Vrije Universiteit Brussel

Investigating the origins of redistributive behaviour in non-ergodic contexts

Vanhoyweghen, A.; Macharis, C.; Ginis, V.

Published in:
Proceedings of the Royal Society A

DOI:
10.1098/rspa.2025.0598

Publication date:
2025

License:
CC BY

Document Version:
Accepted author manuscript

Link to publication

Citation for published version (APA):
Vanhoyweghen, A., Macharis, C., & Ginis, V. (2025). Investigating the origins of redistributive behaviour in non-ergodic contexts. *Proceedings of the Royal Society A*, *481*(2326), Article 20250598. https://doi.org/10.1098/rspa.2025.0598

Copyright
No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy
If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

Download date: 14. Jul. 2026

PROCEEDINGS A

royalsocietypublishing.org/journal/rspa

Research

Article submitted to journal

Subject Areas:
Complexity

Keywords:
Ergodicity, Human Decision-Making, Redistribution

Author for correspondence:
A. Vanhoyweghen
e-mail: arne.vanhoyweghen@vub.be

Investigating the origins of redistributive behaviour in non-ergodic contexts

A. Vanhoyweghen[1,O1], C. Macharis[2,O2] and V. Ginis[1,3,O3]

[1] Data Analytics Lab, Vrije Universiteit Brussel, 1050 Brussel, Belgium
[2]House of Sustainable Transitions, Vrije Universiteit Brussel, 1050 Brussel, Belgium
[3]School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA
[O1]ORCID:0000-0003-0103-4715
[O2]ORCID:0000-0003-0702-0496
[O3]ORCID:0000-0003-0063-9608

This study explores the behaviour of individuals regarding redistribution in a non-ergodic environment. By incorporating insights from geometric Brownian motions, we set up a stated preference experiment to shed light on decision-making processes in dynamics where time-averages diverge from ensemble averages. By focusing on time-averages, we propose a novel perspective on redistribution, challenging traditional expected utility theory.

To strengthen our findings, we conducted two independent iterations of the experiment, each with 30 multiplicative gambles designed to juxtapose expected gains against the optimisation of time-average growth in the context of redistribution. With these bets, we examine the tendency of human decision-makers (n=200 per experiment) to redistribute potential gains and losses, motivated primarily by self-interested growth and risk management rather than altruistic motives.

© The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License http://creativecommons.org/licenses/by/4.0/, which permits unrestricted use, provided the original author and source are credited.

THE ROYAL SOCIETY PUBLISHING

Our findings support a potential inclination towards sharing, as evidenced by a willingness to adopt strategies that do not align with expected gains. The observed behaviour supports the hypothesis that an element of human decision-making prefers time-average optimisation over expected outcomes. This insight enriches our understanding of decision-making and raises questions about the applicability of conventional rationality models in dynamic environments.

1. Introduction

As social creatures, humans tend to collaborate beyond self-interest [1–3]. But why do humans display such acts of altruism? Traditionally, this tendency for seemingly altruistic collaboration is rationalised through reciprocity [4,5], which states that collaboration results in a net, sometimes delayed, benefit for all parties involved.

In [6], the authors challenge this notion by bringing our attention to the implications of non-ergodicity on the cooperation puzzle. The authors tell us that agents who optimise their time-average rather than the ensemble average are intrinsically motivated by gains and the variation of said gains. When viewed from this angle, collaboration becomes a way for entities to exchange expected gains for reductions in the variation, which decreases the expected value of returns but can increase the time-average growth rate. These cooperating and dually motivated entities might appear puzzling when judged only from the perspective of expected gains (i.e., expected outcomes), needing higher-order explanations to resolve the conundrum, but turn out to be perfectly rational from the standpoint of time-averages. Such a divergence between expected and time-average outcomes is not rare; it arises in many real-world processes where outcomes evolve over time, particularly under multiplicative dynamics or path-dependent conditions, where the experience of one agent over time diverges from the average across many agents.

Intrigued by the interaction between expected values, variances, and the time-average, we used the Geometric Brownian Motion as a baseline dynamic to set up a redistribution scheme for which traditional theories predict collaboration to be detrimental, while collaboration is required to maximise the time-average growth rate. As an illustrative example, imagine a repeated gamble that could increase your capital of 10 euros by 50% (multiplicative factor of 1.5) or decrease it by 40% (multiplicative factor of 0.6), both with equal likelihood. Given the opportunity to redistribute a relative amount of your gains and losses, τ, how much of your gains would you give up for an equal part of your losses to be covered by someone else?

This bet with an expected value greater than one, together with this redistribution scheme, juxtaposes both paradigms. From the perspective of expected gains, redistributing seems like a poor decision as you will always give up more potential gains than potential losses, here you are sacrificing 5τ euros to receive 4τ euros depending on the outcome of the coin flip. In other words, the higher the level of redistribution, the lower the expected gains. Consequently, a tendency for (no) redistribution is then attributed to idiosyncratic biases and heuristics, which are reflected in a person's utility function. From baseline Expected Utility Theory, we can make the following deductions for the proposed redistribution/insurance scheme: a risk-seeking person will not redistribute. In contrast, a risk-averse person will redistribute per the strength of their risk aversion.

However, from the perspective of growth rate optimisation, the tendency to redistribute is not irrational. Consider the formula for time-average growth $g_t = \mu - 1/2\sigma^2$ [7] within a continuous time multiplicative dynamic. Our redistribution/insurance scheme measures, in a sense, how much one is willing to give up on a bet's expected value to reduce the variance in its payouts, which in turn impacts the time-average growth rate. This trade-off has a sweet spot where the variance has been sufficiently reduced while maintaining an expected outcome greater than one. We will call this point optimal τ or the optimal cooperation rate. Figure 1 shows the impact of optimal redistribution on the experienced growth rate, which shows that

redistribution is beneficial for both the individual trajectories as well as for the society as a whole. As a consequence, redistribution can increase long-run growth rates in multiplicative settings, even when it lowers expected returns. Similar conclusions have been reached in theoretical work showing that pooling and sharing improve time-average outcomes and can form stable equilibria [8–10].

In this study, we focus specifically on whether participants' choices approximate the time-average optimum, by redistributing their gains and losses. We do not address the distinct question of whether people would sacrifice time-average growth in favour of redistributive motives.

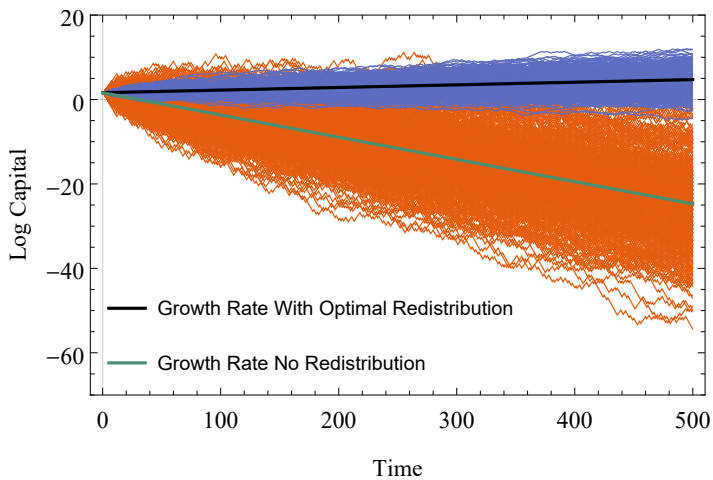

Figure 1: A Monte Carlo simulation of the wealth trajectories of both optimally redistributing (blue) and non-redistributing (red) agents (n=500) subject to the multiplicative gamble, which either increases wealth with a factor of 1.5 or decreases with a factor of 0.6. The non-redistributing trajectories experience a predominantly negative trajectory, losing approximately 5% of their capital for each time step. While the optimally redistributing agents display trajectories that have both less variance and have a slight upward trend, i.e. the median agent is slowly improving their capital position with a time growth of 0.6%. In doing so, the redistributing agents generally experience better wealth trajectories as individuals and as a collective.

Consequently, we use the multiplicative dynamic to set up a stated preference experiment which had the following three main characteristics: First, we introduced all 200 respondents to 30 multiplicative gambles, which would affect their capital proportionally based on the outcome of a coin flip and their respective multiplicative factors ($p = 0.5$, Outcome 1 = a) and ($p = 0.5$, Outcome 2 = b). All 30 bets had an expected multiplicative factor greater than 1, i.e. $(a + b)/2 > 1$. We then tasked respondents with maximising their capital gains by (partially) opting in or out of a proportional redistribution scheme. This proportional redistribution/insurance scheme allowed respondents to redistribute some or all of the gains and losses from a given gamble depending on the value they assigned to $\tau (\tau \in [0, 1])$, see figure 2. Mathematically, our adapted multiplicative dynamic can be described as follows:

$$x(t+1) = \begin{cases} (o_1 + (1 - o_1)\tau)x(t) & \text{w.p. } p \qquad \text{with } 0 \leq \tau \leq 1 \\ (o_2 + (1 - o_2)\tau)x(t) & \text{w.p. } 1 - p \quad \text{with } 0 \leq \tau \leq 1. \end{cases}$$

Intuitively, this redistribution scheme asks what proportion of your gains you will give up to cover an equal proportion of your losses. Note that the redistribution we introduced differs from the one most often discussed in the emerging field of Ergodicity Economics [7]. The type of cooperation considered in [6,11,12] consists of three steps: a wealth update step, a pooling

step, and a sharing step. In the first step, the agent simply plays the proposed multiplicative gamble and has their wealth updated in agreement with the stochastic outcome of the bet, after which everyone shares an amount with the population's wealth pool, and lastly, wealth gets distributed depending on how much they shared with the pool, as shown in figure 3. Notice that our redistribution scheme affects only the potential future gains and losses, not the entire capital. Ergo, it does not have the same (im)balancing effects. Parallels exist between this redistribution scheme and the Kelly criterion [13]. We can attribute this to the fact that, from a different viewpoint, one could interpret this type of redistribution as determining how much of your capital you're willing to risk for a particular bet. More on the link between the Kelly criterion and ergodicity breaking can be found in [14]. Moreover, our cooperation mechanism can also be framed as a special case of a fair insurance contract: each round, the chosen τ reduces both upside and downside by the same proportion, thereby lowering payoff variance while leaving expected value unchanged. Classical insurance setups, however, typically involve a wealth discrepancy between the insured and the insurer, with risk exchanged between unequal parties. By contrast, in our experiment, the redistribution was implemented symmetrically and without clear wealth heterogeneity. We therefore adopt the terminology of "cooperation," both for consistency with the framing presented to participants and to underline the conceptual link to risk-sharing in multiplicative dynamics.

Second, to further isolate humans' baseline intuitive tendencies within non-ergodic dynamics, all respondents played independently from one another. This separation eliminates any possibility of cheating, reputation seeking, or exchanging information about other participants. By design, it also brackets alternative drivers of redistribution widely discussed in the literature, such as warm-glow giving, kin selection, and reputation- or norm-based cooperation [1,15,16]. Respondents were informed only that if they redistributed a portion of their gains, the same proportion of their losses would be covered. At first glance, from the perspective of short-term expected values, trust in others seems unnecessary: in the immediate context of the game, choosing not to redistribute consistently leads to the highest personal payoff. As a result, under these conditions, the typical dynamics of trust and cooperation that give rise to prisoner's dilemmas or strategies like tit-for-tat [17,18] do not come into play. However, when examined through the lens of time-averages, where the long-term accumulation of wealth matters, trust and collaboration start to play a crucial role. Fant et al. [10] provide a theoretical framework for 'cheating' under non-ergodic dynamics, showing that collaboration remains evolutionarily stable over extended time horizons despite the immediate incentives to cheat.

Thirdly, every wager had an optimal redistribution rate, τ_{optimal}. Looking at it from the viewpoint of expected gains, the optimal redistribution rate was invariably no redistribution, $\tau = 0$. In the same paradigm, the worst choice was to redistribute everything, $\tau = 1$. Shifting our perspective to the lens of time-average growth, the worst possible amount to redistribute was either nothing, i.e., $\tau = 0$ if the geometric mean of the gamble is smaller than 1, $(a^p b^{1-p} \leq 1)$, or everything, i.e. $\tau = 1$ if the geometric mean is greater than 1. The time optimal choice to redistribute followed from the derivation described in section 2(a).

We find that respondents generally chose to redistribute, though not in a necessarily time-optimal way. These deviations may reflect the noisy signal of finite sequences, the particular tournament-style reward structure, or broader behavioural factors. Importantly, cooperation emerged in our setting even though channels such as fairness concerns or warm-glow giving were deliberately excluded, and redistribution was suboptimal from the perspective of the expected value.

2. Methods

This experiment took place on the online platform of Prolific to ensure a large enough response size and to reach a diverse group of respondents. In what follows, we will describe the procedure that generated bets, how the optimal value for τ was derived, and the employed sampling procedure.

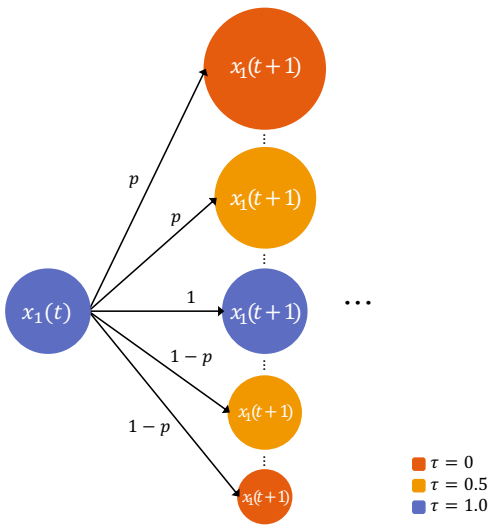

Figure 2: An illustration of the redistribution scheme in which τ represents what percentage of your gains you are willing to redistribute to offset an equal part of the potential losses. The closer τ is to zero, the more you stand to gain and lose; the circle's size reflects potential gains and losses. Conversely, the closer tau gets to 1, the smaller the bet variance. When $tau = 1$, this is equivalent to not playing the gamble.

(a) Defining the bets and their parameter space

To ensure measurably distinct outcomes between the ensemble average and time-average approach, we arranged the parameters in a specific way. All bets were designed to have an arithmetic mean greater than 1, which inherently makes redistribution unappealing from an expected benefit standpoint. We also ensured that all bets had a variance different from zero. This restriction ensures the uncertain nature of the payoffs and guarantees that redistribution could improve the time-average of the bets. The following formula gives the time-average growth rate: $g_t = (o_1 + (1-o_1)\tau)^p (o_2 + (1-o_2)\tau)^{(1-p)}$, which is a concave function, from which we obtain τ_{optimal}:

$$0 = \frac{\delta(g_t)}{\delta\tau} = p\,(1-o_1)\left(\frac{o_1 + (1-o_1)\,\tau_{\text{optimal}}}{o_2 + (1-o_2)\,\tau_{\text{optimal}}}\right)^{p-1} + (1-p)\,(1-o_2)\left(\frac{o_1 + (1-o_1)\,\tau_{\text{optimal}}}{o_2 + (1-o_2)\,\tau_{\text{optimal}}}\right)^{p},$$

$$0 = p\,(1-o_1) + (1-p)\,(1-o_2)\left(\frac{o_1 + (1-o_1)\,\tau_{\text{optimal}}}{o_2 + (1-o_2)\,\tau_{\text{optimal}}}\right),$$

$$\tau_{\text{optimal}} = \frac{o_1 - o_1 o_2 - p o_1 + p o_2}{-1 + o_1 + o_2 - o_1 o_2}.$$

We note that the expression we maximised above is the geometric mean of one-round growth factors. Taking logs yields the time-average (ergodic) objective for multiplicative wealth dynamics with i.i.d increments ($M(\tau)$):

$$g(\tau) = \mathbb{E}[\log M(\tau)] = p\,\log\big(1 + (1-\tau)(o_1 - 1)\big) + (1-p)\,\log\big(1 + (1-\tau)(o_2 - 1)\big).$$

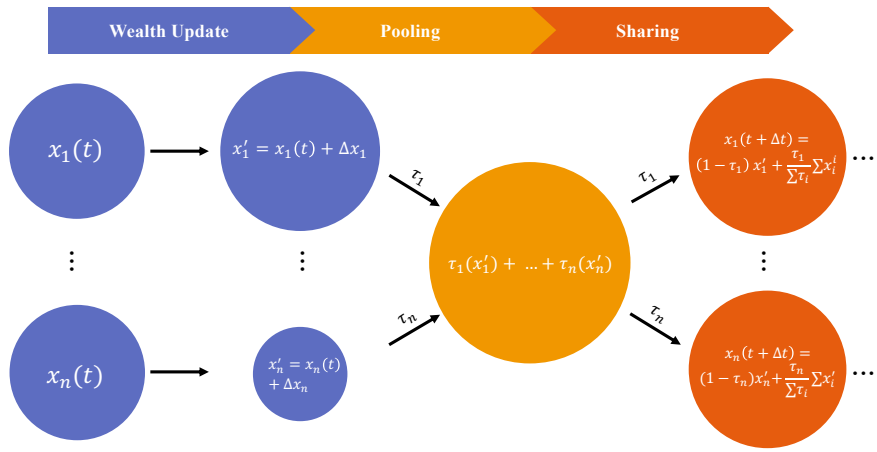

Figure 3: An illustration of the redistribution scheme for multiplicative dynamics as discussed in [6,11,12]. In this redistribution scheme, τ represents how much you contribute to the wealth pool in the pooling phase, which determines how much you receive in the sharing phase. The size of the circles reflects the amount of wealth.

The τ_{optimal} obtained above is identical to the optimiser of expected log-growth. In this approach the logarithm is used as the ergodic transform that renders the increments ergodic in a multiplicative process [19], not an idiosyncratic utility function.

A limitation on τ in our experiment is that it is impossible to have a redistribution rate higher than 1 or less than 0, which imposes the following restriction on the two outcomes (here parameterized as a function of o_1): $o_2 \in \left[\frac{1-po_1}{1-p}, \frac{o_1(1-p)}{o_1-p}\right]$. The final constraint was to set o_1 in the range $[1.1, 2]$ to avoid potential outcomes from becoming excessively extreme. The resulting parameter space is visualised in figure 4, from which we pulled 30 random bets. The formulated bets bear a strong resemblance to insurance risk protection, except that we opted to use gambles akin to a coin toss, where each outcome has a probability of 50% of occurring, to facilitate and simplify the communication towards respondents. Using gambles characterised by small but likely gains and huge but unlikely losses can introduce potential confounders, biases, and complications as described by prospect theory [20]. An overview of the employed bets can be found in Appendix 6(c).

(b) Respondents, implementation, and rewards

We targeted 200 respondents living in Europe, ensuring they had a comprehensive understanding of the euro's value. We presented each of the 30 bets to all participants in a random order. For participants' responses to be deemed suitable for subsequent analysis, they needed to express their preference for all bets and pass three attention checks. Out of the 200 participants, 187 satisfied these requirements. To confirm all respondents grasped the implications of

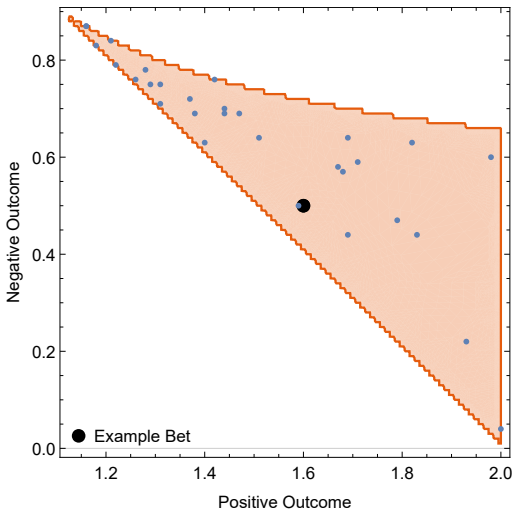

Figure 4: The parameter space defined by the constraints on the redistribution rate τ and outcomes o_1 and o_2. The shaded region represents the feasible parameter space, adhering to the requirements that ensure bets have an expected value greater than 1, a nonzero variance, and τ values that are bounded between 0 and 1. The 30 randomly selected bets are shown in blue, the larger black dot shows an example bet with as multiplicative outcomes 1.5 and 0.6.

redistribution, we provided a detailed written example and made a video available for further clarification, if needed. At the experiment's onset, we assigned each respondent to either be subject to extra time pressure or not, similar to the methodology used in [21,22]. The time pressure was imposed through a timer situated next to the bet, which started incrementing the moment a new bet was displayed. This timer functioned solely as a psychological instrument to spur intuitive decision-making, with no associated consequences, meaning any time pressure derived from this timer was purely perceived. All bets were accompanied by a slider going from 0% to 100%, on which respondents could indicate their preferred redistribution rate. In the original design of the experiment, we kept the starting position of the redistribution, the position at which the slider starts when presented to the respondent, fixed at 0%. This fixed starting position was changed after community feedback in a second iteration of the experiment, see section 2(c).
Respondents received a monetary reward for their participation in the experiment, in the form of an hourly wage of 11 pounds per hour weighed by the forecasted time required to complete the experiment, which was 15 minutes. 5 additional prizes were distributed to the respondents who managed to grow their capital the most over the course of the experiment, this to encourage deliberate responses from respondents rather than arbitrarily selecting a redistribution rate. To ensure a fair evaluation of capital growth, we fixed the outcomes for each coin flip for all respondents. While this reward mechanism increases attention and attaches real consequences to each choice, it might introduce higher levels of risk-seeking behaviour, resulting in lower levels of redistribution than would have been displayed using a different reward mechanism. We simulated the impact of tournament style rewards in (Appendix 6(b)). We find that this winner-take-most scheme mainly advantages more risk-seeking allocations when the rest of the field is highly risk-averse; otherwise, lowering τ does not improve or in fact lowers the odds of falling in the top-5.

We registered coin outcomes and analysis code; data/code on Zenodo. See Data Accessibility statement for more details.

(c) Community feedback and second iteration of the experiment

We presented our preliminary results at the conference Ergodicity Economics 2023, where we received community feedback regarding our experiment and setup. Specifically, they highlighted concerns about the possible anchoring effects [23,24] resulting from the redistribution slider's fixed starting position. To address this concern, we conducted the experiment a second time, this time with the redistribution slider's starting position being randomised. The second iteration of the experiment remained the same in all other aspects; 200 new respondents were recruited through Prolific, with none having participated in the original experiment. Out of the 200 respondents, 198 passed the attention checks and were included in the analysis.
In all subsequent sections, we will show both the results from the original setup, with a fixed starting redistribution rate of 0%, and the setup with a randomised starting redistribution rate.

3. Results:

(a) General redistribution patterns

We examined the redistribution behaviour exhibited by our respondents. In Figure 5 (left), we have collated all the revealed preferences for redistribution per bet. The blue and grey regions denote the 1st, 2nd, and 3rd quantiles of τ selected by respondents for a given bet in the original experiment and the second iteration, respectively. In addition, we have depicted the optimal τ value for each gamble, represented by the orange fences. The central point indicates the precise value for optimal τ, while the fences account for a 2.5% tolerance on this value. This 2.5% tolerance function works as follows: Given the entire spectrum of growth rates achievable through (no)redistribution yields as the maximum difference the delta between the minimal and maximal possible growth rates. We establish the fences by computing 2.5% of this total delta and subtracting it from the maximum growth rate. These fences signify the 2.5% tolerance level, indicating a range where the growth rate is no more than 2.5% less than the maximum potential growth rate achieved through optimal redistribution, given the full range of possible growth rates within our redistribution scheme.

Firstly, Respondents redistributed on average $23.8\%^{+4\%}_{-2.8\%}$ of their gains and losses in both iterations of the experiment, with the respondents in the second iteration of the experiment redistributing on average $27.6\%^{+4\%}_{-3.5\%}$. In both experiments, the chosen redistribution rate seems largely uninfluenced by the value for the optimal τ value. A substantial deviation occurs in the case of bet 14. This bet was the most volatile and polarising bet, featuring outcomes of 2 and 0.04. In the original experiment, the median respondent redistributed 40% of their earnings for this bet, improving the bet's time-average by a 0.5 net increase in growth rate. In the second iteration of the experiment, the median respondent redistributed 50%, improving the time-average growth rate by 0.6. Secondly, we highlighted each bet's least desirable τ value from the perspective of time-averages, which could be 0 or 1. The least desirable value for τ is 0 when the time-average for the given bet is less than 1. In contrast, the worst value for τ is 1 when the time-average of the gamble without redistribution exceeds 1.

(b) Comparing Respondents and Synthetic Agents

In Figure 5 (top right), we set out the normalised capital evolution from respondents within our experiment, together with the evolution of two synthetic agents, one that follows the perspective of optimising gains, i.e., the 'non-distributer' (red) and one that redistributes the time optimal amount the 'optimal distributer' (green). Additionally, we plot the theoretical expected evolution and the theoretical time-average evolution of capital for non-distributers, as described in [7], here in black and purple. Because of our dynamic's non-ergodicity, the non-distributor's experienced evolution falls short of attaining its expected evolution. On the contrary, the theoretical time-average growth predicts its capital evolution rather well, particularly when time increases.

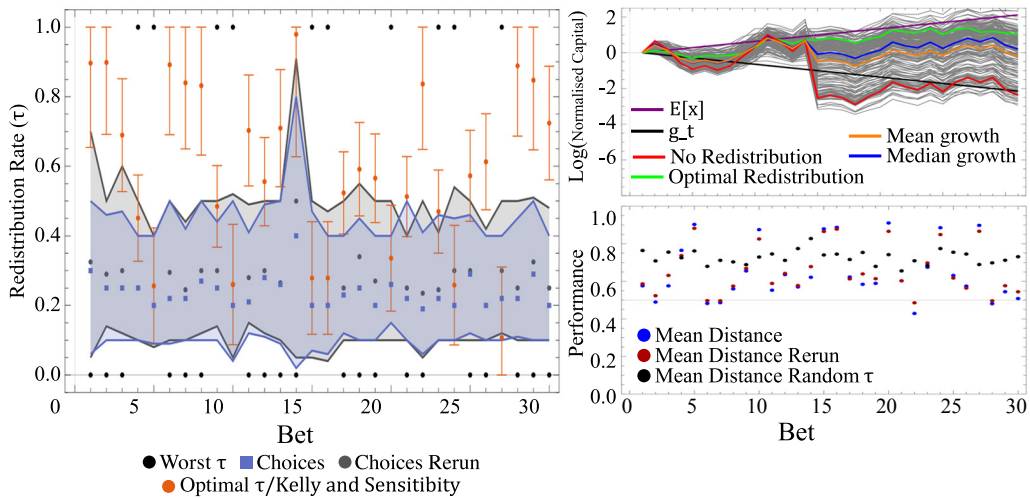

Figure 5: **The left** figure shows the distribution of redistribution rates (τ) chosen by participants across the two iterations of the experiment, highlighted in blue (original) and grey (rerun), blue/grey bands are interquartile range; dots are medians ($n_{\text{original}} = 187$, $n_{\text{rerun}} = 198$). Orange markers indicate the optimal τ for each bet, with the error bars representing a 2.5% tolerance range from the optimal. Choices stay remarkably stable and follow the same trend over both iterations of the experiment. These findings demonstrate a consistent preference for redistributing about 25% of gains and losses, largely unaffected by the optimal τ, except for the most volatile bet where a significant increase in redistribution rate was observed. The panel on **the top right** presents the evolution of capital from participants in the experiment, alongside the trajectories of two synthetic agents: the 'non-distributor' who does not engage in redistribution (red) and the 'optimal distributor' who redistributes the theoretically optimal amount (green). Also plotted is the arithmetic mean of the expected returns (black) and the geometric mean of the time-average growth rate for the non-distributor (purple), highlighting the discrepancy between expected and time-average growth. The trajectories of the participants' normalised capitals reveal a general tendency towards redistribution, with the median trajectory outperforming that of the non-distributor. The panel on **the bottom right** figure shows how well the average respondent redistributed, i.e., their performance, in the original experiment (blue) and in the rerun of the experiment(red). Performance is defined as the normalised distance between the growth rate achieved with their chosen τ and the worst possible choice, as shown in Equation 3.1. For comparison, the black dots show the performance of agents that choose their redistribution rate at random.

The respondents' trajectories show us that most respondents redistributed to a certain extent. Moreover, the median path indicates that most respondents redistributed in such a way that enabled them to remarkably outperform the 'non-distributor'. This preliminary result shows the propensity of most respondents to redistribute beyond the angle of expected benefits. A side-by-side comparison of the original and rerun samples is provided in Appendix 6(a); patterns are closely aligned with minor shifts consistent with sampling variability.

Given that most respondents redistribute far from the optimal amount 3(a), it might seem surprising that the median growth obtained by respondents comes remarkably close to the capital evolution obtained when redistributing optimally. To explain this, let's have a closer look at a respondent's performance or how close the redistribution behaviour of our respondents gets

toward optimal redistribution. We define a respondent's (j) performance for a bet (i) as follows:

$$\text{Performance}_{i,j} = \frac{g_{\tau\ \text{chosen},i,j} - g_{\tau\ \text{worst},i}}{g_{\tau\ \text{optimal},i} - g_{\tau\ \text{worst},i}} \tag{3.1}$$

Performance is a normalised distance measure that returns the distance of the obtained growth rate with respect to the worst and best obtainable growth rates for a given bet. A performance of 0 indicates that redistribution happened in the most sub-optimal way possible, and a performance of 1 indicates optimal redistribution. Figure 5 (bottom right) shows the average performance of the respondents in both iterations of the experiment. In both iterations of the experiment, the average performance of respondents is surprisingly similar, predominantly above 0.5, and consequently closer to optimal redistribution than the worst redistribution choice.

Although the metric is intuitive, it has an important limitation: scoring 0 is intrinsically hard. The worst redistribution rate is a single point, so almost any deviation from it—however small—raises the numerator of (3.1) and can already yield a substantial fraction of the total possible improvement. As a result, even unguided or random choices rarely land exactly on that worst point, driving expected scores well above zero. This effect can be seen in the high mean scores of the synthetic agents (black) that redistribute at random. These synthetic agents reach consistently high-performance rates and tend to outperform human respondents. In the absence of a widely accepted alternative, however, this normalised distance remains the most transparent and intuitive way to compare how close each redistribution choice comes to the optimum across heterogeneous bets.

Alternative Performance Metrics: regret and percent of optimal log-growth

To complement the normalised-distance performance score (Eq. 3.1), we also report two scale-free summaries of how close respondents' choices come to the time-optimal redistribution rate.

Regret. For each respondent we compute the regret in expected–log units,

$$\text{Regret}_j = g(\tau_{optimal}) - g(\tau_j),$$

where $g(\tau) = \mathbb{E}[\log X(\tau)]$ is the expected log-growth (time-average growth rate). As shown in Figure 6 (left), regrets remain well below 0.10 log-units for the vast majority of respondents, confirming that deviations from $\tau^\star$ translate into relatively small growth penalties.

Percent of optimal log-growth captured. We also compute

$$\text{Share}_j = \frac{g(\tau_j)}{g(\tau_{optimal,\,j})} \in [0, 1],$$

and aggregate at the respondent level as the overall proportion of the optimal log-growth attained. Figure 6 (right) shows that most respondents capture between 20–50% of the maximum attainable log-growth over the full experiment, with medians close to 40% in both the original and rerun samples. This view complements our earlier findings, by showing that while respondents do not redistribute optimally, their behaviour tends to secure a meaningful share of the available long-run growth, especially when considered w.r.t. to worst possible choices they could have made.

(c) The impact of τ_{optimal} on redistribution rates

We investigated whether the optimal redistribution rate (τ_{optimal}) influenced the displayed redistribution rate by ordering bets based on increasing τ_{optimal}. In the original experiment, we found a positive correlation between τ_{optimal} and both mean and median redistribution

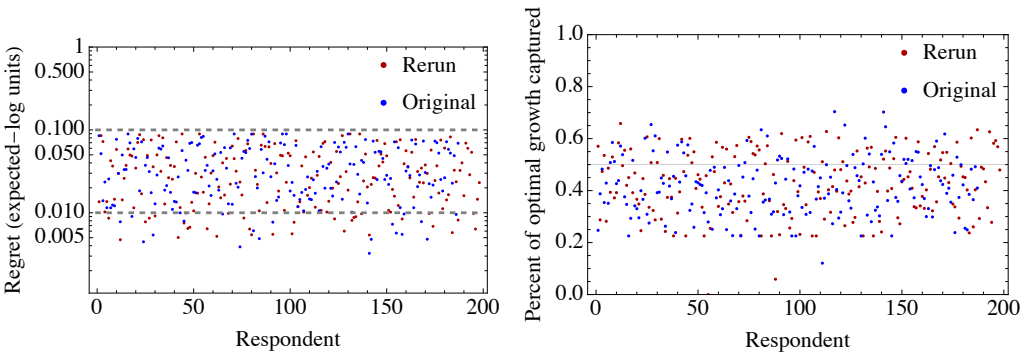

Figure 6: **Left:** Regret in expected–log units, defined as $g(\tau_{optimal}) - g(\tau_i)$, the shortfall relative to the bet-wise optimum. Most respondents show regrets well below 0.10 log-units, indicating that departures from $\tau_{optimal}$ entail only modest penalties on the log-growth scale. **Right:** Percent of optimal log-growth captured, $g(\tau_i)/g(\tau_{optimal})$, which expresses each choice as a fraction of the maximum achievable time-average growth. Respondents typically secure between 20–50% of the optimum, with medians around 40% in both samples. Together, the two panels show that while redistribution choices are rarely optimal, they consistently achieve a meaningful share of the available long-run growth. Red points denote the rerun sample, blue the original.

tendencies, with Spearman's rank correlations of 0.40 ($p = 0.03$) and 0.47 ($p < 0.01$), respectively. However, in the second iteration, this effect was no longer significant. Additionally, no clear relationship was found between τ_{optimal} and the decision to abstain from redistribution in either experiment, suggesting no consistent effect of τ_{optimal} on redistribution behaviour across both iterations.

(d) The impact of time pressure on respondent behaviour

Contrary to the findings of [21] and in line with the findings of [22], we didn't observe any influence of time pressure on response time, in either iteration of our study. Surprisingly, we also observed no effect of time pressure on redistribution patterns, which contrasts previous studies that found that time pressure [25,26] and intuitive responses promote cooperation [27–29]. In the iteration where the starting position was fixed, and timing was applied, the median response time recorded was 204 seconds for non-timed participants and 191 seconds for those timed. The Kolmogorov-Smirnov test yielded a p-value of 0.44, indicating no significant effect of time pressure on completion time. A similar trend was observed in the subsequent iteration, which involved a randomised starting value of the redistribution slider. Here, the median completion times were 197 seconds for participants not subjected to timing and 204 seconds for those who were, with the Kolmogorov-Smirnov test yielding a p-value of 0.95. Similar to [22], we attribute this to the inherent time pressure experienced by all respondents, which stems from the compensation structure adopted by Prolific. Prolific compensates its respondents based on the estimated survey completion time, irrespective of their actual completion duration. This model could induce a natural incentive for faster survey completion. We believe this could have overwhelmed the impact of the additional subliminal time pressure we imposed on respondents.

(e) The impact of starting value on redistribution tendency

We ran the experiment twice, once with a fixed starting redistribution rate and once with randomised initial redistribution rates for each bet. Our findings revealed no significant relationship between the initial position of the slider and the redistribution amounts selected by participants, as indicated by a Spearman Rank test with a p-value of 0.35.

4. Discussion

This manuscript aims to contribute to the growing body of research aimed at understanding human behaviour in non-ergodic dynamics. Previous experiments [21,22,30,31] within this body of work have been focused on human sensitivity towards ergodicity breaking as postulated in [14]. In [6], the authors illustrate that, explanations for collaboration do not need to go beyond its mathematical optimality when considering non-ergodic dynamics. In this new paradigm, the authors highlight that when viewed from the perspective of time-average, collaboration leads to higher time-average growth rates, while when judged from the expected values, naked collaboration leads to net losses, which need to be offset by second-order effects. We continue on this insight by exploring the redistribution tendencies of humans in a multiplicative environment, by means of a stated choice experiment. In this experiment, we introduced respondents to a bet which was expected to increase their wealth and asked them to indicate how big of a percentage of their potential gains they were willing to redistribute in order to offset an equal percentage of potential losses. We structured the experiment in this way to remove any possibility of unfair practices, such as choosing not to share gains or losses, by having participants engage with a virtual insurer that guaranteed all agreed-upon redistribution rates. This approach effectively neutralised considerations of strategic game behaviour (such as tit-for-tat or the Prisoner's Dilemma [32]), as well as factors like kin selection [16], the emotional satisfaction from giving (warm-glow) [15], and efforts to cultivate a good reputation [1]. This framework enabled a direct comparison of two perspectives, making the act of collaboration seem less appealing from the standpoint of expected values, since all forms of redistribution would diminish expected returns, while simultaneously highlighting its appeal through the lens of time-averages, where redistribution can enhance the time-average growth rate. It is possible that this contrast was further magnified by our reward structure, which rewarded the five respondents with the highest capital at the end of the experiment; this could have provided extra incentive towards high levels of risk-seeking and, as such, an aversion towards redistribution, which naturally lowers variance. Turning our attention to the empirical outcomes, we observe that in both iterations of the experiment, the majority of respondents redistribute even without the potential benefits tied to external factors. This finding contrasts with the expected benefit hypothesis, which suggests redistribution is a consequence of net benefit seeking. This is particularly noteworthy, considering our experiment was structured to make redistribution appear disadvantageous for optimising expected gains. Furthermore, when we compare both iterations of the experiment, we remark that respondents redistribute in a consistent but non-optimal way. The median respondents redistributed about 25% of their gains and losses and did so with surprising consistency over all proposed gambles, regardless of the optimal redistribution rate. A plausible explanation could be the diminishing signal-to-noise ratio when redistributing, where a modest amount of redistribution secures most of the variance reduction, after which the marginal benefit of further fine-tuning is shallow. In such flat regions, small calculation costs, judgment noise, and limited attention readily dominate, making coarse, portable rules attractive. This is consistent with satisficing, i.e. settling on a good-enough fraction rather than recomputing the exact τ_{optimal} per bet), and related heuristics. Importantly, these mechanisms are compatible with time-average reasoning but do not depend on it. Similar patterns could emerge from broader risk-management habits (aversion to drawdowns, simple variance caps), loss-sensitive preferences, anchoring or fairness norms that favour stable sharing rates across tasks. An interesting avenue for future work is to test more extreme bets, as suggested by bet 14, where the signal for redistribution is stronger and may overcome satisficing/anchoring, i.e., when the risk is sufficiently high. This would align with observations that the willingness to insure/pool increases with gamble volatility. [33]

A further feature of this study is that collaboration occurred between a human decision-maker and an ideal distributor who consistently reciprocated the respondent's decisions. In standard redistribution contexts, choices are often explained by perceptions of inequality, mobility, and government capacity (see Stantcheva [34]), or by social motives such as fairness norms, deservingness, and warm-glow giving [15]. Because respondents only interacted with

an ideal distributor, our design deliberately excludes such channels, yet participants still opted to redistribute. While this simplification helps manage interaction dynamics, it compromises realism. This limitation highlights the importance of a follow-up experiment where respondents can interact directly with each other. Such a more realistic setup introduces additional confounders, including game-theoretical adaptations like rewarding and punishing [17,32] and social effects such as reputation seeking.

A second limitation concerns the winner-take-most incentive scheme: only the five highest final capitals were rewarded. This could have incentivised respondents to display unusually high risk-seeking choices, which would, in turn, have tilted choices away from redistribution. As detailed in Appendix 6(b), this is mainly a concern for our setup when other respondents are (perceived as) highly risk-averse.

Moreover, this study used 50/50 gambles to facilitate communication with respondents. While this approach simplifies the experimental design, it might make the scenarios less reflective of real-world situations. In reality, many decisions involve probabilities that are not as straightforward as a coin toss. Future studies should consider using a wider range of probabilistic outcomes to better mimic real-life decision-making contexts. This could help us understand how people handle more complex risk scenarios and whether the observed cooperative behaviours persist under different probabilistic structures.

5. Conclusion

The findings in this paper support prior research on human decision-making in non-ergodic environments by suggesting that people are not as irrational as often thought. Instead, it highlights that human decision-making predominantly occurs in a non-ergodic realm, leading to the reproducible cooperative nature of intuitive decision makers [25,27–29]. Specifically, our results show that, on average, humans indeed tend to redistribute wealth 25% of their potential gains and losses even when warm-glow, kinship, and reputation/fairness channels are removed by design via a virtual insurer. This redistribution behaviour is, however, generally unaffected by the time-optimal redistribution rate, except in one notable case where the most volatile bet shifted responses closer to the optimal rate. This outlier highlights how future experiments with more extreme bets and models that account for the potential influence of biases and heuristics could provide further insights into cooperative behaviour in humans. Finally, this study offers a new perspective on assessing redistributing behaviour observed in humans within a non-ergodic dynamic.

Ethics. Prior to the initiation of the study, we sought and received ethical approval from the Ethics Committee for Human Sciences of the Vrije Universiteit Brussel (ECHW). The ethical approval reference number for our study is ECHW_356.02. Before starting the experiment, respondents gave their informed consent to the authors.

Data Accessibility. The computer code and data have been made available at `https://doi.org/10.5281/zenodo.17278074`

Authors' Contributions. A.V. and V.G. conceived the experiment, A.V. conducted the experiment and analysis, A.V. and V.G. analysed the results. A.V. and V.G. wrote the manuscript.

Competing Interests. The authors declare that they have no competing interests.

Funding. We acknowledge funding of the Vrije Universiteit Brussel

References

1. Fehr E, Fischbacher U. 2003 The nature of human altruism. *Nature* **425**, 785–791.
2. Melis AP, Semmann D. 2010 How is human cooperation different?. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**, 2663–2674.
3. Vlerick M. 2021 Explaining human altruism. *Synthese* **199**, 2395–2413.
4. Rand DG, Nowak MA. 2013 Human cooperation. *Trends in Cognitive Sciences* **17**, 413–425.

5. Nowak MA. 2006 Five rules for the evolution of cooperation. *Science* **314**, 1560–1563.
6. Peters O, Adamou A. 2022 The ergodicity solution of the cooperation puzzle. *Philosophical Transactions of the Royal Society A* **380**, 20200425.
7. Peters O. 2019 The ergodicity problem in economics. *Nature Physics* **15**, 1216–1221.
8. Yaari G, Solomon S. 2010 Cooperation evolution in random multiplicative environments. *The European Physical Journal B* **73**, 625–632.
9. Liebmann T, Kassberger S, Hellmich M. 2017 Sharing and growth in general random multiplicative environments. *European Journal of Operational Research* **258**, 193–206.
10. Fant L, Mazzarisi O, Panizon E, Grilli J. 2023 Stable cooperation emerges in stochastic multiplicative growth. *Physical Review E* **108**, L012401.
11. Berman Y, Peters O, Adamou A. 2020 Wealth inequality and the ergodic hypothesis: Evidence from the United States. *Journal of Income Distribution (forthcoming)*.
12. Stojkoski V, Utkovski Z, Basnarkov L, Kocarev L. 2019 Cooperation dynamics in networked geometric Brownian motion. *Physical Review E* **99**, 062312.
13. Kelly JL. 1956 A new interpretation of information rate. *The Bell System Technical Journal* **35**, 917–926.
14. Peters O. 2011 Optimal leverage from non-ergodicity. *Quantitative Finance* **11**, 1593–1602.
15. Andreoni J. 1990 Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal* **100**, 464–477.
16. Foster KR, Wenseleers T, Ratnieks FLW. 2006 Kin selection is the key to altruism. *Trends in Ecology & Evolution* **21**, 57–60.
17. Perc M, Jordan JJ, Rand DG, Wang Z, Boccaletti S, Szolnoki A. 2017 Statistical physics of human cooperation. *Physics Reports* **687**, 1–51.
18. Axelrod R, Hamilton WD. 1981 The evolution of cooperation. *Science* **211**, 1390–1396.
19. Hulme O, Vanhoyweghen A, Connaughton C, Peters O, Steinkamp S, Adamou A, Baumann D, Ginis V, Verbruggen B, Price J et al.. 2023 Reply to" The Limitations of Growth-Optimal Approaches to Decision Making Under Uncertainty". *Econ Journal Watch* **20**, 335–348.
20. Tversky A, Kahneman D. 1988 Rational choice and the framing of decisions. In *Decision Making: Descriptive, Normative, and Prescriptive Interactions* , pp. 167–192. Springer.
21. Vanhoyweghen A, Verbeken B, Macharis C, Ginis V. 2022 The influence of ergodicity on risk affinity of timed and non-timed respondents. *Scientific Reports* **12**, 3744.
22. Vanhoyweghen A, Ginis V. 2023 Human decision-making in a non-ergodic additive environment. *Proceedings of the Royal Society A* **479**, 20230544.
23. Tversky A, Kahneman D. 1974 Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131.
24. Furnham A, Boo HC. 2011 A literature review of the anchoring effect. *The Journal of Socio-Economics* **40**, 35–42.
25. Cone J, Rand DG. 2014 Time pressure increases cooperation in competitively framed social dilemmas. *PLoS ONE* **9**, e115756.
26. Capraro V, Cococcioni G. 2015 Social setting, intuition and experience in laboratory experiments interact to shape cooperative decision-making. *Proceedings of the Royal Society B: Biological Sciences* **282**, 20150237.
27. Rand DG, Greene JD, Nowak MA. 2012 Spontaneous giving and calculated greed. *Nature* **489**, 427–430.
28. Cornelissen G, Dewitte S, Warlop L. 2011 Are social value orientations expressed automatically? Decision making in the dictator game. *Personality and Social Psychology Bulletin* **37**, 1080–1090.
29. Lotz S. 2015 Spontaneous giving under structural inequality: Intuition promotes cooperation in asymmetric social dilemmas. *PLoS ONE* **10**, e0131562.
30. Meder D, Rabe F, Morville T, Madsen KH, Koudahl MT, Dolan RJ, Siebner HR, Hulme OJ. 2021 Ergodicity-breaking reveals time optimal decision making in humans. *PLoS Computational Biology* **17**, e1009217.
31. Skjold B, Steinkamp S, Hulme OJ, Peters O, Connaughton C. 2023 Are risk preferences optimal?. *OSF Preprints*.
32. Perc M, Gómez-Gardenes J, Szolnoki A, Floría LM, Moreno Y. 2013 Evolutionary dynamics of group interactions on structured populations: a review. *Journal of the Royal Society Interface* **10**.
33. Gollier C, Pratt JW. 1996 Risk vulnerability and the tempering effect of background risk. *Econometrica: Journal of the Econometric Society* pp. 1109–1123.

34. Stantcheva S. 2021 Perceptions and preferences for redistribution. Technical report National Bureau of Economic Research.

6. Appendix

(a) Wealth Evolution Rerun

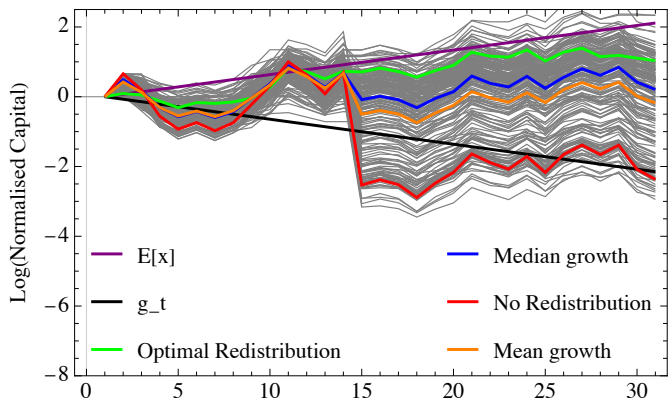

Figure 7: Capital evolution in the second iteration of the experiment. Coin tosses and bets were kept homogeneous with those used in the first iteration of the experiment; see figure 5. The only differences between both experiment iterations were the respondents and the fixed or variable nature of the starting redistribution rate.

(b) Impact of Tournament Style Rewards

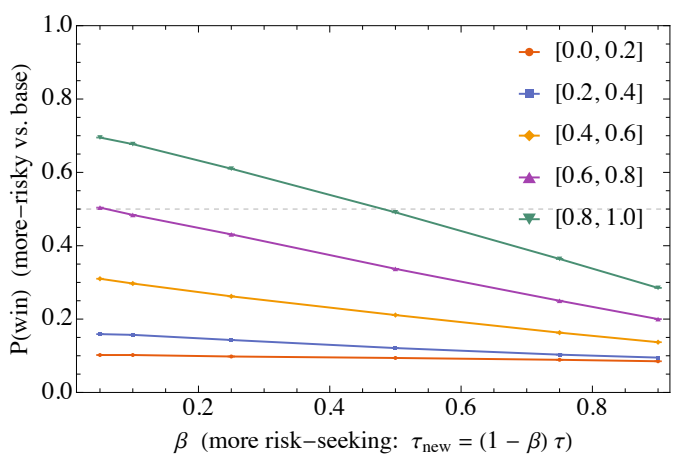

Figure 8: Each point shows the mean fraction of games (across agents) in which a more-risky clone of a baseline agent wins (higher final wealth). Baseline agents draw at random one τ per bet from five bins: $[0, 0.2], [0.2, 0.4], [0.4, 0.6], [0.6, 0.8], [0.8, 1.0]$. The more risky clone uses $\tau' = (1 - \beta)\tau$ (clipped to [0,1]) $n_{agents} = 500$, $n_{games} = 1000$. We observe that the tournament-style reward system, in which only the 5 best-performing respondents receive bonus payments, primarily benefits relative risk-seeking when the other respondents are highly risk-averse.

(c) Bets

Bet(i)	Outcomes	Bet(i)	Outcomes	Bet(i)	Bet (Right)	Bet (Wrong)
1	1.93 0.22	16	1.16 0.87	31*	200 -180	200 -200
2	0.63 1.4	17	1.44 0.69	32*	190 -202	-202 150
3	1.79 0.47	18	0.64 1.51	33*	-200 210	120 -200
4	1.47 0.69	19	0.72 1.37	34*	192 -200	192 -230
5	1.21 0.84	20	1.69 0.64	35*	-204 250	-204 230
6	0.79 1.22	21	1.28 0.78			
7	1.26 0.76	22	0.83 1.18			
8	1.69 0.44	23	1.44 0.70			
9	0.59 1.71	24	0.63 1.82			
10	0.6 1.98	25	1.68 0.57			
11	1.38 0.69	26	1.31 0.75			
12	1.67 0.58	27	0.76 1.42			
13	1.84 0.44	28	1.31 0.71			
14	0.04 2.00	29	0.5 1.59			
15	1.16 0.87	30	1.59 0.75			

Table 1: All 30 bets shown to respondents as well as the 5 attention checks, 31-35. Here, the respondent was required to select the most optimal bet. The outcome for all bets depended on a fair coin toss.