 443–478.
- LAURY, S. K., M. M. MCINNES, & J. T. SWARTHOUT (2012): "Avoiding the Curves: Direct Elicitation of Time Preferences." *Journal of Risk and Uncertainty* **44(3)**: pp. 181–217.
- LAWRANCE, E. C. (1991): "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* **99(1)**: pp. 54–77.
- LEY, E. & M. F. STEEL (2009): "On the Effect of Prior Assumptions in Bayesian Model Averaging with Applications to Growth Regression." *Journal of Applied Econometrics* **24(4)**: pp. 651–674.
- LOCEY, M. L., B. A. JONES, & H. RACHLIN (2011): "Real and Hypothetical Rewards." *Judgment and Decision Making* **6(6)**: pp. 552–564.
- LOEWENSTEIN, G. (1987): "Anticipation and the Valuation of Delayed Consumption." *The Economic Journal* **97(387)**: pp. 666–684.
- LOEWENSTEIN, G., D. READ, & R. F. BAUMEISTER (2003): *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. Russell Sage Foundation, New York: NY.
- MADIGAN, D., J. YORK, & D. ALLARD (1995): "Bayesian Graphical Models for Discrete Data." *International Statistical Review* **63(2)**: pp. 215–232.
- MALESZA, M. (2019): "The Effects of Potentially Real and Hypothetical Rewards on Effort Discounting in a Student Sample." *Personality and Individual Differences* **151(C)**: p. 108807.
- MARWELL, G. & R. E. AMES (1981): "Economists Free Ride, Does Anyone Else? Experiments on the Provision of Public Goods." *Journal of Public Economics* **15(3)**: pp. 295–310.
- MAZUR, J. E. (1984): "Tests of an Equivalence Rule for Fixed and Variable Reinforcer Delays." *Journal of Experimental Psychology: Animal Behavior Processes* **10(4)**: pp. 426–436.
- MCCLOSKEY, D. N. & S. T. ZILIAK (2019): "What Quantitative Methods Should We Teach to Graduate Students? A Comment on Swann's Is Precise Econometrics an Illusion?" *Journal of Economic Education* **50(4)**: pp. 356–361.
- MCCLURE, S. M., K. M. ERICSON, D. I. LAIBSON, G. LOEWENSTEIN, & J. D. COHEN (2007): "Time Discounting for Primary Rewards." *Journal of Neuroscience* **27(21)**: pp. 5796–5804.
- MEIER, S. & C. D. SPRENGER (2010): "Present-Biased Preferences and Credit Card Borrowing." *American Economic Journal: Applied Economics* **2(1)**: pp. 193–210.
- MEIER, S. & C. D. SPRENGER (2013): "Discounting Financial Literacy: Time Preferences and Participation in Financial Education Programs." *Journal of Economic Behavior & Organization* **95(C)**: pp. 159–174.
- MEIER, S. & C. D. SPRENGER (2015): "Temporal Stability of Time Preferences." *The Review of Economics and Statistics* **97(2)**: pp. 273–286.
- MEYER, A. G. (2015): "The impacts of elicitation mechanism and reward size on estimated rates of time preference." *Journal of Behavioral and Experimental Economics* **58(C)**: pp. 132–148.
- NELSON, J. & J. MORAN (2020): "Effects of Alcohol Taxation on Prices: A Systematic Review and Meta-Analysis of Pass-Through Rates." *The B.E. Journal of Economic Analysis & Policy* **20(1)**: pp. 1–21.
- NEWELL, R. G. & J. SIKAMAKI (2015): "Individual Time Preferences and Energy Efficiency." *American Economic Review* **105(5)**: pp. 196–200.
- NOBEL, A., S. LIZIN, R. BROUWER, S. B. BRUNS, D. I. STERN, & R. MALINA (2020): "Are biodiversity losses valued differently when they are caused by human activities? A meta-analysis of the non-use valuation literature." *Environmental Research Letters* **15(7)**: pp. 1–23.
- OLIVOLA, C. Y. & S. W. WANG (2016): "Patience Auctions: The Impact of Time vs. Money Bidding on Elicited Discount Rates." *Experimental Economics* **19(4)**: pp. 864–885.
- PERCOCO, M. & P. NIJKAMP (2009): "Estimating Individual Rates of Discount: A Meta-Analysis." *Applied Economics Letters* **16(12)**: pp. 1235–1239.

- PHELPS, E. S. & R. A. POLLAK (1968): "On Second-Best National Saving and Game-Equilibrium Growth." *The Review of Economic Studies* **35(2)**: pp. 185–199.
- RAFTERY, A. E., D. MADIGAN, & J. A. HOETING (1997): "Bayesian Model Averaging for Linear Regression Models." *Journal of the American Statistical Association* **92(437)**: pp. 179–191.
- READ, D. & N. L. READ (2004): "Time Discounting Over the Lifespan." *Organizational Behavior and Human Decision Processes* **94(1)**: pp. 22–32.
- RUSNAK, M., T. HAVRANEK, & R. HORVATH (2013): "How to Solve the Price Puzzle? A Meta-Analysis." *Journal of Money, Credit and Banking* **45(1)**: pp. 37–70.
- SAMUELSON, P. (1937): "Note on Measurement of Utility." *Review of Economic Studies* **4(2)**: pp. 155–161.
- STANLEY, T. D. (2008): "Meta-Regression Methods for Detecting and Estimating Empirical Effects in the Presence of Publication Selection." *Oxford Bulletin of Economics and Statistics* **70(1)**: pp. 103–127.
- STANLEY, T. D. & H. DOUCOULIAGOS (2010): "Picture this: A simple Graph that Reveals Much Ado about Research." *Journal of Economic Surveys* **24(1)**: pp. 170–191.
- STANLEY, T. D. & H. DOUCOULIAGOS (2014): "Meta-Regression Approximations to Reduce Publication Selection Bias." *Research Synthesis Methods* **5(1)**: pp. 60–78.
- STANLEY, T. D., S. B. JARRELL, & H. DOUCOULIAGOS (2010): "Could It Be Better to Discard 90% of the Data? A Statistical Paradox." *The American Statistician* **64(1)**: pp. 70–77.
- SUTTER, M., M. G. KOCHER, D. GLATZLER, & S. T. TRAUTMANN (2013): "Impatience and Uncertainty: Experimental Decisions Predict Adolescents' Field Behavior." *American Economic Review* **103(1)**: pp. 510–531.
- TAKEUCHI, K. (2011): "Non-Parametric Test of Time Consistency: Present Bias and Future Bias." *Games and Economic Behavior* **71(2)**: pp. 456–478.
- TANAKA, T., C. F. CAMERER, & Q. NGUYEN (2010): "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam." *American Economic Review* **100(1)**: pp. 557–571.
- THALER, R. (1981): "Some empirical evidence on dynamic inconsistency." *Economics Letters* **8(3)**: pp. 201–207.
- TOKUNAGA, M. & I. IWASAKI (2017): "The Determinants of Foreign Direct Investment in Transition Economies: A Meta-analysis." *The World Economy* **40(12)**: pp. 2771–2831.
- TOL, R. S. J. (1999): "Time Discounting and Optimal Emission Reduction: An Application of FUND." *Climatic Change* **41(3-4)**: pp. 351–362.
- TSUKAYAMA, E. & A. L. DUCKWORTH (2010): "Domain-Specific Temporal Discounting and Temptation." *Judgment and Decision Making* **5(2)**: pp. 72–82.
- TVERSKY, A. & D. KAHNEMAN (1981): "The Framing of Decisions and the Psychology of Choice." *Science* **211(4481)**: pp. 453–458.
- UBFAL, D. (2016): "How general are time preferences? Eliciting good-specific discount rates." *Journal of Development Economics* **118(C)**: pp. 150–170.
- UGUR, M., S. A. CHURCHILL, & H. M. LUNONG (2020): "What do we know about R&D spillovers and productivity? Meta-analysis evidence on heterogeneity and statistical power." *Research Policy* **49(1)**.
- UGUR, M., S. A. CHURCHILL, & E. SOLOMON (2018): "Technological Innovation And Employment In Derived Labour Demand Models: A Hierarchical Meta-Regression Analysis." *Journal of Economic Surveys* **32(1)**: pp. 50–82.
- VALICKOVA, P., T. HAVRANEK, & R. HORVATH (2015): "Financial Development And Economic Growth: A Meta-Analysis." *Journal of Economic Surveys* **29(3)**: pp. 506–526.
- VOORS, M. J., E. E. M. NILLESEN, P. VERWIMP, E. H. BULTE, R. LENSINK, & D. P. V. SOEST (2012): "Violent Conflict and Behavior: A Field Experiment in Burundi." *American Economic Review* **102(2)**: pp. 941–964.
- WANG, M., M. O. RIEGER, & T. HENS (2016): "How Time Preferences Differ: Evidence from 53 Countries." *Journal of Economic Psychology* **52(C)**: pp. 115–135.
- WARNER, J. T. & S. PLEETER (2001): "The Personal Discount Rate: Evidence from Military Downsizing Programs." *American Economic Review* **91(1)**: pp. 33–53.
- WORLD BANK (2020): "Median Monthly Per Capita Expenditure (Or Income)." *World Bank's global database of household surveys (PovcalNet) March 2020 Update*, Global Poverty Monitoring, Washington, DC: World Bank.
- XUE, X., W. R. REED, & A. MENCLOVA (2020): "Social capital and health: a meta-analysis." *Journal*

of *Health Economics* **72(C)**: p. 102317.
 ZAUBERMAN, G., B. K. KIM, S. A. MALKOC, & J. R. BETTMAN (2009): "Discounting Time and Time Discounting: Subjective Time Perception and Intertemporal Preferences." *Journal of Marketing Research* **46(4)**: pp. 543–556.
 ZEUGNER, S. & M. FELDKIRCHER (2015):

"Bayesian Model Averaging Employing Fixed and Flexible Priors: The BMS Package for R." *Journal of Statistical Software* **68(4)**: pp. 1–37.
 ZIGRAIOVA, D. & T. HAVRANEK (2016): "Bank Competition And Financial Stability: Much Ado About Nothing?" *Journal of Economic Surveys* **30(5)**: pp. 944–981.

Appendices

A Robustness Checks to Tests of Publication Bias

Table A1: Funnel asymmetry tests with standard errors clustered at the level of authors

	OLS	Fixed effects	Instrument	Precision
Standard error (<i>publication bias</i>)	0.535*** (0.0331)	0.875*** (0.0146)	0.316 (0.194)	1.031** (0.455)
Constant (<i>effect beyond bias</i>)	0.518*** (0.125)	0.341*** (0.00762)	0.633*** (0.180)	0.259*** (0.0391)
Observations	927	927	927	927
Clusters	31	31	31	31

Notes: The table reports the results of regression $\delta_{ij} = \delta_1 + \gamma_1 \cdot SE(\delta_{ij}) + u_{ij}$, where δ_{ij} denotes the i -th annualized discount rate estimated in the j -th study, and $SE(\delta_{ij})$ denotes its standard error. The table shows estimation by OLS, study-level fixed effects, instrumental variables (where the instrument for the standard error is the inverse of the square root of the number of observations in a study), and precision weighting (where estimates are weighted by the inverse of their standard error). Standard errors, clustered at the level of authors, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Funnel asymmetry tests for medians of individual-specific discounting

	OLS	Fixed effects	Instrument	Precision
Standard error (<i>publication bias</i>)	0.535*** (0.0282)	0.875*** (0.0154)	0.535*** (0.0282)	1.012** (0.453)
Standard error * Median (<i>additional bias in median estimates</i>)	0.373 (0.259)	-1.093*** (0.0518)	0.373 (0.259)	0.417 (0.619)
Constant (<i>effect beyond bias</i>)	0.509*** (0.118)	0.369*** (0.00817)	0.509*** (0.118)	0.258*** (0.0376)
Observations	927	927	927	927

Notes: The table reports the results of regression $\delta_{ij} = \delta_1 + \gamma_1 \cdot SE(\delta_{ij}) + \gamma_2 \cdot SE(\delta_{ij}) \cdot Median_{ij} + u_{ij}$, where δ_{ij} denotes the i -th annualized discount rate estimated in the j -th study, $SE(\delta_{ij})$ denotes its standard error, and $Median$ is a dummy variable that equals 1 if the estimate of the discount rate is a median of individual-specific discounting. The table shows estimation by OLS, study-level fixed effects, instrumental variable (where the instrument for the standard error is the inverse of the square root of the number of observations in a study), and precision weighting (where estimates are weighted by the inverse of their standard error). Standard errors, clustered at the study level, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Excluding estimates with unidentified discounting type

PANEL A: Linear models				
	OLS	Fixed effects	Instrument	Precision
Standard error (<i>publication bias</i>)	1.112 ^{***} (0.210)	0.852 ^{**} (0.359)	-0.233 (1.598)	2.814 ^{***} (0.684)
Constant (<i>effect beyond bias</i>)	0.384 ^{***} (0.0745)	0.414 ^{***} (0.0403)	0.535 ^{***} (0.192)	0.194 ^{***} (0.0302)
Observations	507	507	507	507
PANEL B: Non-linear models				
	WAAP of Ioannidis <i>et al.</i> (2017)	Stem-based method of Furukawa (2020)	Selection model of Andrews & Kasy (2019)	Endogenous kink of Bom & Rachinger (2019)
Effect beyond bias	0.305 ^{***} (0.016)	0.067 [*] (0.040)	0.218 ^{***} (0.130)	0.145 ^{***} (0.004)
Observations	507	507	507	507

Notes: The table reports the results of regression $\delta_{ij} = \delta_1 + \gamma_1 \cdot SE(\delta_{ij}) + u_{ij}$, where δ_{ij} denotes the i -th annualized discount rate estimated in the j -th study, and $SE(\delta_{ij})$ denotes its standard error. Estimates for which the discounting model is not explicitly stated are omitted from estimations in this table. Panel A shows estimation by OLS, study-level fixed effects, instrumental variables (where the instrument for the standard error is the inverse of the square root of the number of observations in a study), and precision weighting (where estimates are weighted by the inverse of their standard error). Panel B shows the recently developed non-linear estimation techniques; WAAP stands for the Weighted Average of the Adequately Powered estimates. Standard errors, clustered at the study level, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Funnel asymmetry tests in absolute value

	OLS	Fixed effects	Instrument	Precision
Standard error (<i>bias in positive estimates</i>)	0.534 ^{***} (0.0304)	0.872 ^{***} (0.0158)	0.534 ^{***} (0.0304)	1.040 ^{**} (0.456)
Standard error * Negative (<i>bias in negative estimates</i>)	-2.104 ^{***} (0.371)	-0.610 (0.730)	-2.104 ^{***} (0.371)	-2.306 ^{***} (0.743)
Constant (<i>effect beyond bias</i>)	0.523 ^{***} (0.114)	0.344 ^{***} (0.00899)	0.523 ^{***} (0.114)	0.260 ^{***} (0.0374)
Observations	927	927	927	927

Notes: The table reports the results of regression $|\delta_{ij}| = \delta_1 + \gamma_1 \cdot SE(\delta_{ij}) + \gamma_2 \cdot SE(\delta_{ij}) \cdot Negative_{ij} + u_{ij}$, where δ_{ij} denotes the i -th annualized discount rate estimated in the j -th study, $SE(\delta_{ij})$ denotes its standard error, and $Negative$ is a dummy variable that equals 1 if the estimate of the discount rate is negative. The table shows estimation by OLS, study-level fixed effects, instrumental variables (where the instrument for the standard error is the inverse of the square root of the number of observations in a study), and precision weighting (where estimates are weighted by the inverse of their standard error). Standard errors, clustered at the study level, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

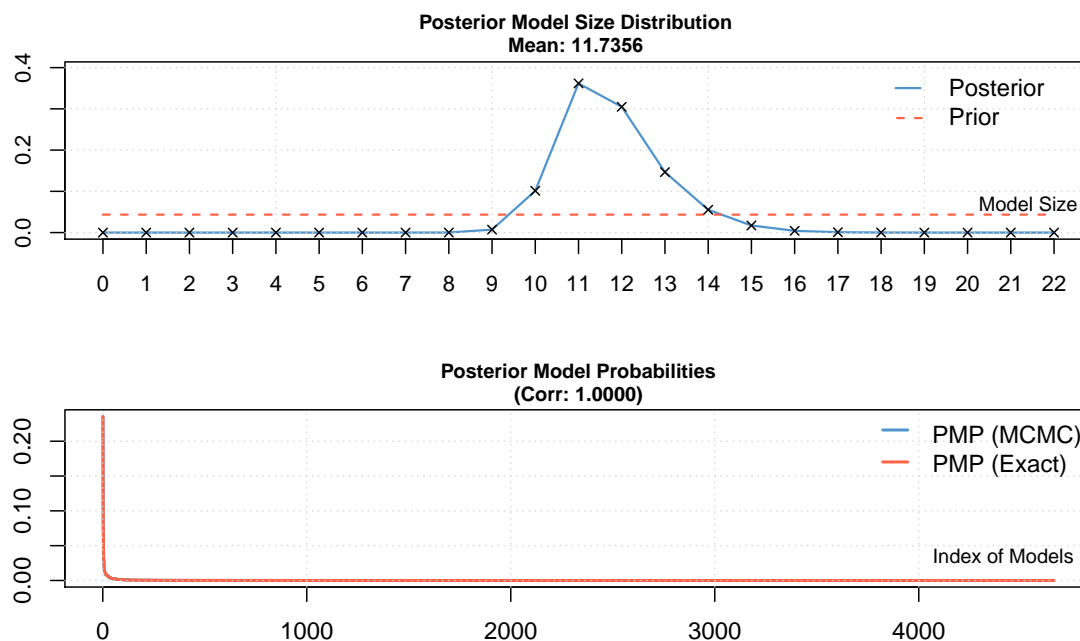
B Robustness Checks and Additional Statistics to BMA

Table B1: Summary of the benchmark BMA estimation

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
11.7356	$2 \cdot 10^6$	$1 \cdot 10^6$	2.350162 mins	402,090
<i>Modelspace</i>	<i>Models visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$4.19 \cdot 10^6$	9.60%	100%	1.0000	927
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random/11	UIP	Av = 0.9989		

Notes: We employ the priors recommended by Eicher *et al.* (2011) and George (2010), the unit information prior (the prior has the same weight as one observation in the data) and the dilution prior (accounting for potential collinearity). The results of this BMA exercise are reported in Table 6

Figure B1: Model size and convergence for the benchmark BMA model



Notes: The figure depicts the posterior model size distribution and the posterior model probabilities of the BMA exercise reported in Table 6.

Table B2: Alternative BMA priors and frequentist model averaging

Variable:	Bayesian model averaging (BRIC)			Frequentist model averaging		
	Post. mean	Post. SD	PIP	Mean	SE	p-value
Constant	-0.244	NA	1.000	-0.393	0.140	0.005
Standard error	0.549	0.021	1.000	0.572	0.024	0.000
<i>Estimation characteristics</i>						
Hyperbolic discounting	0.039	0.061	0.351	0.132	0.062	0.035
Exponential discounting	0.006	0.029	0.074	0.089	0.075	0.235
Delay	0.000	0.002	0.040	-0.002	0.009	0.843
Front-end delay	0.013	0.041	0.141	0.109	0.064	0.089
Lab experiment	0.156	0.101	0.777	0.124	0.075	0.100
<i>Experimental characteristics</i>						
Real reward	-0.005	0.026	0.075	-0.031	0.067	0.648
Matching task	0.017	0.045	0.160	0.017	0.066	0.791
Health domain	0.346	0.088	0.993	0.317	0.091	0.001
Other domain	0.441	0.069	1.000	0.424	0.072	0.000
Negative framing	-0.148	0.106	0.735	-0.139	0.077	0.073
Neutral framing	0.003	0.030	0.045	0.017	0.089	0.851
Stakes						
<i>Subject pool characteristics</i>						
Sample size	0.075	0.014	1.000	0.084	0.017	0.000
Students	0.877	0.111	1.000	0.825	0.132	0.000
Students * Lab experiment	-0.753	0.144	1.000	-0.670	0.160	0.000
Males only	0.013	0.052	0.089	0.098	0.110	0.374
Females only	-0.001	0.022	0.040	0.000	0.012	1.000
North America	0.012	0.040	0.125	0.113	0.066	0.085
Asia	0.385	0.103	0.991	0.384	0.095	0.000
Africa	3.170	0.118	1.000	3.295	0.137	0.000
<i>Publication characteristics</i>						
Citations	-0.003	0.011	0.094	-0.014	0.022	0.527
Publication year	0.121	0.026	1.000	0.104	0.029	0.000
Observations	927			927		
Studies	56			56		

Notes: Response variable = annualized estimates of the discount rate. SD = standard deviation, PIP = Posterior inclusion probability, SE = standard error. The first specification from the left uses Bayesian model averaging with an alternative model prior, the beta-binomial prior advocated by Ley & Steel (2009) and Zellner's g prior BRIC according to Fernandez *et al.* (2001). The second specification, frequentist model averaging, applies Mallows's model averaging estimator (Hansen, 2007) using the orthogonalization of covariate space suggested by Amini & Parmeter (2012) to reduce the number of estimated models. All variables are described in Table 5.

Table B3: Alternative specifications of the baseline BMA model

Variable:	Bayesian model averaging (without SE)			Bayesian model averaging (with stakes)			Bayesian model averaging (known model)		
	P. mean	P. SD	PIP	P. mean	P. SD	PIP	P. mean	P. SD	PIP
Constant	0.790	NA	1.000	0.180	NA	1.000	-0.704	NA	1.000
Standard error				0.567	0.023	1.000	0.856	0.110	1.000
<i>Estimation characteristics</i>									
Hyperbolic discounting	-0.383	0.068	1.000	0.001	0.013	0.039			
Exponential discounting	-0.505	0.084	1.000	0.004	0.023	0.055	0.000	0.012	0.043
Delay	-0.010	0.018	0.306	0.004	0.011	0.122	-0.098	0.017	1.000
Front-end delay	-0.403	0.062	1.000	0.050	0.069	0.398	0.185	0.111	0.808
Lab experiment	0.278	0.148	0.855	0.097	0.122	0.445	0.311	0.073	0.997
<i>Experimental characteristics</i>									
Real reward	0.166	0.140	0.664	-0.041	0.077	0.267	-0.001	0.020	0.051
Matching task	0.335	0.117	0.972	0.007	0.032	0.071	0.002	0.024	0.056
Health domain	0.110	0.144	0.442	0.979	0.173	1.000	0.382	0.095	0.996
Other domain	0.031	0.075	0.201	0.646	0.097	1.000	0.420	0.083	1.000
Negative framing	-0.409	0.092	0.999	-0.033	0.074	0.201	-0.030	0.075	0.179
Neutral framing	0.017	0.076	0.101	0.010	0.059	0.049	0.002	0.035	0.044
Stakes				-0.478	0.094	1.000			
<i>Subject pool characteristics</i>									
Sample size	-0.050	0.027	0.856	0.120	0.018	1.000	0.142	0.029	1.000
Students	0.933	0.193	1.000	0.398	0.296	0.755	-0.007	0.043	0.075
Students * Lab experiment	-0.684	0.254	0.960	-0.395	0.339	0.643	-0.001	0.044	0.066
Males only	0.005	0.042	0.071	0.016	0.061	0.092	0.015	0.063	0.085
Females only	-0.006	0.043	0.073	0.000	0.020	0.029	0.004	0.034	0.050
North America	0.005	0.030	0.093	-0.002	0.018	0.049	0.146	0.111	0.704
Asia	0.306	0.175	0.835	0.073	0.146	0.244	0.351	0.108	0.975
Africa	2.570	0.155	1.000	3.242	0.134	1.000			
<i>Publication characteristics</i>									
Citations	0.003	0.012	0.100	-0.041	0.045	0.511	-0.001	0.009	0.059
Publication year	0.374	0.038	1.000	0.017	0.036	0.232	0.013	0.034	0.173
Observations	927			777			507		
Studies	56			51			32		

Notes: Response variable = annualized estimates of the individual discount rate. P. mean = posterior mean, P. SD = posterior standard deviation, PIP = posterior inclusion probability. We employ Bayesian model averaging (BMA) using unit information prior (Eicher *et al.*, 2011) and the dilution prior suggested by George (2010) which accounts for collinearity. In the first specification from the left, we exclude the variable *Standard error*; in the second specification we introduce variable *Stakes* into the model (which reduces the number of observations to 777); in the third specification we use only those observations where the type of discounting can be explicitly identified. All variables are described in Table 5.

Table B4: BMA specifications accounting for non-linearity and exact delay

Variable:	Bayesian model averaging (money * non-linearity)			Bayesian model averaging (exact delay)		
	Post. mean	Post. SD	PIP	Post. mean	Post. SD	PIP
Constant	-0.242	NA	1.000	-0.748	NA	1.000
Standard error	0.549	0.021	1.000	0.611	0.027	1.000
<i>Estimation characteristics</i>						
Hyperbolic discounting	0.039	0.068	0.326			
Exponential discounting	0.005	0.028	0.068	0.057	0.094	0.334
Delay	0.000	0.002	0.037	0.002	0.008	0.134
Front-end delay	0.013	0.040	0.132	0.173	0.115	0.775
Lab experiment	0.153	0.102	0.766			
<i>Experimental characteristics</i>						
Real reward	-0.005	0.026	0.072	-0.004	0.053	0.114
Matching task	0.016	0.045	0.152	0.158	0.113	0.747
Health domain	0.346	0.089	0.992	0.304	0.147	0.889
Other domain	0.442	0.071	1.000	0.658	0.100	1.000
Money domain * non-linearity correction	0.001	0.035	0.090			
Negative framing	-0.146	0.107	0.724	-0.007	0.034	0.098
Neutral framing	0.003	0.030	0.042	0.053	0.152	0.165
<i>Subject pool characteristics</i>						
Sample size	0.075	0.014	1.000	0.166	0.032	1.000
Students	0.877	0.111	1.000	1.184	0.178	1.000
Students * Lab experiment	-0.752	0.144	1.000	-0.992	0.142	1.000
Males only	0.012	0.050	0.082	0.081	0.152	0.284
Females only	-0.001	0.021	0.037	0.011	0.062	0.086
North America	0.011	0.039	0.116	0.037	0.076	0.260
Asia	0.382	0.104	0.989	0.288	0.169	0.831
Africa	3.169	0.117	1.000	3.146	0.200	1.000
<i>Publication characteristics</i>						
Citations	-0.003	0.011	0.089	0.000	0.013	0.080
Publication year	0.121	0.027	1.000	0.151	0.041	0.994
Observations	927			568		
Studies	56			28		

Notes: Response variable = annualized estimates of the discount rate. SD = standard deviation, PIP = Posterior inclusion probability, SE = standard error. We employ Bayesian model averaging (BMA) using unit information prior (Eicher *et al.*, 2011) and the dilution prior suggested by George (2010), which accounts for collinearity. In the first specification we include variable *Money domain * non-linearity correction*, interaction of *Money domain* with a correction for non-linearity of utility functions; in the second specification we estimate a model on a subsample of estimates for which the exact time horizon is coded (which reduces the number of observations to 568 and eliminates variables *Hyperbolic discounting* and *Lab experiment* due to high collinearity). All variables are described in Table 5.