

Measuring the distribution of spitefulness

Citation for published version (APA):

Kimbrough, E. O., & Reiss, J. P. (2012). *Measuring the distribution of spitefulness*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 039
<https://doi.org/10.26481/umamet.2012039>

Document status and date:

Published: 01/01/2012

DOI:

[10.26481/umamet.2012039](https://doi.org/10.26481/umamet.2012039)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Erik O. Kimbrough, J. Philipp Reiss

**Measuring the Distribution of
Spitefulness**

RM/12/039

METEOR

Maastricht University School of Business and Economics
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands

Measuring the Distribution of Spitefulness*

Erik O. Kimbrough[†] J. Philipp Reiss[‡]

June 25, 2012

Abstract

Spiteful, antisocial behavior undermines the moral and institutional fabric of society, producing disorder, fear and mistrust. Previous research demonstrates the willingness of individuals to harm others, but little is understood about how far people are willing to go in being spiteful or their consistency in spitefulness across repeated trials. Our experiment is the first to provide individuals with repeated opportunities to spitefully harm anonymous others when the decision entails zero cost to the spiter and cannot be observed by the object of spite. This method reveals that the majority of individuals exhibit consistent (non-)spitefulness over time and that the distribution of spitefulness is bipolar: when choosing whether to be spiteful, most individuals either avoid spite altogether or impose the maximum possible harm on their unwitting victims.

*We thank Björn Bartling, Dan Friedman, workshop participants at Simon Fraser University, two anonymous referees, and an associate editor for helpful comments. This research was funded by grants from Maastricht University's METEOR research program in Economic Behavior, Theory and Computing.

[†]Department of Economics, Simon Fraser University, Burnaby, BC, Canada; e-mail: ekimbrough@gmail.com

[‡]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands; e-mail: p.reiss@MaastrichtUniversity.nl

Introduction

The Stanford Prison Experiment revealed the startling facility with which individuals lapse into antisocial and sadistic behavior when given the means and opportunity [1, 2]. Given the chance, individuals may readily abandon the peaceable character common to their daily lives and systematically and brutally mistreat others. Similar destructive tendencies are evident following breakdowns in traditional mechanisms of social control, as in the looting and indiscriminate vandalism that often follow in the aftermath of natural disasters and political demonstrations [3]. Under less extreme conditions, scattered acts of spitefulness nevertheless occur frequently, from locals pettily misleading tourists to children bullying their smaller peers.

Explanations for these (mis-)behaviors variously emphasize the importance of situational factors (e.g. dehumanization, differences in relative power, anonymity, action unobservability) [1, 2, 4, 5, 6] and individual characteristics such as personal history and personality [7, 8, 9]. However, evidence from twin studies also suggests that anti-social behavior has a strong genetic component [10].

When we call an action “spiteful”, we mean that it directly imposes harm on another and provides no immediate benefit to the spiteful actor. Our notion of spite differs from that typically employed by evolutionary biologists in that the latter require the spiter to undertake an (expected) cost when reducing the relative fitness of the other [11, 12]. Costly spite has been observed in some non-human species, e.g. social insects [13] as well as in humans [14, 15, 16, 17, 18].

However, many empirical studies of spiteful behavior with human subjects suffer from identification problems. In the case of the classic social-psychology research [1, 2], it is unclear to what degree spiteful actions are undertaken as a result of implicit or explicit experimenter demand as opposed to individual desire to do harm. In ultimatum games, there is debate over the motivations for decisions to reject non-zero offers (see e.g. [19]). However, since the decision to reject an offer is all or none, even if rejections represent spite (as claimed in [16]), it is still impossible to measure the *extent* of spitefulness. In the public goods games reported in [17], it is unclear whether the observed behavior is spite or merely an attempt to signal current dissatisfaction with the goal of promoting future cooperation. Recent experiments on costly ‘antisocial punishment’, in which some individuals actively punish cooperative others has been observed in a broad range of

cultures and environments, improve upon these studies because they directly measure the extent to which individuals are willing to endure costs to impose harm on others [20, 15], and recent models suggest that such behavior may be a result of selection [21]. However, in environments where spite is costly, individuals face an unobservable tradeoff between the costs and benefits of being spiteful. The presence of this tradeoff complicates inference about spiteful strategies because measured spite will be sensitive to the relative costs of spite to the spiter and to the target.

One advantage of studying spite in auctions, particularly second price auctions, is that spite can be measured in the intentional increase of the price that another bidder must pay. This element was recognized in previous studies of spite in auctions. For example, [22] report an experiment in which subjects in two-bidder, asymmetric second-price and ascending bid auctions with complete information on other bidders' values. They observe that lower value subjects overbid their values more frequently than higher value subjects. Similarly, in the auctions reported in [18], overbidding one's value may be explained by spite, but in both experiments, subjects submit bids without knowledge of the current high bid so that these bids still imply some risk of winning the auction. This feature inhibits measurement of spite as some overbidding may also be explained by judgment errors and/or a desire to win the item, regardless of cost.¹

Evidence from mosquitos indicates that when the costs are borne only by the target, spiteful behaviors can and will persist [24]. In our experiment, spite is also costless to the spiter, so that we can directly observe the underlying willingness to do harm.

Although (non-)spitefulness is a prominent behavioral pattern, little is known about how *observed* harm compares to the maximum harm that could have been done or to what degree (non-)spitefulness is stable within individuals over repeated trials. To isolate these aspects of spite, we report an incentivized laboratory experiment, holding situational factors constant, in which we can observe spiteful behavior, and we ask: when individuals have means and opportunity to anonymously harm others at zero personal cost, what is the prevalence, extent and individual consistency of spitefulness? Is there heterogeneity of spitefulness across individuals?

The remainder of the paper is organized as follows: In the next section, we introduce the experimental design. The following sections report our results and a discussion of their

¹See [23] for an earlier paper modeling spiteful preferences in auctions which also potentially explains overbidding in first-price auctions.

implications. Then, we provide a detailed description of our experimental procedures, and an appendix provides a copy of our instructions.

Experimental Design

To generate opportunities to observe spiteful behavior and measure the extent of spitefulness, subjects participate in a sequence of 16 market periods in each of which they attempt to buy a single unit of a fictitious item. In each of the 16 periods, one unit of supply is available for purchase in a two-stage auction, and unlimited supply of an identical item is available at a fixed price after the auction ends. In the first stage of the auction, bidders are informed about their value and the fixed price at which they can buy after the auction, and they submit an initial bid for the auctioned item. In the second stage, everyone is informed about the highest initial bid, and subjects submit their final bids. The item is allocated to the highest bidder at a price equal to the *second highest* final bid. Subjects must submit a final bid, and it must be at least as high as their initial bid, but they are permitted to resubmit the same bid.² Those subjects who do not purchase an item at auction can then purchase it in the aftermarket at the fixed price.

We call a bidder *spiteful* if she increases the final auction price with no intention of winning the auction. By submitting a final bid higher than her own initial bid but also lower than the *highest* initial bid, a spiteful bidder can increase the price paid by the winning bidder (thereby reducing that person's earnings) at no cost to herself. Typically in auctions, spiteful bidding may be inhibited by bidders' uncertainty about the current highest bid. Without this information, making a bid with the goal of driving up the auction price entails some risk of winning the auction and paying the high price as, e.g., on eBay. We eliminate this risk of unintentionally winning the auction by announcing publicly - just before bidders submit their final bids - the highest bid submitted in the first stage of the auction. This feature of our experiment not only enables us to observe spiteful behavior *per se* but also allows us to quantify the extent of feasible spitefulness: the maximally spiteful bid is the largest bid that avoids winning the auction.³

We informed subjects only about the highest initial bid and *not* the second-highest

²These auctions are stylized versions of those employed by the auction website eBay [25, 26].

³One crucial element distinguishing our setup from [22], is that since bidders in their experiments know others' values but remain uncertain about their *bids* until the auction ends. Thus, a spiteful bid in their environment still entails some risk and does not permit measurement of relative spitefulness.

initial bid (which would determine the price if the auction ended then), so they could not differentiate between auction prices that were set by competitive bidding in the initial stage or by spiteful bidding in the final stage. This prevented subjects from conditioning their behavior on having been harmed because they could not even know if they were the object of spite. This design element rules out, e.g., positive and negative reciprocity, which were argued to drive behavior in [22].

After reviewing the highest initial bid, a subject who is not the high bidder in the first stage can choose whether to engage in spite (by driving up the price) or not (by keeping his bid constant), but he may also choose to increase his bid above the highest initial bid in an attempt to win the auction. To measure the extent of spitefulness, it is essential that subjects frequently encounter a decision where the latter action is undesirable, or else we would rarely observe a choice between spiteful and non-spiteful bidding. To increase the likelihood that a subject may decide whether to be spiteful, we provide each bidder not winning the auction with the opportunity to buy an item - identical to the auctioned item - at a fixed price after the auction. Thus if the highest initial bid meets or exceeds the (expected) fixed price, a bidder should never choose to buy at auction and will instead wait to buy the identical item later on, but nevertheless, each bidder has to submit a final bid.

We collected all initial bids before providing subjects with any feedback and before beginning the second stage of any auction. This removes the possibility that spitefulness can be justified as a way to teach other subjects that submitting unreasonably high initial bids can be a costly mistake. We provided feedback on the auction outcome after all bidders submitted final bids. This provided subjects with many opportunities to observe how spiteful bids affected the winning bidders' earnings so that repeated submission of spiteful bids by the same subject cannot be dismissed by inadvertency.

We designed this experiment as part of a research program on price formation in auctions, unrelated to spiteful behavior (See [27]). For this reason, our design incorporates an individual choice "real effort" task between the auction stages and an extra decision, prior to submitting the final bid, in which subjects may choose whether to recall the fixed price at which they may purchase in the aftermarket or to earn money in the "real effort" task.

Each experimental session consists of three stages: an "Opening" stage in which they submit initial bids in the auction; an "Effort" stage in which they earn money by com-

pleting a real effort task; and a “Closing” stage in which subjects submit a final bid in the auction and then subjects who are unable to purchase an item at auction.

In the “*Opening*” stage, each subject submits a sequence of 16 initial bids, one for each auction. For each auction, subjects observe their induced value for the item and the fixed price at which they will later be able to purchase. After receiving this information for the first auction, they submit their first bid. Then the process repeats until they have submitted initial bids for all 16 auctions.

Following the initial bids, subjects enter the “*Effort*” stage. Here each subject participates in 3, two-minute periods of a real-effort ‘slider task’ [28] with a break of one minute in between. Subjects observe a screen with 48 sliders, each representing a scale from 0 to 100. The sliders are initially set to “0”, and subjects receive a payment $k^* = 5$ for each slider that they set to “50” by the end of the two minutes. (We expect subjects being capable of correctly placing 15-20 sliders in a two-minute period leading payments of 1.50-3.00 EUR over the three tasks.)

At the end of the Effort stage, subjects enter the third and final stage of the session in which each auction ends in the same sequence in which the subject submitted bids in the *Opening* stage. The “*Closing*” stage is divided into two sub-stages *for each auction*, which we call the “*Market*” and the “*Aftermarket*”. In the *Market* sub-stage, subjects observe their bid and whether they are currently the highest bidder. Then, they may submit a new bid if they desire. However, if they want to recall the fixed price before making their bid, they must agree to forgo the opportunity to participate in another minute of the slider task. Hence, by varying the value of k in the slider task, we vary the opportunity cost of recalling the fixed price, and we can observe the effect of opportunity cost on the probability of overbidding in the auction. We do not discuss the findings related to this hypothesis here, but see also [27], for a discussion.

After each subject submits a new bid (or resubmits the original bid), those subjects who did not choose to observe the fixed price participate in one minute of the slider task in which each correctly placed slider yields a return of k_{group} . Those who choose to forgo the slider task must wait quietly. At the end of the minute, the auction immediately concludes with the highest bidder paying the second highest bid. Each subject learns whether they were able to purchase the good at auction, and those subjects who were unsuccessful are then given the opportunity to purchase in the *Aftermarket* sub-stage. In the *Aftermarket*, the fixed price is revealed to all the remaining buyers at no cost, and

subjects simply choose whether or not to buy at the revealed price. At the end of the first aftermarket, the second auction enters the *Closing* stage, and so on until all the auctions have closed.

Each auction consists of $n = 3$ bidders chosen from a matching group of size $N = 12$. For each auction, each subject receives a randomly drawn integer value, $v \sim U[500, 1000]$, and we draw the fixed price, $p \sim U[300, 500]$, so that buyers will always buy in the aftermarket even if they are unable to buy at auction. Thus, we can observe spite in the absence of concerns about being unable to purchase an item. Over the sequence of 16 auctions, all four groups face the same values and fixed price in each auction. Our design also varies the value of k in the slider task (though this variation does not matter for the results we report here). Specifically, in each auction, one group of 3 bidders faces each of $k_{group} \in \{0, 5, 15, 30\}$. We rematch groups to ensure that each bidder faces each value of k four times.

Subjects are rematched across auctions and receive no additional information about the other bidders in their auction. The bidder’s role is framed as that of a seller facing the opportunity of buying several commodities in the auctions for resale to the experimenters, one commodity per auction. One advantage of this framing is that it better motivates participation in multiple auctions. In the instructions, we inform subjects that they will learn the fixed price in the *Opening* stage and that they “will have the option to review the fixed price before submitting a new bid in the *Closing* stage.” This language avoids the implication that the opportunity is *freely* available without directly revealing our treatment variation. We also inform subjects that “if [they] are unable to buy the item in the auction, [they] will be able to purchase an identical item during the *Aftermarket*.” See the online supplementary information for a complete set of instructions.

Results

We define *potentially spiteful* bids as all final bids that were not directed at winning the auction (i.e. were lower than the highest initial bid). Subjects submitted potentially spiteful bids in 383 out of 768 instances, so we have ample opportunity to identify spite. We call a final bid *spiteful* if it is *both* greater than the bidder’s own initial bid and less than the highest initial bid so that it necessarily increases the price paid by the auction winner. We call a final bid *maximally spiteful* if it is exactly one bidding increment less

than the highest initial bid (e.g. if the highest initial bid was 100, a maximally spiteful bid would be 99). Such a bid maximizes the loss a spiteful bidder can impose on the auction winner without risk of winning the auction. We observe abundant spite as 67.6% of all potentially spiteful final bids are actually *spiteful*. Furthermore, 31.1% of potentially spiteful bids are maximally spiteful. The frequency of observing a spiteful bid is roughly constant over trials, but the frequency of maximally spiteful bids is increasing (Figure 1).

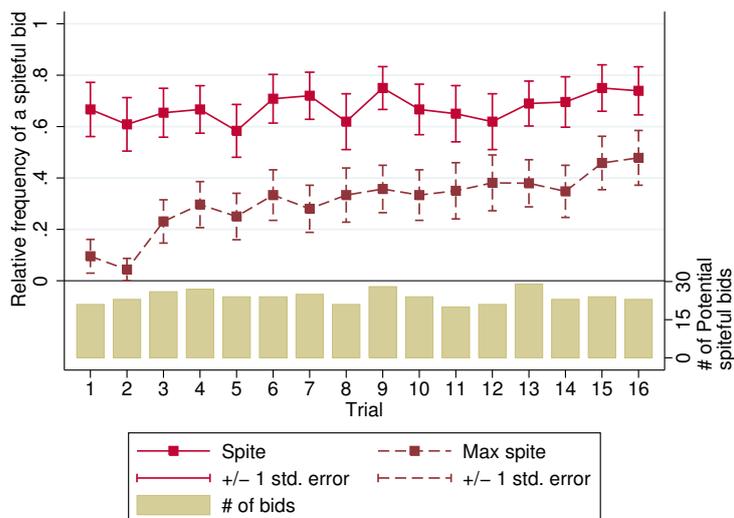


Figure 1: **Time series of the relative frequencies of spiteful and maximally spiteful bids, conditional on bids being potentially spiteful.** The solid line indicates the observed relative frequency of subjects making a spiteful bid in each trial, and the dashed line indicates the observed relative frequency of a maximally spiteful bid. Capped spikes display the standard errors of the mean in each trial. The subgraph displays the number of opportunities to observe a spiteful bid in each trial, and suggests that the relative frequency of observing spite is not related to the number of opportunities to be spiteful.

To facilitate inter- and intra-subject comparisons of the *extent* of spitefulness, we compute the share of observed harm imposed (price increase) by subject i in trial t out of the maximum possible harm in that trial. Denote the initial bid and final bid by $b_{i,t}^{1st}$ and $b_{i,t}^{2nd}$ and the highest initial bid by b_t^{max} . Spitefulness is defined as the actual bid increase

divided by the maximum possible bid increase that implies no risk of winning the auction,

$$S_{i,t} = \frac{b_{i,t}^{2\text{nd}} - b_{i,t}^{1\text{st}}}{(b_t^{\text{max}} - 1) - b_{i,t}^{1\text{st}}}$$

conditional on a subject not attempting to win the auction, so that $S_{i,t} \in [0, 1]$. The distribution of spitefulness indicates that we not only observe spiteful bids with striking frequency, the distribution of spitefulness is also highly skewed to both tails, indicating that most bids are either not spiteful at all or maximally spiteful (Figure 2). Note that for a bidder who has decided *not* to try to win the auction, any bid less than the first stage high bid is, strictly speaking, a weak best response. However, if subjects were adopting such a strategy, we would expect them to choose each value between their current bid and the current high bid with equal probability. Clearly from the u-shaped distribution of spitefulness in Figure 2, this is not the case.

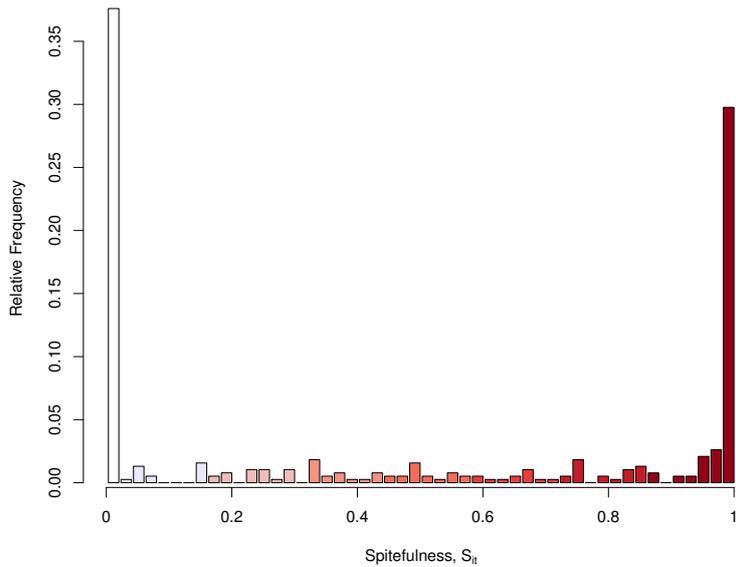


Figure 2: **Histogram of measured spitefulness ($S_{i,t}$), pooled over all subjects and trials.** The height of each bar represents the relative frequency of observing spitefulness in each of 50 intervals of length 0.02. Nearly 70% of the weight of the distribution is in the extreme values zero and one, indicating that, given the opportunity, subjects are either maximally spiteful or not spiteful at all.

Our measure of spitefulness $S_{i,t}$ does not distinguish between a maximally spiteful bid

that raised the final price by 2 and another that raised the final price by 200, although the extent of spitefulness may depend on the maximal harm that can be imposed. However, the data suggest that the decision to be maximally spiteful, as measured by $S_{i,t}$, is independent of the potential harm done (Figure 3). We present the data separately for men and women to control for a potential gender effect as there is mixed evidence on gender differences in antisocial behavior [10, 29]. Our data in both panels suggest the absence of a gender effect.

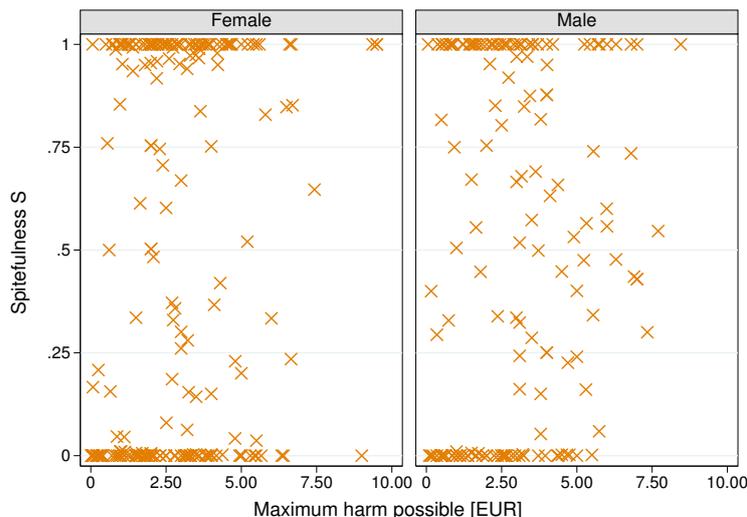


Figure 3: **The effects of the magnitude of maximum harm possible and gender on spitefulness.** The panels present the data separately for males and females. Each data point represents one observation of spitefulness relative to the magnitude of maximum harm possible.

Panel regression analysis ($n = 383$) confirms the absence of both magnitude and gender effects. We regress spitefulness on the magnitude of potential harm, a gender dummy, an interaction between female and magnitude of potential harm, the aftermarket price, and a time (trial) trend along with a constant. The estimates show that spitefulness is affected by neither the level of possible harm (p -value = 0.187) nor by gender or the interaction term (p -values = 0.615 and 0.708). We include the aftermarket price because it may potentially affect spitefulness. For example, when the aftermarket price is high, those bidders who do not buy in the auction may want to raise the price paid by auction winners

so that their earnings are not substantially different.⁴ However, we find that spitefulness is unaffected by the aftermarket price (p -value = 0.357); full regression output is available in Table 1. Consistent with the increase in spitefulness over time (Figure 1), the estimated time trend is positive and highly significant (p -value < 0.01).

Independent Variable	Coefficient	Rob. Std. Err.	z	p -value	95% conf. interval
Maximum Possible Harm	0.00015	0.00011	1.32	0.187	[-0.00007, 0.00036]
Aftermarket Price	-0.00020	0.00022	0.92	0.357	[-0.00063, 0.00023]
Female	-0.05767	0.11480	-0.50	0.615	[-0.28267, 0.16733]
Female·(Max. Harm.)	-0.00005	0.00014	-0.37	0.708	[-0.00033, 0.00022]
Trial	0.02105***	0.00652	3.23	0.001	[0.00827, 0.03383]
Constant	0.38353***	0.11886	3.23	0.001	[0.15058, 0.61649]

Table 1: **Random effects estimation of relative spitefulness** $S_{i,t}$. Significance levels are denoted by: ***0.001, **0.01, *0.05. The random effects error structure is included for individual subjects to control for repeated measurement. A positive and significant coefficient of Trial (t) indicates increasing spitefulness over time. Note also the insignificant coefficients on Maximum Harm, Aftermarket Price, Female and their interaction.

Finally, we examine the heterogeneity of spitefulness between and within individuals. While 5 subjects are maximally spiteful at every opportunity (i.e. $S_{i,t} = 1$ whenever spite was possible), 6 of them are never spiteful (i.e. $S_{i,t} = 0$ whenever spite was possible). Similarly, we find that 13/45 subjects are maximally spiteful at least 50% of the time, and 14/45 subjects are not at all spiteful at least 50% of the time. This is markedly similar to the observed distribution of spitefulness in Figure 2.

More generally, many subjects display striking consistency in their level of (non-)spitefulness (Figure 4). To further evaluate individual consistency, we estimate a simple linear regression, separately for each subject, where the dependent variable is spitefulness $S_{i,t}$ and the independent variable is trial t . We classify a subject as behaviorally inconsistent if the regression yields a significant estimate of the trial coefficient. By this criterion 73.3% of our subjects display consistent levels of (non-)spitefulness. Note that the individual stability of spitefulness appears to be inconsistent with the significant aggregate time trend noted in Figure 1 and Table 1; the reason we observe increasing average spitefulness over time is that, among the individuals whose level of spitefulness is *not* consistent over time, 10 out of 12 show an increasing trend. This creates an increase in the aggregate level of spitefulness, despite the stability of most individuals.

⁴[30] observe that in a modified dictator game, when individuals receive low offers, some subjects are willing to incur a small cost to substantially reduce the payoff of the other.

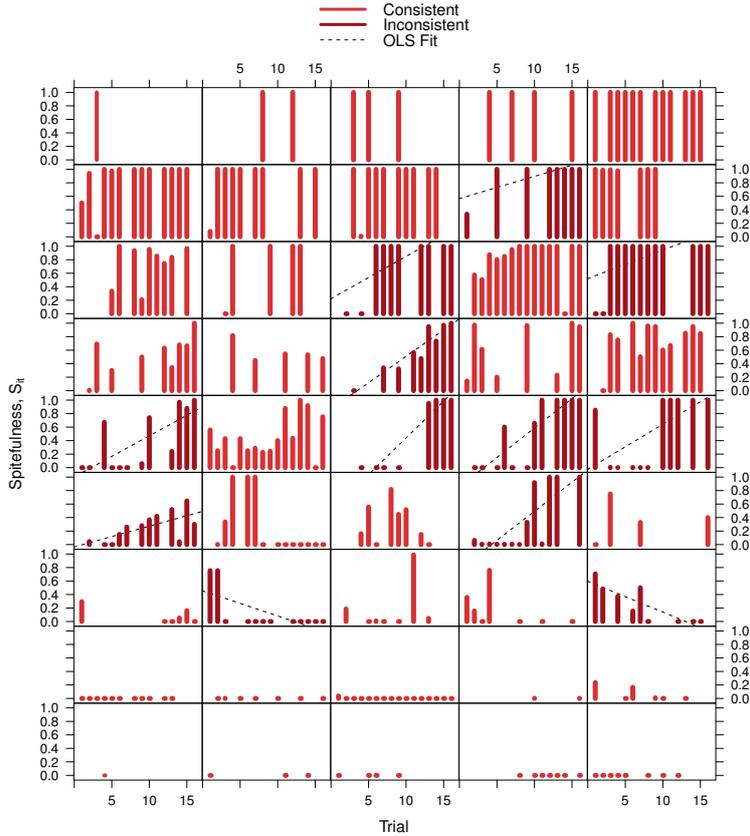


Figure 4: **Barplot time series of individual spitefulness ($S_{i,t}$)**. Each panel displays the spite data for a single subject. Panels are sorted from top to bottom by average spitefulness. Each bar represents the measured spitefulness of a subject in the indicated trial; if the bar has the length of a dot, the subject had the opportunity to act spitefully but chose to be non-spiteful. Each dashed line shows a linear fit to the data for subjects classified as behaviorally inconsistent, i.e. if the estimated coefficient of trial (t) is insignificant at a level of 5%. Out of 48 subjects, 45 had the opportunity to submit spiteful bids. The data reveal both considerable heterogeneity in spitefulness across individuals and consistency within individuals.

Discussion

One potential concern with our method of measuring spite is that some bids which we label spiteful may result from alternative bidding strategies. For example, bidders who discover that the initial high bid is greater than their value have a weak best response to bid their value, and if their initial bid is less than their value, this strategy will produce

a positive measure of spitefulness. Similarly, bidders who know the posted aftermarket price have a dominant strategy to bid the minimum of their value and the posted price; here too, for a sufficiently low initial bid, a final bid that follows this strategy will be measured as spiteful. Figure 5a provides a scatter plot of potentially spiteful bidders' final bids against their values, and Figure 5b plots potentially spiteful bidders' final bids against the aftermarket price. Neither values nor aftermarket prices appear to account for observed bids.

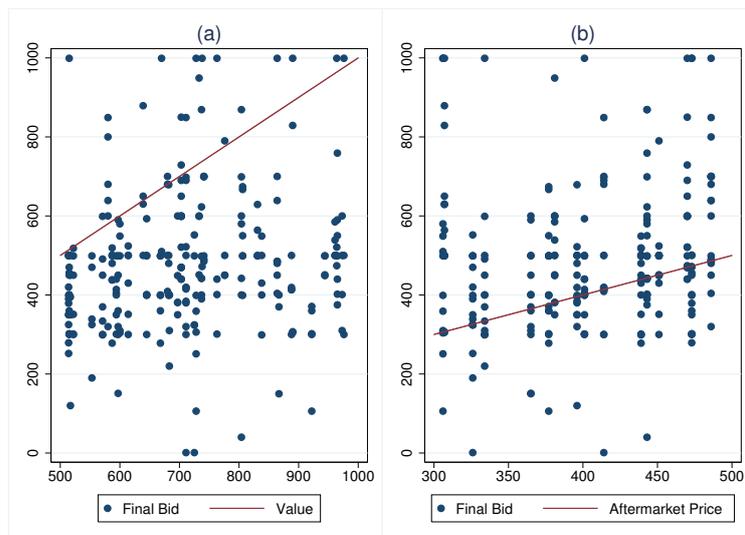


Figure 5: **Scatterplots of final bids submitted by potentially spiteful bidders.** Panel (a) shows the relationship between potentially spiteful final bids and values, and panel (b) shows the relationship between such bids and the aftermarket price. Each point represents a single bid, and the lines provide a reference showing where the bid equals each value or aftermarket price.

Our experiment places subjects in repeated situations in which they choose whether to spitefully harm an anonymous other at zero cost and in which spiteful acts are not revealed as such during or after the interaction. Because we have access to information not only about how much harm was imposed, but also how much harm *could have been* imposed, we are able to directly measure each individual's level of spitefulness. In this setup we find that spitefulness widely prevails, but its distribution is bipolar; typically we observe either zero or maximum spitefulness. The shape of the distribution is neither accidental nor generated by arbitrary behavior on the part of subjects. Instead the large

majority of subjects exhibit consistent spitefulness across trials suggesting the existence of a stable individual characteristic of (non-)spitefulness, conditional on circumstances.

Methods

Ethics Statement

All experiments were conducted with the informed consent of 48 healthy adult subjects who were free to withdraw from participation at any time. Only individuals who voluntarily entered the experiment recruiting database were invited, and informed consent was indicated by electronic acceptance of an invitation to attend an experimental session. The experiments were conducted following the peer-approved procedures established by Maastricht University's Behavioral and Experimental Economics Laboratory (BEElab). Our study was approved by the BEElab at a public ethics review and project proposal meeting that is mandatory for all scholars wishing to use the BEElab facilities.

Experimental Procedures

The experiments were conducted at the BEElab of Maastricht University with 48 students. Subjects' decisions were fully incentivized using Experimental Currency Units (ECUs). Their ECU-profits were converted to EUR at a rate of $100 \text{ ECU} = 1 \text{ EUR}$ and these from two randomly selected auctions were paid to them in cash at the end of the experiment. In order to avoid the influence of wealth effects, at the end of the experiment, we randomly select two auctions for each subject for payment. Then subjects receive private cash payments including a 4 EUR payment for arriving to the experiment on time, their earnings from the Effort stage, and their earnings from the two randomly selected auctions (including what they earned in the slider task, if they participated).

In total we ran 4 sessions of 12 subjects drawn from the undergraduate population of Maastricht University (Average age = 22.6, 46% Female). On average, subjects earned 15.32 EUR for a 60-minute session ranging from a low of 6.93 EUR to 26.65 EUR, including show-up payment. All sessions were conducted in May 2011.

Our experiment was programmed using the z-Tree software package [31]. Some graphics and data analysis were performed using *R*, an open-source statistical software [32].

References

- [1] Haney C, Banks WC, Zimbardo PG (1973) Study of prisoners and guards in a simulated prison. *Nav Res Rev* 9: 1-17.
- [2] Haney C, Banks WC, Zimbardo PG (1973) Interpersonal dynamics in a simulated prison. *Int J Criminol Penol* 1: 69-97.
- [3] Tilly C (2003) *The Politics of Collective Violence*. Cambridge: Cambridge University Press.
- [4] Merton R (1938) Social structure and anomie. *Am Sociol Rev* 3: 672-682.
- [5] Milgram S (1963) Behavioral study of obedience. *J Abnormal Soc Psychol* 67: 371-378.
- [6] Abbink K, Sadrieh A (2009) The pleasure of being nasty. *Econ Lett* 105: 306-308.
- [7] Robins L (1978) Sturdy childhood predictors of antisocial behavior: replications from longitudinal studies. *Psych Med* 8: 611-622.
- [8] Patterson G, DeBaryshe B, E R (1989) A developmental perspective on antisocial behavior. *Am Psych* 44: 329-335.
- [9] Miller JD, Lynham D (2001) Structural models of personality and their relation to anti-social behavior: a meta-analytic review. *Crim* 39: 765-798.
- [10] Rhee SH, Waldman ID (2002) Genetic and environmental influences on antisocial behavior: A meta-analysis of twin and adoption studies. *Psychol Bull* 128: 490-529.
- [11] Hamilton WD (1970) Selfish and spiteful behavior in an evolutionary model. *Nature* 228: 1218-1220.
- [12] Knowlton N, A PG (1979) An evolutionarily stable strategy approach to indiscriminate spite. *Nature* 279: 419-421.
- [13] Foster KR, Wenseleers T (2000) Spite in social insects. *Trends Ecol Evol* 11: 469-470.
- [14] Levine D (1998) Modeling altruism and spitefulness in experiments. *Review of economic dynamics* 1: 593-622.

- [15] Fehr E, Hoff K, Kshetramade M (2008) Spite and development. *Am Econ Rev* 98: 494-499.
- [16] Pillutla MM, Murnighan JK (1996) Unfairness, anger, and spite: Emotional rejections of ultimatum offers. *Org Behav Hum Dec Proc* 68: 208-224.
- [17] Cason T, Saijo T, Yamato T (2002) Voluntary participation and spite in public good provision experiments: an international comparison. *Experimental Economics* 5: 133–153.
- [18] Cooper DJ, Fang H (2008) Understanding overbidding in second price auctions: An experimental study. *Econ J* 118: 490-529.
- [19] Hoffman E, McCabe K, Shachat K, Smith V (1994) Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior* 7: 346–380.
- [20] Herrmann B, Thoni C, Gaechter S (2008) Antisocial punishment across societies. *Science* 319: 1362-1367.
- [21] Rand DG, Nowak MA (2011) The evolution of antisocial punishment in optional public goods games. *Nat Commun* 2: 1-7.
- [22] Nishimura N, Cason TN, Saijo T, Ikeda Y (2011) Spite and reciprocity in auctions. *Games* 2: 365–411.
- [23] Morgan J, Steiglitz K, Reis G (2003) The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy* 2: 1102–1127.
- [24] Hurst LD (1991) The evolution of cytoplasmic incompatibility or when spite can be successful. *J Theor Biol* 148: 269-277.
- [25] Ariely D, Ockenfels A, Roth AE (2005) An experimental analysis of ending rules in internet auctions. *RAND J Econ* 36: 890-907.
- [26] Roth AE, Ockenfels A (2002) Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *Am Econ Rev* 92: 1093-1103.
- [27] Kimbrough EO, Reiss JP Opportunity cost explains overbidding in auctions. Mimeo.

- [28] Gill D, Prowse V (Forthcoming) A structural analysis of disappointment aversion in a real effort competition. *Am Econ Rev* .
- [29] Hyde J (1984) How large are gender differences in aggression? a developmental meta-analysis. *Dev Psychol* 20: 722-736.
- [30] Houser D, Xiao E (2010) Inequality-seeking punishment. *Economics Letters* 109: 20–23.
- [31] Fischbacher U (2007) z-tree: Zurich toolbox for ready-made economic experiments. *Exp Econ* 10: 171-178.
- [32] Team RDC (2011) R: A Language and Environment for Statistical Computing. Vienna: R Foundation for Statistical Computing.

Supplementary Information: Experiment Instructions

General information

You are now participating in a decision making experiment. If you follow the instructions carefully, you can earn a considerable amount of money depending on your decisions and the decisions of the other participants.

This set of instructions is for your private use only. During the experiment you are not allowed to communicate with anybody. In case of questions, please raise your hand. Then we will come to your seat and answer your questions. Any violation of this rule excludes you immediately from the experiment and all payments. The research organisation METEOR has provided funds for conducting this experiment.

During the experiment the outcome of your decisions will be measured in ECU (Experimental Currency Units) instead of Euro. We will convert your total earnings into Euro at a rate of $100 \text{ ECU} = 1 \text{ Euro}$ at the end of the experiment and pay you in cash privately.

Outline of the Experiment

This experiment consists of 16 auctions of fictitious items, each of which will have some value to you; you can think of your value as the amount of money (in ECU) that the experimenter will pay you for the item if you buy it. So, if you have a value of 800 ECU

for the item, and you buy it in the auction, you will earn 800 minus the price ($800 - P$) ECU.

In each auction, your value will be a random number between 500 and 1000 ECU. Each value is equally likely. You will have a new random value in each auction.

There are three (3) bidders in each auction, and the two (2) other people in each auction will change randomly. Other people also receive random values between 500 and 1000 ECU for each item, and all values are determined independently of one another. You will only know your own value, and the other people will only know their own values.

At the end of each auction, the person who submitted the highest bid receives the item and pays for it a price equal to the second highest bid. If there is a tie for the high bid, the item is assigned to one of the bidders by chance, and the price paid is equal to the bid.

If you do not submit the highest bid, so that you do not receive the item in the auction, you can buy an identical item after the auction at a fixed price between 300 and 500 ECU, called the aftermarket price.

Anyone who did not buy the item in the auction can buy a unit of the item at the aftermarket price. This item is identical to the auctioned item so that you have the same value for it as for the auctioned item.

So, if the aftermarket price is 400 ECU and you choose to buy the item after the auction, you would receive $800 - 400 = 400$ ECU.

Each of the 16 auctions is divided into 3 stages:

- (1) *Opening Stage* each person submits a first bid for the item.
- (2) *Task Stage* each person participates in a task in which they can earn money by manipulating objects on the computer screen.
- (3) *Closing Stage* each person submits a second and final bid in the auction, and this bid cannot be smaller than their first bid. Everyone will learn whether they bought item at auction, and the people who did not buy at auction will have the opportunity to buy an identical item at the aftermarket price.

The instructions that follow will give you detailed information about stage 1. After completing stage 1, you will receive detailed instructions for stage 2 and stage 3.

At the end of the experiment, we will pay you the total of your earnings from the task in stage 2 plus your earnings from 2 randomly chosen auctions. In addition, you begin the experiment with an initial endowment of 400 ECU for arriving on time.

Stage 1: The Opening Stage

To begin the experiment, you will participate in the Opening stage for all 16 auctions.

In the Opening stage of each auction, you will see your value for the item and the aftermarket price at which, after the auction, you will be able to buy an identical item if you are not the high bidder in the auction.

After reviewing this information, you submit a first bid for the item. Your bid can be thought of as being the price that you want to pay for the item in the auction.

In the center of the screen, you will see a box in which you can type your first bid. Once you type in a first bid that you are happy with, you will click the button labelled “Submit First Bid”.

At the same time, the other people in your auction will submit their own first bids. Once all bids are submitted for the first auction, the Opening stage will begin for the second auction.

In the second auction, you will again see your value and the aftermarket price for a new item, and as before you will submit a first bid.

This process will repeat until you have submitted a first bid for all 16 auctions.

Later, in the 3rd stage of the experiment, you will see the highest first bid in each of the 16 auctions, and you will have the opportunity to submit a second and final bid in each auction. Note, your final bid cannot be smaller than your first bid, but you can submit the same bid again if you wish.

Stage 2: The Slider Task Stage

The Slider Task stage allows you to earn money from a computerised task.

You will now participate in 3 rounds of an identical task lasting 120 seconds per round with a break of 60 seconds between each round. Rounds 1 and 2 allow you to practice the task and will not affect your earnings. The score that you will have obtained in round 3 will be used for determining your total earnings from the Slider Task Stage.

The task will consist of a screen with 48 sliders. Each slider is initially positioned at 0 and can be moved as far as 100. Each slider has a number to its right showing its current position.

You can use the mouse or the keyboard's arrow keys in any way you like to move each slider. To use the arrow keys you have to select the slider that you want to adjust first by clicking it with the mouse. You can readjust the position of each slider as many times as you wish. You adjust the sliders to earn points. Specifically, your score in the task will be the number of sliders positioned at exactly 50 at the end of the 120 seconds.

Your score in round 3 will be converted to ECU at a rate of 1 point = 5 ECU. For example, if after 120 seconds you have correctly placed 6 sliders at 50, then you have earned 6 points, which is converted to 30 ECU and will be paid to you at the end of the experiment.

Thus your Slider Earnings are equal to (Your Score in round 3 x 5 ECU)

Are there any questions?

Stage 3: The Closing Stage

The Closing stage of the 16 auctions consists of two parts, the Market and the Aftermarket. In the Market part, each person submits a second and final bid for each of the 16 auctions. Before submitting your final bid, you will be informed of your value, your first bid, and the highest first bid. After reviewing this information, you will have the option to either:

- 1) Review the aftermarket price before submitting your bid

OR

- 2) Participate in another 30-second round of the Slider Task to earn ECUs.

If you choose to review the aftermarket price, you will not be able to participate in the Slider Task and you will have to wait quietly for 30 seconds. During the 30 seconds you will be able to see the aftermarket price at which the item will be available in the Aftermarket part.

On the other hand, if you choose to participate in the Slider Task, you will earn additional ECU based on your score from correctly adjusting sliders to 50, but you will not be able to review the aftermarket price before submitting your final bid. The amount of ECU that you receive per correctly adjusted slider will change over the course of the 16 auctions and you will be informed about it before you choose to participate, or not to participate - in the task.

At the end of 30 seconds, everyone will submit a final bid. On the bid submission screen, everyone will see their value, their first bid, and the highest first bid. After you review the available information, type in your final bid and click the button labelled “Submit Final Bid”.

After everyone in your auction submits a final bid, you will be informed of the outcome of the auction. Remember, the person who submits the highest final bid receives the auctioned item and pays a price equal to the second highest final bid. Therefore, your Market Earnings for each auction are computed as follows:

- if you are the high bidder:
 - Your Market Earnings = (Your value - 2nd highest final bid) ECU
- if you are not the high bidder:
 - Your Market Earnings = 0 ECU

If you are the high bidder so that you receive the auctioned item, then your income in this round is equal to your Market Earnings and you do not participate in the Aftermarket. However, if you are not the high bidder, then you will have the opportunity to buy an identical fictitious item in the Aftermarket as described in the following.

In the Aftermarket, you will be reminded of your value for the item (which is exactly equal to your value for the auctioned item) and the aftermarket price. If you want to buy the item at the aftermarket price, then you click the button labelled “Buy”.

Your Aftermarket Earnings are determined as follows:

- if you choose to buy the item in the aftermarket:
 - Your Aftermarket Earnings = (your value - the aftermarket price) ECU
- if you choose not to buy the item in the aftermarket:
 - Your Aftermarket Earnings = 0 ECU

At the end of the Aftermarket, the Closing stage is over and the auction is complete. Then we repeat the Closing stage until all 16 auctions are completed.

Your Earnings

Any auction can end in exactly six ways, and your Auction Earnings for an individual auction for each of these cases are as follows:

1. You win the item in the auction, and you participate in the Slider Task
Auction Earnings = (your value - 2nd highest final bid + Slider Task Earnings) ECU
2. You win the item in the auction, and you do not participate in the Slider Task
Auction Earnings = (your value - 2nd highest final bid) ECU
3. You do not win the item at auction, you participate in the Slider Task, and you choose to buy an item in the Aftermarket.
Auction Earnings = (your value - aftermarket price + Slider Task Earnings) ECU
4. You do not win the item at auction, you do not participate in the Slider Task, and you choose to buy an item in the Aftermarket.
Auction Earnings = (your value - aftermarket price) ECU
5. You do not win the item at auction, you participate in the Slider Task, and you choose not to buy an item in the Aftermarket.
Auction Earnings = Slider Task Earnings ECU
6. You do not win the item at auction, you do not participate in the Slider Task, and you choose not to buy an item in the Aftermarket.
Auction Earnings = 0 ECU

Total Earnings

At the end of the experiment, we will randomly draw 2 out of the 16 auctions for payment and sum your earnings from the 2 auctions. If your auction earnings are positive, then your ECU balance increases by this amount. If your auction earnings are negative, you made a loss and your ECU balance decreases by the amount of the loss. If your auction earnings are 0 ECU, then your ECU balance does not change.

Finally, we will add your earnings from the Slider Task stage to your ECU balance and we will convert it to Euro at a rate of 100 ECU = 1 Euro. When the experiment is over, please wait quietly until we call you for payment.