Creative Industries
Policy and Evidence Centre

Led by Newcastle University Business School with RSA

Discussion paper

A Silver Lining?

Consumer Preferences Over Human-Made and AI-Generated Production in Two Domains

Tom Lane, Harry Pickard, Matthew J. Walker & Timo Heinrich

DOI:10.5281/zenodo.22111802

DP Number: 2026/1

September 2026

Creative PEC is funded by: UKRI Arts and Humanities Research Council

About the Authors

Tom Lane, Newcastle University Business School, tom.lane@newcastle.ac.uk
Harry Pickard, Newcastle University Business School, harry.pickard@newcastle.ac.uk
Matthew J. Walker, Newcastle University Business School, matt.walker@newcastle.ac.uk
Timo Heinrich, Hamburg University of Technology, timo.heinrich@tunh.de

About the Creative Industries Policy and Evidence Centre

The Creative Industries Policy and Evidence Centre (Creative PEC) supports growth of the UK's creative industries through providing independent and authoritative evidence and policy advice.

Led by Newcastle University, with the Royal Society of Arts (RSA) and funded by the UKRI via Arts and Humanities Research Council, Creative PEC comprises a core consortium of Newcastle University, Work Advance, the University of Sussex and the University of Sheffield.

Disclaimer

This discussion paper is published by the Creative Industries Policy and Evidence Centre. The Creative PEC is funded by the UK Arts and Humanities Research Council (AHRC) and is hosted by Newcastle University in partnership with the Royal Society of Arts.

The views and interpretations expressed in this paper lie solely with the authors and may not be shared by the Director of the PEC, the editor of the Discussion Paper series, Newcastle University, the RSA and the AHRC. Readers who wish to challenge the evidence and/or interpretations provided are encouraged to do so by directly contacting the lead author and/or by writing to the Series Editor Prof Giorgio Fazio at giorgio.fazio@pec.ac.uk

Citation

If the information in this report is used in any subsequent research and/or publications, please cite as follows:

Lane, T. et al. (2026) A Silver Lining? Consumer Preferences over Human-Made and AI-Generated Production in Two Domains. Creative PEC Discussion Paper Series. Available at: https://doi.org/10.5281/zenodo.22111802.

A Silver Lining? Consumer Preferences over Human-Made and AI-Generated Production in Two Domains

Tom Lane* Harry Pickard* Matthew J. Walker* Timo Heinrich†

July 17, 2026

Abstract

Artificial intelligence can now generate outputs that were once the preserve of human labour. Yet well-documented algorithm aversion raises the possibility of consumer bias in favour of human-made work. This paper uses incentivised experimental auctions and state-of-the-art generative AI models to compare consumer valuations for human-made versus AI-generated images and poetry, sold respectively as physical prints and on mugs. In a broad population sample from the UK, we observe a premium for human-made work in consumer valuations. However, among the youngest participants we observe no such premium, mirroring the findings of a laboratory study with student participants. Demand for AI-generated outputs varies with individuals' attitudes toward AI. Whether the age pattern reflects ageing or cohort replacement will determine whether consumer preferences can help sustain human labour.

JEL classification: C91; D44; J24; O33; Z11
Keywords: Artificial intelligence; Algorithm aversion; Auction; Experiment; Labour market displacement; Willingness to pay

---

*Newcastle University Business School, Newcastle upon Tyne, United Kingdom, NE1 4SE (Tom.Lane@newcastle.ac.uk, corresponding author; Harry.Pickard@newcastle.ac.uk; Matt.Walker@newcastle.ac.uk).

†Hamburg University of Technology, 21073 Hamburg, Germany (timo.heinrich@tuhh.de).

The lab experiment reported in this paper was pre-registered at AEARCTR-0014469 on September 24, 2024. The online experiments were pre-registered at AEARCTR-0018880 on June 07, 2026. Ethics approval was obtained from the Newcastle University Faculty of Humanities and Social Sciences Ethics Committee (#50559/2023 and #50090/2024). Funding from Hamburg University of Technology and Newcastle University is gratefully acknowledged. We thank the editor Steffen Huck, four anonymous reviewers and Lijia Wei for helpful feedback, along with audiences in presentations at Durham University, Newcastle University, the University of Amsterdam, the 2025 CeDEx-China Workshop in Experimental and Behavioural Economics, the 2026 Wuhan Cherry Blossom Workshop on Experimental Economics, and the 2026 Royal Economic Society Conference. We are grateful to colleagues at Newcastle University's Fine Art Department (especially Nigel Villalard and Theresa Easton), the Newcastle University Library Special Collections Team and Katja Riemer at Kunsthalle Bremen for their assistance in sourcing artworks.

1 Introduction

Artificial intelligence is a defining narrative of our times. After decades of gradual development, it rose precipitously into the public consciousness following the release of *ChatGPT* in late 2022. The current advance of AI is dizzying, leaving few areas of economic life untouched and sparking both excitement and anxiety. It has been estimated that current technologies alone are theoretically able to automate activities taking up 57% of US labour hours (Yee et al., 2025). Generative AI marks a departure from prior waves of automation, which largely impacted routine or physical tasks while leaving creative and cognitive work comparatively insulated (Autor, 2015); it is today able to perform creative, tacit-knowledge and design work once thought resistant to automation (Noy and Zhang, 2023). Free-to-use AI tools can now write prose and poetry that raters judge superior to human output (Herbold et al., 2023; Porter and Machery, 2024) and achieve levels of persuasion similar to humans (Goldstein et al., 2024). AI can even generate satirical content of comparable appeal to that of professional comics (Gorenz and Schwarz, 2024).

As with historical mechanisation, economic theory would predict that AI will create some jobs and destroy others (Acemoglu and Restrepo, 2018). While there are positive productivity effects from the new technology that should contribute to labour demand, these may be insufficient to counter the adverse forces on employment in sectors where AI can strongly substitute for human labour (Acemoglu and Restrepo, 2019). In the instance that AI does not merely augment a worker's efficiency in completing specific tasks but directly manufactures their finished product at equivalent quality and a fraction of the cost, the worker's prospects would appear bleak. The social consequences of worker skills honed over many years being treated by markets as largely obsolete could be substantial.

There is, however, a potential saving grace for human labour: buyers might attach higher value to its output *simply because it is human-made*. That is, goods known to be

produced by human labour might attract higher willingness-to-pay (WTP) than goods of otherwise comparable appeal that are known to be generated by AI. This hypothesis is in line with prior evidence that physical products are valued more highly when made by hand rather than by machine (Fuchs et al., 2015). It is also consistent with the well-documented phenomenon of *algorithm aversion*, wherein information and advice provided by humans is preferred over that from algorithms (e.g., Dietvorst et al., 2015; Castelo et al., 2019; Burton et al., 2020; Jussupow et al., 2020; Mahmud et al., 2022; Dargnies et al., 2026).

This paper tests for such a silver lining. Using incentivised Vickrey auctions, we measure consumers' WTP for visual artworks and poetry-inscribed mugs. We compare bids for items known to feature human-made images or poetry with bids for items known to feature AI-generated content of similar baseline appeal (i.e., valued similarly by consumers unaware of each item's origin). This approach identifies whether consumers attach intrinsic value to human origin under controlled conditions. We conducted three experiments, spanning two populations and two content domains. First, we ran laboratory auctions for artworks with a UK student sample. Next, we replicated these auctions online with a broad UK population sample, and in parallel ran auctions for poetry mugs on a further sample drawn from the same general population. Relative to the artworks, the poetry mugs extend our tests to the domain of literary content, thus broadening the scope of our research.

We selected these items because they represent useful test cases for the identification of consumer preferences for work to be of human origin. It would be difficult to identify such preferences cleanly by studying an informational product, such as those addressed in the algorithm aversion literature. The quality of such products is generally not perfectly observable to consumers, who may base their valuations on perhaps biased perceptions of the perceived trustworthiness or reliability of a human producer relative to AI. Such considerations are not of direct relevance to value judgments over the products we study,

nor are our items valued primarily on usefulness or accuracy, as utilitarian outputs such as translations or computer code would be. This allows cleaner measurement of raw preference over comparable human-made and AI-generated output. This is true for both item types in our experiments, since we vary only the origin of the content featured on each item (the image or the poetry) and not of the physical product itself.

The products we study (visual artworks and poetry mugs) are not works of art, but anonymous, reproducible copies. Such products have traditionally been heavily reliant on human input but AI can now generate substitutes at low marginal cost.[1] Absent intrinsic preferences for human production, human workers whose outputs are used in these products may be at risk of displacement. Indeed, in this study's research process, one of the authors – lacking formal training and, according to colleagues, devoid of discernable artistic or literary taste – was able to generate a portfolio of hundreds of images in the space of one afternoon with the assistance of the AI tool *Midjourney*, and in the space of another a collection of hundreds of haikus with help from the Large Language Model (LLM) *Claude*. Observers, to whom the source of these works was not disclosed, deemed many of them indistinguishable in origin and appeal to the works of professional artists or poets.[2]

---

[1]It is important to note that human-made images are usually technically mediated somewhere in the production process. AI models add a new order of magnitude to the technical mediation based on recognising and applying statistical patterns. We thank an anonymous referee for their helpful comments on this point.

[2]The commercialisation of AI-generated art, and art-adjacent products, has accelerated rapidly. In November 2024, a portrait of Alan Turing generated by an AI robot was auctioned by Sotheby's for $1 million (Pope, 2024), just months after the announcement of the world's first AI art museum in Los Angeles (Beckett, 2024). Institutional acceptance of AI-generated art has continued to expand, with Christie's holding the first major auction dedicated to AI art in early 2025, although the sale prompted thousands of creatives to protest that many of the underlying models had been trained on copyrighted works without permission. At the same time, legal disputes over AI-generated imagery have intensified, with landmark copyright actions brought against AI image-generation firms (Hill, 2025). Meanwhile, Merriam-Webster named *slop* - defined as "digital content of low quality that is produced usually in quantity by means of artificial intelligence" -

Our results provide some hope for those at risk of displacement by AI. In our experiments using broad population samples, we find a significant WTP premium on human output for both product types. Notably, this effect is driven by older consumers; no such premium is found among younger participants. Our student lab sample also demonstrate indifference between human-made and AI-generated artworks.

We find that WTP for AI products is strongly sensitive to buyers' attitudes towards AI. It is plausible that such heterogeneity in valuation could originate from two sources. First, an individual's perception of AI as a threat to their occupation: Eloundou et al. (2024) demonstrate that occupations vary in their exposure to LLM capabilities. Second, differences in adoption rates across demographic and labour market characteristics: Bick et al. (2024) document substantial variation in AI usage patterns that closely mirror adoption of the personal computer. Those in our experiments who are relatively hostile to AI tend to value AI work less. Workers' displacement may therefore depend on how such sentiment evolves, particularly among younger cohorts who hold more favourable attitudes towards AI and who, according to our data, demonstrate no intrinsic preference for human-made output.

Our research contributes to the nascent – yet rapidly growing – economics literature on the impact of AI for the future of work (Acemoglu and Restrepo, 2019; Ernst et al., 2019; Lu and Zhou, 2021; Acemoglu et al., 2022; Zhu and Zou, 2023; Charness and Grieco, 2024; Locci and Massoni, 2024; Acemoglu, 2025). Various recent studies have investigated the impacts of workers adopting AI tools, demonstrating mixed findings (e.g., Babina et al., 2024; Brynjolfsson et al., 2025; Kanazawa et al., 2026). Whereas such research focuses on the supply side of markets for AI-produced outputs, we contribute to a complementary strand of the literature which addresses the demand side by studying consumers' reactions to AI adoption. Up to now, no published study has provided, as we do, controlled

as its word of 2025 (https://www.merriam-webster.com/wordplay/word-of-the-year).

measurements of demand for human- versus AI-produced goods using methods widely accepted in economic research.

Earlier published research in this area has relied on hypothetical measurement techniques and generally used deceptive practices to control the baseline appeal of products presented to subjects as AI- or human-made. To date, such studies have produced inconsistent results, with some finding significant preferences in favour of human output (e.g., Granulo et al., 2021; Horton et al., 2023; Millet et al., 2023; van Waning, 2024) but others obtaining significant results in the opposite direction (Zhang et al., 2022). There is evidence that preferences for human output are strengthened when observers evaluate it in close comparison with AI output, rather than evaluating each in isolation (Horton et al., 2023); indeed, Neef et al. (2024) and Jacobs et al. (2026) find the premium on human output appears only when using a within-subject rather than between-subject elicitation, a design feature also found to exert a moderating effect in a recent meta-study (Qin et al., 2025). Another meta-study points towards a lower premium on human production in studies with greater ecological validity (Zehnle et al., 2025).

We build on the prior published literature by eliciting incentivised rather than hypothetical decisions, with no deception, to measure WTP for genuinely AI-generated and human-made artworks, the baseline appeal of which is controlled through a rigorous pretesting selection procedure.[3] Importantly, when auctioning human-made artworks online, we took steps to avoid using images in the public domain that subjects could freely search for and access outside the experimental platform, thereby sidestepping a potential confound that may influence treatment effects in related designs.[4] Indeed, one contribution

[3]One published experiment which includes an incentivised (but deceptive) valuation of human versus AI production is by Feng and Song (2026), who find a significant premium on human output for artworks.

[4]We are aware of three relevant unpublished papers, all reporting online experiments using incentivised methods. Abel and Johnson (2025) find no significant effect of AI production on WTP, but only investigate one product, a short story labelled as either human-written or AI-generated. The design relies on deliber-

of our paper is the provision of a novel methodological approach, which could easily be replicated by economists working on related projects in the future who wish to obtain high levels of experimental control while remaining within the no-deception convention of the discipline. Moreover, our tests are conducted over a range of artworks and poetry, thereby mitigating a concern that the results may be sensitive to the peculiarities of any single item under examination.

2 Methods

Our objective is to cleanly identify the *ceteris paribus* effect of AI production on the demand for goods that were once the preserve of human labour.[5] To achieve this goal, we designed three studies to elicit WTP for items which are either made by humans or generated by AI, but are otherwise comparable in their baseline appeal: (i) an initial laboratory experiment with student subjects where we auctioned artworks; (ii) a follow-up online experiment with a general population sample where we auctioned artworks; and (iii) another follow-up online experiment with a general population sample where we auctioned

---

ately misdirecting subjects in one treatment to believe that the AI-generated product is human-made (this is implicitly signalled by displaying an image of a person on a typewriter). Mandel and Imas (2026) elicit WTP for exclusively human-made artworks – labelled as either human or AI – finding a significant premium on human-labelled outputs, which weakens when the artworks are presented as mass-produced. Friedrichsen et al. (2026), who study WTP for a small number of musical pieces, avoid deception but do not control baseline appeal between human and AI songs – instead they provide controlled measurements of the effect of disclosing to participants that a given song is human-made or AI-generated, with disclosure of the AI source shown to lower WTP. We regard our study as complementary to these three, while also offering certain advantages. We investigate a wider product range, with a variety of sample pools and settings, including laboratory measurements.

[5]The *ceteris paribus* qualification is important as the overall value for outputs generated by AI may be larger or smaller than outputs made by humans if (1) the two outputs differ in baseline appeal, or (2) if the provenance of AI-generated outputs are clearly discernible and consumers place a higher or lower subjective value on one of the two production modes. Our tests exclude by design the first channel and focus purely on the second.

mugs with haiku-style poetry printed on them (henceforth referred to as poetry mugs). Survey evidence suggests that the adoption of generative AI among UK university students was already normalised by 2024 (the year we conducted the lab experiment) and is now nearly universal (Stephenson and Armstrong, 2026). Hence, to increase the generalisability of our results, the follow-up experiments were designed to elicit valuations across a wide range of demographic strata, as well as different item types. All three studies were pre-registered on the AEA RCT Registry.[6]

In each study, to obtain the final samples of artworks or poetry for use in the experiments, we followed a rigorous pre-testing procedure to ensure that the source (human or AI) of each work is undetectable (unless disclosed) and that, in absence of knowledge about the source of the outputs, their range of baseline appeal is similar between the two production modes. By comparing WTP not merely for a pair of single goods but for two pairs of product sets (10 human-made and 10 AI-generated artworks or poetry mugs), we enhance the informativeness and generalisability of our study. Studying a range of items mitigates the concern that results may be driven by the idiosyncrasies of any single one. It also allows us to separately explore the effects of AI production for relatively lower or higher value outputs.

2.1 Selection of Test Items

2.1.1 Artworks for the lab experiment

We arrived at our two sets of visual artworks for the lab experiment through a multi-step process. First, we selected initial samples of 100 human-made and 100 AI-generated artworks. The human-made images were randomly selected from the open access art aggregator Rawpixel, conditional upon (1) the images having a CC0 license (public domain)

[6]For details, see Lane et al. (2024) and Lane et al. (2026).

which permits unrestricted use of the images, including resale; (2) the images not containing artists' signatures; and (3) the images not being immediately recognisable to the researchers as famous works.

The AI images were generated using Midjourney, a premium AI art platform which turns simple text prompts into images. Our images were generated using one of two approaches. The first employed Anthropic's LLM, Claude, to generate an image description, which we then fed to Midjourney as a prompt. The second "re-imagined" human-made images by randomly selecting a human-made image to upload to Midjourney, and asking the software to describe the subject of this image in text. The text was then fed back into Midjourney as a prompt, along with the "Style reference" option.[7]

Next, we ran an online survey to identify the sub-sample of artworks for which it is undetectable (on average) whether the image was AI-generated or human-made. Each participant was randomly presented (in random order) 50 of the 200 images. They were informed only that all images were drawn from a database where half the images were human-made and half were AI-generated. For each image, participants were asked to evaluate the chances, in percentage terms, that the image was made by a human artist and that it was generated by AI. The two percentages were required to sum to 100. The accuracy of these evaluations were not incentivised, because the human-made images were sourced from the public domain.[8]

[7]Our results remain qualitatively unchanged when excluding the two "re-imagined" images eventually used in the lab experiment. More details and examples of the image generation process are given in Appendix A.

[8]The merits of incentivising belief elicitation in experimental economics remain subject to debate, since in principle it could either increase or reduce accuracy (Schotter and Trevino, 2014; Trautmann and van de Kuilen, 2015; Danz et al., 2022). In our case, we did not want to induce information seeking and there is no obvious reason for subjects to deliberately misrepresent their beliefs in the absence of incentives. In the pre-test for our second online experiment we find no difference in the mean-error rate of predictions between incentivised and unincentivised belief elicitation (see Section 2.1.2). Hence, we believe that not incentivising beliefs is unlikely to have been a source of bias in the pre-testing of any of the items.

We recruited a sample of 528 respondents to complete the survey through Prolific.[9] To maintain consistency in demographic characteristics with our eventual lab experiment sample, we limited participation to current students located in the UK. We retained in our sub-sample all images for which we failed to reject the null hypothesis, at the 5% significance level, that the perceived chance of the image being human-made was 50%, based on a two-tailed t-test.[10] In total, 48 images were retained (19 human-made and 29 AI-generated).

This sub-sample was then taken forward to a second Prolific survey, run on a different set of participants but drawn from the same target population. Its purpose was to elicit the baseline appeal for each image in absence of knowledge about the source. We recruited a sample of 132 respondents to complete the second survey. Each respondent was presented with all 48 images, in random order, and in each case was required to report their hypothetical WTP for an A5 print of the image. No information was provided regarding the origin of each image. Participants were presented with a multiple price list of amounts in increments of £1 ranging from £0 to £15, and were asked to select the highest amount they would be willing to pay.

The responses from the second survey were used to select the final samples of 10 human-made and 10 AI-generated artworks. To ensure the two sets of images had comparable but wide-ranging distributions of baseline appeal, we implemented the following procedure. For the human-made artworks, we selected the images with the highest and

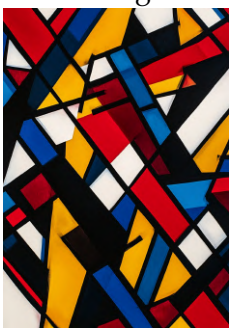
[9] In all of our pre-test surveys for the lab experiment, we dropped from analysis the 5% of respondents with the fastest response times, on the grounds that such respondents may not have been paying sufficient attention to generate reliable data. To further improve the data quality, we restricted eligibility to respondents with a Prolific approval rate above 90% and we included a Captcha to exclude bots and an attention check at a random point in the survey (we excluded observations from any respondent who failed the attention check).

[10] Based on our number of usable observation, this test has 80% power to detect at the 5% significance threshold a difference of 6 percentage points from a mean perceived chance of 50%, with Cohen's d = 0.25.

lowest average (hypothetical) WTP; calculated eight evenly spaced points on the monetary interval between them; and selected as the remaining images those with the average WTP closest to each point. Then, we chose the 10 AI-generated images closest in average WTP to each selected human-made image. The final sets of images for the lab experiment are presented in Figure 1; it is clear that both the human and AI images span a range of artistic styles (abstract and representational) and visual characteristics.

Figure 1: Artworks in Lab Experiment

AI Image 1



AI Image 2



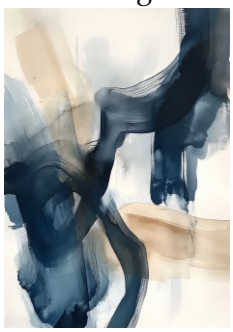
AI Image 3



AI Image 4



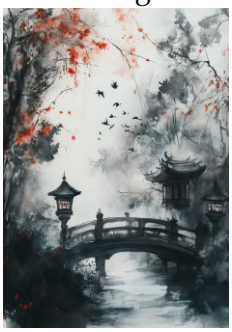
AI Image 5



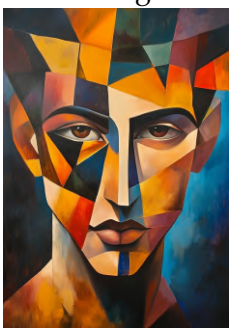
AI Image 6



AI Image 7



AI Image 8



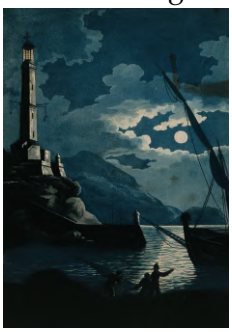
AI Image 9



AI Image 10



Human Image 1



Human Image 2



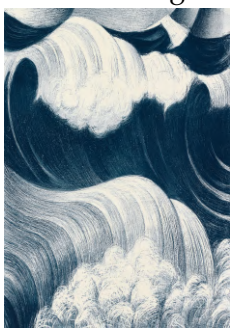
Human Image 3



Human Image 4



Human Image 5



Human Image 6



Human Image 7



Human Image 8



Human Image 9



Human Image 10



Finally, we ran a third Prolific survey with 132 new participants, in which every respondent rated each of the images in the final sample on an aesthetic rating scale from 1 to 5, where 1 is "very unappealing" and 5 is "very appealing".[11] As in the second survey, participants were not informed about the origins of the images and the order of images was randomised.

2.1.2 Artworks for the online experiment

The process for selecting the items to be auctioned in the online experiments closely follows the pre-testing procedures described for the lab experiment. One necessary difference was the process for sourcing the human-made images. As the auctions were to be conducted online, we required that the images were not in the public domain so as to avoid a potential interaction with the treatment effect. If images are circulating publicly, mass-produced physical copies could be available cheaply, so bids for such items are capped by the market price rather than reflecting WTP. This would affect only the human-made images and not newly AI-generated images, biasing the WTP for human-made artworks downwards in relative terms.[12] This concern was not present for the lab experiment, where we could ensure that participants were unable to navigate out of the experiment software or access personal electronic devices.

To ensure that the new artworks spanned a range of artistic styles as in the lab experiment, we collected images from different offline sources. First, we collected an initial sample of images made by art students. Students were invited to submit original images

[11] A similar aesthetic rating scale is used by Kirk et al. (2009). In all of our pre-test surveys for the lab experiment, respondents received a fixed participation fee (£2.25 in the first two surveys, and £1.05 in the third survey. These fees were 50% higher than originally planned, because a soft launch of the first survey revealed longer than expected completion times).

[12] We generated a new set of AI images for the online experiment and so these images were also not in the public domain (full details of the image generation process can be found in Appendix A).

in A5-sized portrait format using a range of artistic styles (abstract and representational) and visual characteristics. We collected these images through a call for submissions and required, as conditions of inclusion, that (1) the images were made by the submitting student, who owned the rights to the work and agreed that the research team may use, share, reproduce, and resell the work for research purposes; (2) the images did not contain signatures or any other text; (3) the images could not be found online using TinEye reverse image search;[13] and (4) the images were not immediately recognisable to the researchers as famous works or close derivatives thereof.

Second, we sourced additional images from Newcastle University's library archives. Instead of condition (1), we required images to hold no copyright; we maintained conditions (2), (3) and (4). Finally, we sourced additional images from the collection of Kunsthalle Bremen. Again, we required images to hold no copyright and maintained conditions (2), (3) and (4).

This sourcing process left us with an initial sample of 100 human-made and 100 AI-generated images. The steps for obtaining the final sample of 10 human-made and 10 AI-generated artworks to be auctioned then followed the same rigorous pre-testing procedure as described in Section 2.1. To maintain consistency in demographic characteristics with our eventual online experiment sample, we used Prolific's regional representative sample feature for the adult population in the United Kingdom, which samples participants based on age, gender and regional quotas derived from UK Census data.[14]

---

[13] tineye.com.

[14] We also introduced additional data quality checks to account for recent developments in bot and AI agent capabilities. This included employing a video attention check designed to screen out AI agents (Celebi et al., 2026) and Prolific's recently introduced features for bot and LLM screening (prolific.com/resources/introducing-authenticity-checks-beta-ensure-genuine-human-responses-in-the-age-of-ai), in addition to the classic multiple-choice attention check, Captcha and exclusions based on response time implemented previously. For full details on the protocol and power calculations used to determine the sample sizes for the online experiment pre-tests, see our pre-registration.

Because none of the images we used were in the public domain, we could compare incentivised and unincentivised belief elicitation in our source detection survey. In a first wave of data collection, we randomly assigned participants to one of two treatment arms. In the first arm, source evaluations were elicited without incentives (as in the pre-test for the lab experiment). In the second arm, source evaluations were elicited under incentives using a binarized quadratic scoring rule (e.g., Healy and Leo, 2025). Specifically, participants provided their belief p that an image was AI-generated; they then received a lottery that paid a bonus of £0.75 with probability $1-(1-p)^2$ if the image was AI-generated, or with probability $1-(0-p)^2$ if it was human-made. We recruited 535 participants across the two arms. This enabled us to assess whether incentivisation leads to systematically different source evaluations.

We found no statistically significant difference in the mean error-rate of predictions that the image was human-made between the two treatments.[15] Hence, we proceeded to recruit another 510 participants for the source detection survey without incentives (for a total sample size of 765). In total, we failed to reject the null hypothesis that the perceived chance of the image being human-made was 50% at our pre-registered significance threshold of 5% for 33 images (11 human-made and 22 AI-generated).

This sub-sample was again taken forward to a second Prolific survey, run on a different set of participants, to elicit the hypothetical WTP for each image in absence of knowledge about the source. The final sample of 10 human-made and 10 AI-generated artworks was then selected using the same method as for the lab experiment to ensure the two sets of images had comparable but wide-ranging distributions of baseline appeal (see Appendix Figure A.3). Unlike for the lab experiment, the average hypothetical WTP in the final sample of human-made artworks for the online experiment was somewhat lower than for the AI-generated artworks (difference in means of £0.58, p<0.01 based on a paired t-

[15]For brevity, we do not report the test results here. These are available on request.

test). This divergence occurred because the matching process was constrained by only 11 human-made artworks passing the first pre-test as indistinguishable in origin, and by the absence of AI-generated artworks with similarly low hypothetical WTP to match the lowest-appeal human artworks in the final selection. However, we still ended up with sets of human-made and AI-generated artworks that both spanned a wide range of baseline appeal. For each image in the final sample, we elicited aesthetic ratings in a third survey.[16]

2.1.3 Poetry for the online experiment

We conducted a parallel stream of pre-tests, run on different subjects, for the poetry mugs. The first step was to identify source poetry written by humans and source poetry generated by AI. For this purpose, we collated an initial sample of 100 haiku-style compositions based on extracts from human-written poetry at Project Gutenberg.[17] For ease of reference, we refer to these compositions as *haikus* below. This terminology is used as a convenient shorthand for readability: we recognise that the pieces do not necessarily conform to haiku as a historically and culturally specific form. In the experiment itself, we deliberately referred to them with the more general term *poetry* rather than haikus, so as not to imply a stricter classification than the compositions warranted. We required that the poetry on which the haiku compositions are based is in the public domain,[18] and that the resulting compositions were not immediately recognisable to the researchers as unusually famous works. We used Claude to generate the AI haikus (full details of the poetry generation process can be found in Appendix A).

As with the artworks, we conducted the first pre-test survey to identify a candidate

[16]The number of respondents recruited for the second and third pre-test surveys in the online artwork experiment were 111 and 170, respectively.

[17]projekt-gutenberg.org.

[18]The potential confound that was present for the artworks is not applicable to the poetry mugs because while the human-written poetry is in the public domain, the product itself is bespoke to the experiment.

sample of haikus whose source is undetectable (on average). Participants were shown randomly selected unidentified haikus from the initial sample and evaluated the probability that each was human-written or AI-generated. Unlike in the artwork stream, source evaluations were elicited without incentives only, as the human-written poetry was sourced from the public domain. We recruited 778 participants for the poetry source detection survey using Prolific's representative sample feature. In total, we failed to reject the null hypothesis that the perceived chance of the haiku being human-written was 50% at our pre-registered significance threshold of 5% for 75 haikus (33 human-written and 42 AI-generated). These haikus were taken forward to the second pre-test survey in which we elicited hypothetical WTP for mugs on which they were printed. The final sample of haikus (10 human-written and 10 AI-generated) was selected based on the resulting distributions using the same method as for the artworks (see Figure A.4; the average hypothetical WTP of the two sets of poetry mugs is almost identical). Finally, we elicited aesthetic ratings in a third survey.[19]

2.2 Measuring WTP

2.2.1 Auctions in the lab experiment

To obtain incentivised estimates of the WTP for each artwork in the lab, we used Vickrey auctions (Vickrey, 1961). In each session, subjects participated in a series of 10 independent second-price auctions, either for the human-made or for the AI-generated items. The auctions were conducted in groups of four bidders, which remained fixed across all rounds.

Each subject was endowed with a budget of £12 which they could use for bids, with

[19]The number of respondents recruited for the second and third pre-test surveys in the online poetry experiment were 263 and 160, respectively.

any unspent money retained as take-home income from the experiment. In each round, an A5 print of the respective image was auctioned; subjects were shown the physical print for sale and could also view the artwork on their computer screens. They entered sealed bids, with any amount allowed, at intervals of £0.01, between £0 and £12. The order of lots was randomised in each session. After all rounds were complete, one round was selected at random separately for each group. The subject who submitted the highest bid in the group in the selected round's auction won the artwork and paid an amount equal to the second-highest bid in the group. At the end of a session, all payments and distribution of the physical artwork print for a winning bidder to take away were made individually and in private at the back of the laboratory.

To maintain the independence of observations between subjects, no feedback was provided on other participants' bids, or the outcome of each auction round, until the end of the experiment. Providing no feedback between rounds is particularly important when values are likely to be affiliated across subjects (Harrison et al., 2004).[20]

In a Vickrey auction, each participant should theoretically submit a bid equal to their maximum WTP for the item. At the beginning of the experiment, subjects received instructions including the auction rules and an explanation as to why it would be optimal for them to bid their WTP in each round.[21] Providing such information is appropriate in this type of experiment where our objective is to elicit "homegrown values", as helping subjects understand the dominant strategy may reduce measurement error (Harrison et al., 2004). To further facilitate learning, before the main auction rounds commenced, each subject individually undertook a hypothetical training exercise in an induced value setting where they received feedback if they failed to follow the dominant strategy.

[20]For this reason, we decided against employing an open English auction. We also enforced a strict "no personal electronic devices" protocol in all sessions.

[21]We used the task instructions for the Vickrey auction due to Rutström (1998, p.432).

2.2.2 Auctions in the online experiments

To maximise comparability of the design with the lab experiment, we used the same second-price auction protocol for the two online experiments as described in the previous sub-section, with fixed groups of four bidders participating across 10 auction rounds for either visual artworks or poetry mugs. Again, each subject was endowed with a budget of £12 which they could use for bids, with any unspent money retained as take-home income from the experiment. Due to the online implementation, group matching and the realization of auction outcomes for the randomly selected payoff-relevant round was conducted post hoc (within 48 hours of a subject completing the experiment) rather than in real time.

2.3 Treatments and Hypotheses

2.3.1 Lab experiment

In the lab experiment, each subject was randomly assigned to one of four treatments, which differed according to the production source of the items being auctioned and whether or not subjects were informed of this.

In the ***Human No Info*** treatment, we elicited WTP for the human-made artworks but without providing any information about the origin of each work. In the ***AI No Info*** treatment, we elicited WTP for the AI artworks without informing subjects of the works' origin. As our pre-testing suggested subjects would be unable to perceive differences between the human and AI artworks in their likelihood of being human-made, and in absence of information about this would perceive the two sets of images to be of comparable appeal, we expected to observe no difference in average WTP between these two treatments.

The same sets of human-made and AI-generated images were auctioned, respectively, in the ***Human Info*** and ***AI Info*** treatments, but in these treatments subjects were informed

of the artworks' origins. The *Human Info* instructions included the sentence: "All artworks being sold are created by human artists." The equivalent sentence in the *AI Info* treatment read: "All artworks being sold are created by an artificial intelligence (AI) program." To keep treatment variations minimal and avoid experimenter demand effects through our wording, we opted to use the term *created* consistently in both variations. To ensure this information remained salient, participants received an on-screen reminder about the artworks' origins in each round.

Assuming our strategy to select human-made and AI-generated artworks of comparable baseline appeal is successful, the difference in WTP between the *Human Info* and *AI Info* treatments can be interpreted as the *ceteris paribus* effect of AI production on demand for the artworks. Any residual differences between these treatments in baseline appeal, as measured in our pre-testing process, will be controlled in our regression analyses. Based on the evidence discussed in the introduction that people often prefer human-made output and human-provided advice over algorithmic alternatives, we hypothesised a negative effect of AI production on demand, and therefore higher WTP in the *Human Info* treatment than in the *AI Info* treatment, on average. Nevertheless, the average treatment effect may mask substantial heterogeneity across individuals based on their attitudes toward AI, a possibility that we consider in more detail below. We further expected that the knowledge an artwork is AI generated would reduce its valuation by consumers, whereas the knowledge that a work is human-made would increase this valuation. Therefore, we also hypothesised higher WTP in the *Human Info* treatment than in the *Human No Info* treatment, and lower WTP in the *AI Info* treatment than in the *AI No Info* treatment.

To summarise, we pre-registered the following hypotheses for the lab experiment:

Hypothesis 1 *Bids in* ***AI Info*** *will be lower than in* ***AI No Info****.*

Hypothesis 2 *Bids in* ***Human Info*** *will be higher than in* ***Human No Info****.*

Hypothesis 3 *Bids in **Human Info** will be higher than in **AI Info**.*

2.3.2 Online artwork experiment

Subjects in the online artwork experiment were randomly assigned to one of two treatments, which differed according to the production source of the works being auctioned. The online experiments were conducted after the lab experiment, in which we would find no significant difference in WTP between the ***Human No Info*** and ***AI No Info*** treatments (see Section 3.1). Since the sets of human-made and AI-generated artworks in the lab experiment were almost identical in baseline appeal as elicited in our pre-testing process, this result validates our assumption that this measure of baseline appeal closely predicts how uninformed subjects bid for items under incentivised experimental conditions, and that this would also hold true in the online experiments. As such, we opted to exclude the ***Human No Info*** and ***AI No Info*** treatments from our online experiments, in order to maximise statistical power for our comparison of primary interest, between ***Human Info*** and ***AI Info***, under the assumption that this comparison accurately reveals preferences for human-made over AI-generated items, so long as pre-test-measured baseline appeal is controlled.

In the ***Human Info*** treatment, the visual artworks being sold were based on human-made images, and subjects were informed of this. In the ***AI Info*** treatment, the visual artworks being sold were based on AI-generated images, and subjects were informed of this. As in the lab experiment, we opted to use the term *created* consistently in both variants of the instructions to keep treatment variations minimal.

We pre-registered the following hypothesis for the online artwork experiment:

Hypothesis 4 *For visual artworks, average bids are lower in the **AI Info** treatment than in the **Human Info** treatment.*

2.3.3 Online poetry mug experiment

Subjects in the online poetry mug experiment were again randomly assigned to one of two treatments, which differed according to the production source of the works being auctioned. In the ***Human Info*** treatment, the poetry mugs being sold were based on human-written poetry, and subjects were informed of this. In the ***AI Info*** treatment, the poetry mugs being sold were based on AI-generated poetry, and subjects were informed of this. For the same reasons as with the artworks, we did not run ***Human No Info*** and ***AI No Info*** treatments in this experiment.

We pre-registered the following hypothesis for the online poetry mug experiment:

Hypothesis 5 *For poetry mugs, average bids are lower in the **AI Info** treatment than in the **Human Info** treatment.*

2.4 Implementation and Sample Characteristics

2.4.1 Procedures

All pre-tests were conducted using Prolific. For the lab experiment, the pre-testing process took place between September and November 2024. For the online experiment, the pre-testing process took place in June 2026.

The lab experiment was conducted between December 2024 and January 2025 at the Experimental and Behavioural Economics Lab of Newcastle University (United Kingdom). Treatments were randomly assigned at the session level and sessions lasted around 45 minutes. The experimental interface was programmed using oTree (Chen et al., 2016). Subjects were university students recruited via ORSEE (Greiner, 2015).[22] We aimed for at least 50 subjects in each treatment.

[22]The experimental instructions for the lab experiment are available in Appendix B.1, and the screenshots of the auction program in Appendix C.1.

The final sample size for the lab experiment is $N = 224$ ($n = 56$ in *Human Info*, $n = 60$ in *Human No Info*, $n = 56$ in *AI Info* and $n = 52$ in *AI No Info*) across either three or four sessions per treatment.

The online artwork and poetry mug experiments were run in parallel and the combined sample was provided by Pureprofile.[23] The online experiments were implemented using Qualtrics survey software (Qualtrics, Provo, UT). Participants were recruited over an approximately 10-day period commencing at the end of June 2026, using a single link which randomly assigned them to one of the two studies and treatments. The sample for the online experiments was stratified to match the composition of the regional UK representative sample collected on Prolific for the pre-tests, using interlocked quotas for age and gender and independent quotas for region. The auction winner for the payoff-relevant round in each group was mailed a physical copy of their artwork or poetry mug after data collection concluded, with the second-highest bid for that group and round subtracted from their participation bonus of £12.[24]

We targeted a sample size of 252 participants per treatment in each of the online experiments. This target was based on a two-sample two-sided t-test with 80% power to detect a difference at the 5% significance level, assuming Cohen's d = 0.25 and a pooled standard deviation of 2.33 as observed across the *Human Info* and *AI Info* treatments in the lab experiment. Given that the 5% of respondents with the fastest response times would be dropped, we aimed to recruit 267 participants per treatment.

The final sample size for the online artwork experiment was $N = 525$ ($n = 262$ in *Human Info* and $n = 263$ in *AI Info*). The final sample size for the online poetry mug experiment was $N = 541$ ($n = 273$ in *Human Info* and $n = 268$ in *AI Info*).

[23] business.pureprofile.com/.

[24] For details of the procedures employed to acquire high-quality data online, along with a copy of the online experimental instructions and screenshots of the auction program, see Appendices B.2 and C.2.

At the end of each of the lab and online studies, subjects completed a post-experiment questionnaire to elicit standard socio-demographic data, information about art or poetry knowledge and awareness, familiarity with and knowledge of AI tools, and general attitude toward AI using a previously validated measure which we will discuss below.[25]

2.4.2 Sample characteristics

Student sample for the lab experiment. Summary statistics for the sample of student subjects are shown in Table E.1 and the full set of covariate balance tests by treatment in Table E.2. Treatments are generally balanced and a Wald test of the joint significance of all control variables is not significant at the 5% level.[26] Overall, 50% of our student sample are aged 18 to 21, 49% are male, 40% are UK nationals and 41% reported a parental income at age 16 above the country average.

Regional representative sample for the online experiments. Summary statistics for the general population samples are shown in Table F.1 and the full set of covariate balance tests by treatment in Table F.2. Treatments are balanced in each of the online experiment studies and Wald tests of the joint significance of all control variables between *Human Info* and *AI Info* treatments are not significant at the 5% level for either the artwork study or the poetry mug study.[27] Overall, across the two online studies, 11% of our broad population sample are aged 18 to 24, 19% are aged 25 to 34, 17% are aged 35 to 44, 17% are aged 45 to 54 and 36% are aged 55 or over; 50% are male, 94% are UK nationals, 55% either hold or are studying for a university degree, and 25% reported a family income above the country

[25]The post-experiment questionnaires can be found in Appendices D.1 and D.2, respectively.

[26]We note the following exceptions: the proportion of 18-21 year-olds is higher in the *Human Info* treatment than in the *Human No Info* treatment (p=0.038) or *AI Info* treatment (p=0.024); and self-reported art knowledge is higher in the *Human Info* treatment than in the *Human No Info* treatment (p=0.029).

[27]Only 1 out of 22 comparisons is significant at the 5% level, which falls within what random sampling variation would predict.

average. We have proportionate sample coverage across all 12 UK regions.[28]

We will check the robustness of all results to controlling for subject characteristics using multivariate regression analyses.

2.4.3 General Attitude toward AI

Demand for items generated by AI may be sensitive to attitudes towards and perceptions of AI technologies. To capture this heterogeneity across individuals, we adopted the AI Attitude Scale (AIAS-4) as a non-incentivised measure of general attitude toward AI (Grassini, 2023). This survey instrument uses a 10-point Likert response scale to elicit agreement (from 1 "Not at all" to 10 "Completely agree") with each of the following statements: (i) *I believe that AI will improve my life*; (ii) *I believe that AI will improve my work*; (iii) *I think I will use AI technology in the future*; and (iv) *I think AI technology is positive for humanity*. An AI Attitudes index is constructed for each individual in our sample as the within-person average across the four scores.

The resulting distribution of each of the component questions for the student sample in the lab experiment is presented in Figure E.1 and the distribution of the overall index is presented in Figure E.2. We find that the student subjects are most positive about their intention to use AI technologies in the future (mean value of 8.04) followed by the potential for AI to improve their work (mean value of 7.10) or their life (mean value of 6.58). Each of these distributions is negatively skewed, with a tail of responses to the left and most values found to the right of the mean. They are least positive about the overall impact of AI technology on society and humanity as a whole (mean value of 5.85) and the response distribution for this last item is more symmetric. Overall, the mean value of the 4-item AIAS index in our student sample is 6.89 (standard deviation 2.13).

The general population sample display a markedly less positive attitude toward AI on

[28]For a more detailed comparison with our quota targets, see Table F.3.

average than the students and the distributions are consistently less negatively skewed (Figures F.1 and F.2). Again, the participants are most positive about their intention to use AI technologies in the future (mean value of 6.30), but this is lower than the value reported above for the student sample. Average ratings for the three other components are also consistently lower (mean value of 5.31 for the potential of AI to improve their work; 5.54 for whether AI will improve their life; and 5.52 for whether AI is positive for humanity). It is worth pointing out that the final component (whether AI is positive for humanity) is most closely aligned across the two samples. Overall, the mean value of the 4-item AIAS index in our broad population sample is 5.67 (standard deviation 2.64).

We will investigate how a participant's WTP for human-made or AI-generated output relates to their score on this index in all three studies. Though this analysis is exploratory – it does not relate to a formal hypothesis – it was pre-registered.

3 Results

3.1 Main Results: Lab Experiments

Figure 2 presents the distribution of average bids in each treatment estimated at the participant level, based on all 10 auctions. Vertical dashed lines indicate treatment means. The comparison in each panel corresponds to one of our three pre-registered hypotheses.[29]

Contrary to Hypothesis 1 (upper-left panel), revealing AI generation increased rather than decreased valuations. Participants in the *AI Info* treatment bid on average £2.22, which is £0.58 higher than the £1.64 average in the *AI No Info* treatment. In addition to this difference between conditions not being statistically significant (p=0.185), the direction of the effect directly contradicts our hypothesis that disclosure of the AI source would

[29] All comparisons of treatment means are made using two-sample t-tests. All tests in this paper are two-sided. In Table E.3, we present a summary of test statistics for each hypothesis.

reduce consumer valuations. A Kolmogorov-Smirnov test further confirms that the overall distributions of bids do not significantly differ between these two conditions (p=0.641).

For Hypothesis 2 (upper-right panel), we find no evidence that disclosing the artworks are human-made affects valuations. Participants in the *Human Info* treatment bid on average £2.14, which is £0.13 higher than the £2.01 average bid in the *Human No Info* treatment (p=0.770). The distributions of bids in the two treatments are not significantly different at the 5% level (p=0.083). This result supports that disclosing a human source neither enhances nor diminishes valuations of the artworks.

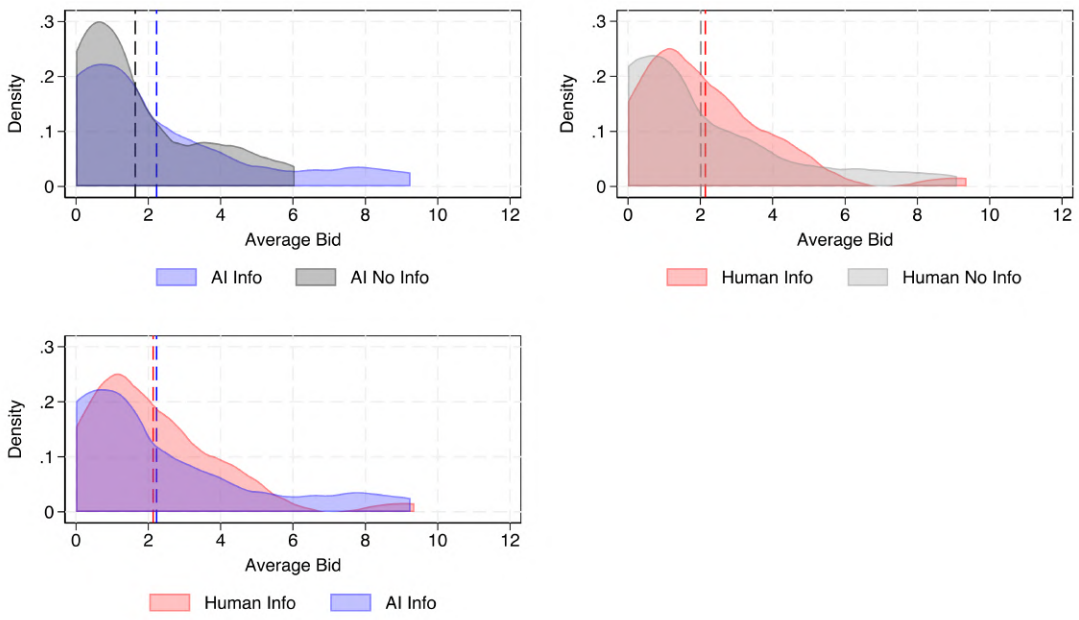
When comparing valuations between the *AI Info* and *Human Info* treatments for Hypothesis 3 (bottom panel), we find virtually identical bidding behaviour. Mean bids were £2.22 for AI-generated works and £2.14 for human-made works (p=0.841), with the Kolmogorov-Smirnov test indicating no significant difference in the two bid distributions (p=0.230). Thus, we reject Hypothesis 3 that participants would place a higher intrinsic economic value on human-made works.[30]

The results are robust to estimating the treatment effects at the bid-level and accounting for the large number of bids at zero by using censored regression (Tobit)[31] procedures and controlling for demographic and artwork characteristics, which also captures any residual differences in the baseline appeal of artworks between production sources as measured by the pre-testing process (see Table 1). These results suggest that, contrary to our expectations, knowledge of the artworks' origin has minimal impact on the valuation of work. Holding constant the baseline appeal of the artwork produced, participants

[30] We can also compare bids between the *AI No Info* and *Human No Info* treatments to validate the effectiveness of our design in presenting participants with AI-generated and human-made artworks of statistically indistinguishable baseline appeal. Figure E.3 demonstrates nearly identical bidding patterns between these conditions (i.e., when the source of the artwork is withheld).

[31] The results are robust to estimating the treatment effects by OLS and to clustering the standard errors at the bidder-level (Table E.4).

Figure 2: Distribution of Bids by Treatment: Lab Experiment



Notes: The figure shows kernel density estimates of the distribution of average bids (£) per individual. Dashed vertical lines show mean bids for each treatment condition. The upper-left panel compares bids for AI artworks when their source is informed vs not informed (H1), the upper-right panel compares bids for human artworks when their source is informed vs not informed (H2), and the lower panel compares bids when participants are informed about AI vs human sources (H3).

reveal no systematic preference for the human over AI source, and if anything value AI-generated works slightly more when informed of their origin.[32]

To further explore our treatment effects across different levels of baseline appeal, we compare among treatments the relationship between average bids and the independent aesthetic ratings obtained from the pre-testing process. Figure 3 presents these relationships for our three treatment comparisons; each point represents a single artwork's average bid plotted against its aesthetic rating from the pre-test survey.

In all panels, we observe that higher aesthetic ratings correspond to higher bids, indicating that participants' valuations aligned meaningfully with independent assessments.[33] Across treatments, the slope of the relationship between aesthetic rating and bids is very similar, suggesting that the artworks' source and the information provided about this had little impact on how baseline appeal translated into WTP. As exploratory analysis, we examine this formally by estimating bid-level models by Poisson, so that the treatment coefficient is a proportional effect, and interacting the treatment with the artwork's pre-test hypothetical WTP or aesthetic rating. This finds no evidence of heterogeneity in treatment effects by baseline appeal: the interaction is close to zero and insignificant in every comparison.[34] It is therefore the case that average bids for artworks of similar aesthetic appeal, whether high or low, are very similar across treatments: thus, the absence of treatment

[32]Bidders may not immediately appreciate that a second-price auction is strategy-proof, in which case bids could drift as they gain experience. We test this in Table E.5, which adds a linear round trend and its interaction with the treatment to a bid-level model with subject fixed effects. Because the order in which images were shown was randomised across sessions, the round number is not confounded with the image on screen. We find no evidence of learning of this kind: the round trend is economically negligible (between £0.01 and £0.03 per round) and statistically insignificant, and the interaction between the round and the treatment is insignificant in all three comparisons. The treatment effect therefore neither emerges nor fades as bidders gain experience, and our results are not an artefact of confusion in the early rounds.

[33]In Figure E.4, which replaces our artwork-specific aesthetic rating with a hypothetical WTP value obtained from a separate Prolific survey, we continue to observe a strong positive relationship with actual WTP (bids), and the figure reveals qualitatively similar patterns across treatments to those visible in Figure 3.

[34]Results are available on request.

effects holds true across the baseline appeal distribution. The upper-left panel reveals that WTP under *AI Info* is consistently higher than under *AI No Info* across the spectrum. However, for each image, this difference in mean bids is not significant.[35]

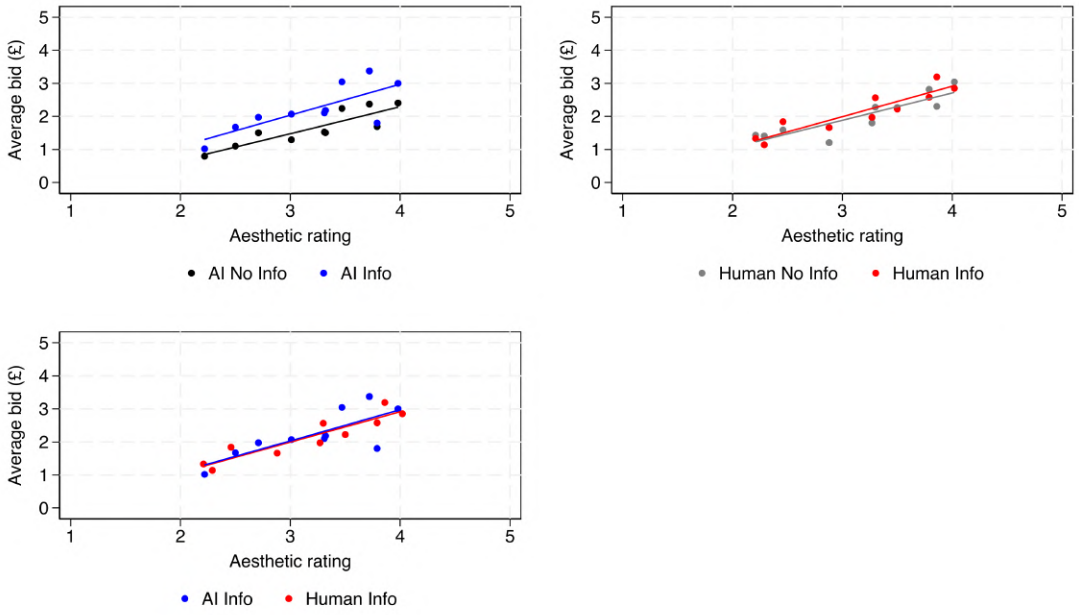
[35]See Table E.6 which reports hypothesis test results for each artwork. Figures E.5 and E.6 show each artwork's distribution of bids by treatment.

Table 1: Random-effects Tobit Regressions: Lab experiment

	Bid														
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
AI Info vs AI No Info	0.654	0.646	0.646	0.647	0.245										
	(0.646)	(0.647)	(0.647)	(0.647)	(0.650)										
Human Info vs Human No Info						0.154	0.163	0.164	0.164	0.003					
						(0.572)	(0.572)	(0.571)	(0.571)	(0.590)					
Human Info vs AI Info											0.056	0.120	0.103	0.097	0.515
											(0.642)	(0.643)	(0.643)	(0.643)	(0.666)
Aesthetic rating (pre-test)		1.529***		-0.649	-0.612		1.366***		0.024	0.025		1.521***		-0.484	-0.443
		(0.161)		(0.435)	(0.444)		(0.116)		(0.628)	(0.627)		(0.144)		(0.490)	(0.501)
Hypothetical WTP (pre-test)			0.589***	0.796***	0.789***			0.544***	0.535**	0.534**			0.599***	0.769***	0.760***
			(0.054)	(0.149)	(0.152)			(0.045)	(0.246)	(0.246)			(0.052)	(0.181)	(0.185)
18-21 years old					-1.164					-0.633					-1.104
					(0.741)					(0.678)					(0.797)
Male					0.562					-0.055					0.464
					(0.646)					(0.572)					(0.649)
UK national					-0.326					-0.034					0.482
					(0.762)					(0.692)					(0.835)
Income above average					-0.608					0.473					-0.791
					(0.681)					(0.564)					(0.680)
Studying Economics or Finance					0.936					-1.118					-0.978
					(0.893)					(0.772)					(0.920)
Art knowledge					-0.133					0.569*					0.032
					(0.325)					(0.324)					(0.318)
Constant	0.468	-4.431***	-2.010***	-0.799	0.066	1.136***	-3.182***	-1.122**	-1.160	-2.056	1.116**	-3.757***	-1.408***	-0.577	-0.399
	(0.467)	(0.706)	(0.527)	(0.965)	(1.405)	(0.399)	(0.547)	(0.444)	(1.076)	(1.354)	(0.456)	(0.657)	(0.513)	(0.983)	(1.379)
Uncensored obs.	654	654	654	654	634	757	757	757	757	757	703	703	703	703	683
Observations	1,080	1,080	1,080	1,080	1,060	1,160	1,160	1,160	1,160	1,160	1,120	1,120	1,120	1,120	1,100
Subjects	108	108	108	108	106	116	116	116	116	116	112	112	112	112	110

Notes: The dependent variable is the bid amount (£). The table shows estimates from random effects tobit regressions, with bids censored below at £0 and above at £12. Columns (1)-(5) compare the *AI Info* and *AI No Info* treatments (H1), columns (6)-(10) the *Human Info* and *Human No Info* treatments (H2), and columns (11)-(15) the *Human Info* and *AI Info* treatments (H3). Within each block, the first column includes the treatment only, the second adds the pre-test aesthetic rating, the third the pre-test hypothetical WTP, the fourth both, and the fifth adds individual-level controls. The two pre-test item scores are close substitutes, so the fourth and fifth columns, in which they enter together, are reported for comparability with the online experiments rather than as our preferred specification. Standard errors are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure 3: Relationship Between Aesthetic Appeal and Auction Bids by Treatment: Lab experiment



Notes: Each point shows the mean bid (£) and aesthetic rating for one of the 10 artworks per treatment condition. The aesthetic rating, elicited in the Prolific pre-testing phase, takes a value from 1 to 5, where 1 is "very unappealing" and 5 is "very appealing". Lines show linear fits. The upper-left panel compares bids for AI artworks when their source is informed vs not informed (H1), the upper-right panel compares bids for human artworks when their source is informed vs not informed (H2), and the lower panel compares bids when participants are informed about AI vs human sources (H3).

3.2 Main Results: Online Experiments

3.2.1 Artwork study

Figure 4 (upper-left panel) presents the distribution of average bids in each treatment of the online artwork experiment, estimated at the participant level based on all 10 auctions. Participants in the *AI Info* treatment bid on average £5.29, which is £0.39 lower than the £5.68 average in the *Human Info* treatment. While the sign of this difference is consistent with Hypothesis 4, its magnitude is modest and it is not statistically significant (p=0.191); the Kolmogorov-Smirnov test indicates the distributions are not significantly

different from one another (p=0.207).[36]

These tests do not account for the fact that, unlike in our lab experiment, the average baseline appeal of human-made artworks – as measured by hypothetical WTP, the variable used to select the final artworks for auction from the candidate sample – was somewhat below that of AI-generated artworks (see Section 2.1.2). Our regression analyses, which control for baseline appeal, therefore take on a greater importance in this experiment. Table 2 reports random-effects Tobit models that account for the censoring of bids at £0 and at the £12 endowment.[37] Consistent with the basic comparison of means, in the first model with no controls for artwork or subject characteristics, the estimated treatment effect is small and imprecise (column 1). However, once we control for our measures of baseline appeal from the pre-tests, the coefficient on *AI info* becomes larger in magnitude and statistically significant, at -£0.73 with a control for aesthetic rating (column 2, p <0.1), -£0.76 with an independent or added control for hypothetical WTP (columns 3 and 4, p <0.05) and -£0.82 after adding the full set of artwork and demographic controls (column 5, p <0.05). The final estimate corresponds to approximately 24% of a standard deviation of average bids. Thus, in contrast to the lab study, after controlling for residual differences in the baseline appeal of artworks between production sources that remain after the pre-tests, we do find a significant premium for human-made artworks, supporting Hypothesis 4.

The lower-left panel of Figure 4 plots each artwork's average bid against its aesthetic rating from the pre-test survey. As in the lab experiment, higher aesthetic ratings correspond to higher bids in both treatments. The *AI info* line of best fit sits below the *Human info* line across the spectrum of aesthetics, consistent with the negative treatment effects in

[36] A summary of the Kolmogorov-Smirnov and t-test statistics, for each hypothesis in the online experiments, is presented in Table F.4.

[37] The results are robust to estimating the treatment effects by OLS and to clustering the standard errors at the bidder-level (Table F.5).

the regressions after holding the baseline appeal constant.[38] In this experiment, the line is also steeper for *Human info*, and Poisson regressions find that the treatment difference significantly increases as either aesthetic ratings or hypothetical WTP increases, suggesting that the premium on human-made items expands for more valuable products.

This aggregate effect masks heterogeneity by age. The left-hand panel of Figure 5 re-estimates the participant-level comparison between treatments for participants aged 18 to 34 – an age range more comparable to our student sample in the lab. The difference is practically zero (£6.23 versus £6.38, p=0.769).[39] Table 3 re-estimates the fully controlled specification separately for participants aged 18 to 34 and those aged 35 and over. Among the younger group, the estimated effect of *AI Info* is small and statistically insignificant (-£0.50, p=0.458). Among participants aged 35 and over, by contrast, the estimated effect is -£0.96 and significant at the 5% level. The demand-side preference for human-made artwork is therefore concentrated among older consumers in our sample: as in the lab experiment conducted with students, we find no evidence of a significant premium among the young.

3.2.2 Poetry mug study

The distributions of bids in the online poetry mug experiment are presented in the upper-right hand panel of Figure 4. Participants in the *AI info* treatment bid on average £4.94 for the poetry mugs, which is £0.70 lower than the £5.64 average in the *Human info* treatment. This difference is statistically significant (p=0.020) and a Kolmogorov-Smirnov test confirms that the two bid distributions differ significantly (p=0.013). Unlike in the artwork online experiment, the average baseline appeal as measured in pre-testing was very simi-

[38] Figure F.3 replaces the aesthetic rating with the hypothetical WTP measure obtained from the pre-testing surveys and reveals qualitatively similar patterns.

[39] A summary of the Kolmogorov-Smirnov and t-test statistics, for each hypothesis in the online experiments, for this age group is presented in Table F.6.

lar for the human and AI-written haikus used in our auctions. These test results therefore suggest that consumers value poetry mugs less when they are known to be AI-generated rather than human-written.

The bid-level estimates reinforce this conclusion and are notably stable across specifications. The coefficient on *AI Info* is significant at the 5% level even without controlling for our pre-test measures of baseline appeal (column 6 of Table 2), and is essentially unchanged by the addition of the pre-test aesthetic ratings and hypothetical WTP (columns 7 to 9). With the full set of controls, the estimated effect is -£1.10 (column 10, p <0.01), corresponding to roughly 20% of the mean bid in the *Human Info* treatment.[40] The lower-right panel of Figure 4 presents the same vertical separation observed for artworks: the relationship between baseline appeal and bids is common to both treatments, but the *AI Info* line lies below the *Human Info* line throughout. As with artworks there is some evidence that the treatment difference expands when baseline appeal is higher; Poisson regressions find the treatment effect significantly interacts with aesthetic ratings, though not with hypothetical WTP.[41]

The age patterns mirror the results for the online artwork experiment and, if anything, are even starker. Among 18 to 34 year olds the regressions find that the effect of *AI Info* is indistinguishable from zero (-£0.07; the participant-level means are £5.79 and £5.81, p=0.963), while among those aged 35 and over it is -£1.56 and significant at the 1% level (see Table 3 and Figure 5).

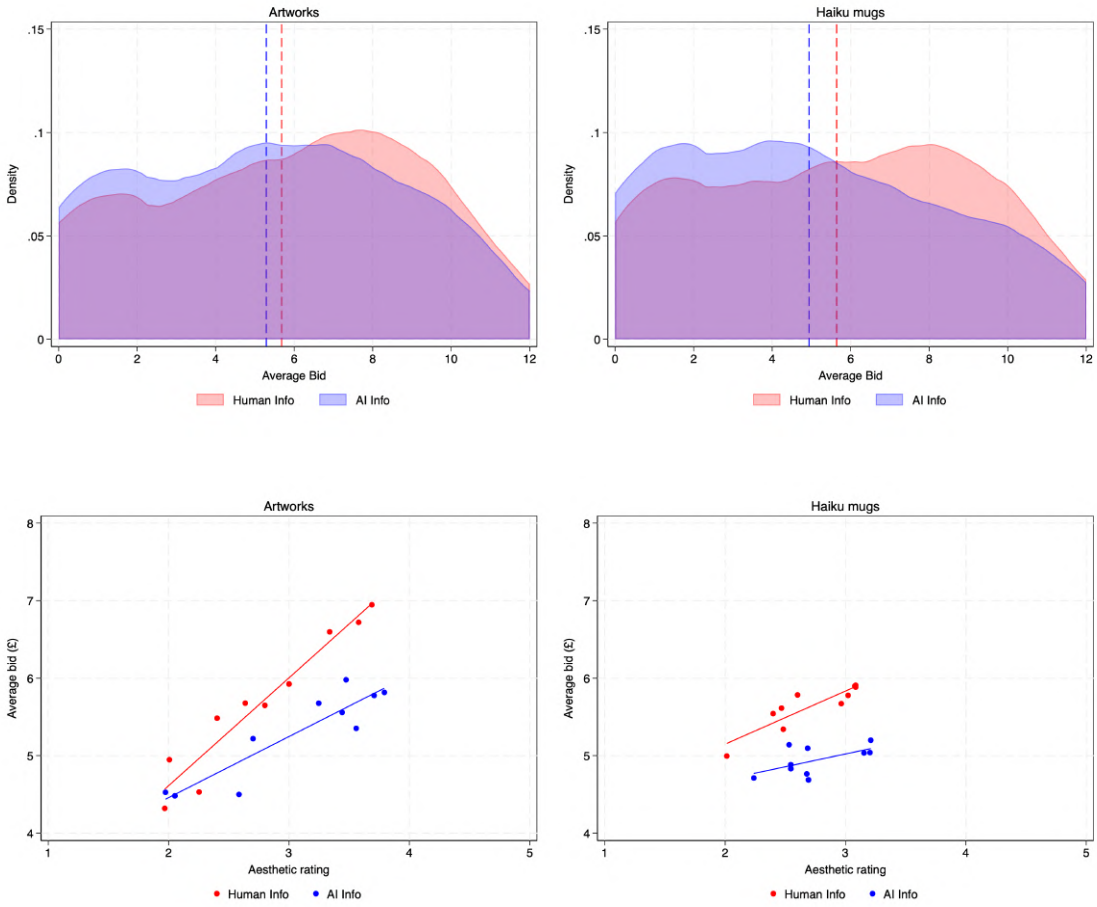
Pooling the two online experiment bid samples together confirms that the penalty for AI-generated content is common to both item types rather than specific to either. Columns 11 to 15 of Table 2 estimate the treatment effect on the combined sample, including a

[40] As in the lab, we check how bids change over the rounds for the online experiments. The results are shown in Table F.7. In neither case does the round interact significantly with the treatment, so the treatment effect again neither emerges nor fades as bidders gain experience.

[41] Full results of all Poisson regressions are available on request.

dummy for the poetry mug and its interaction with *AI Info*. In the full-controls specification, the coefficient on *AI Info* is -£0.76 (p <0.05), while the interaction term is small and statistically insignificant (-£0.33, p >0.10), as it is in every specification. We therefore cannot reject the null that consumers discount AI-generated work by the same amount whether it is an artwork or a poem. The pooled age pattern in Table 3 tells the same story as in the two experiments separately: the baseline AI penalty for artworks is undetectable among 18 to 34 year olds ($-0.37, p = 0.563$) but significant among those aged 35 and over ($-0.93, p < 0.05$), with no significant interaction for poetry mugs in either age group.

Figure 4: Bids in Online Experiments: Distributions and Relationship with Aesthetic Appeal, by Treatment



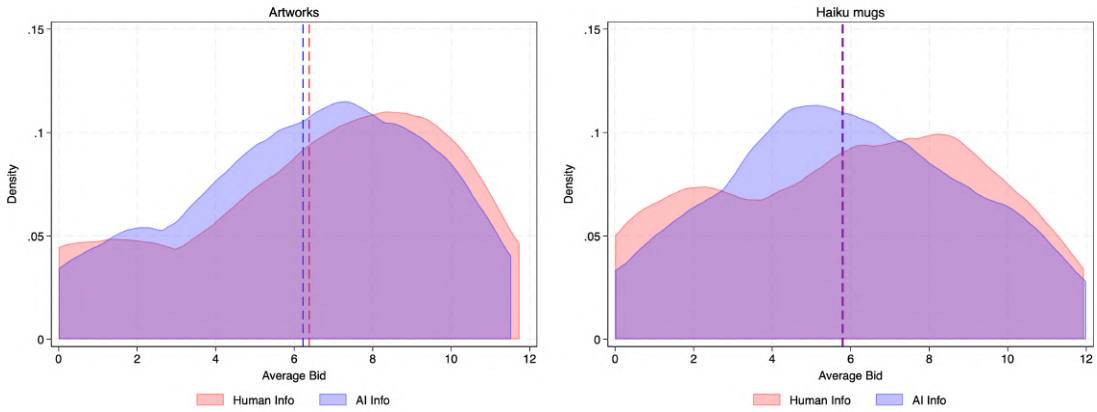
Notes: The left-hand plots show artworks (H4) and the right-hand plots poetry mugs (H5). The upper row shows kernel density estimates of the distribution of average bids (£) per individual, with dashed vertical lines at the mean bid for each treatment condition. The lower row plots the mean bid (£) against the aesthetic rating for each of the 10 artworks or 10 haiku mugs per treatment condition; the aesthetic rating, elicited in the Prolific pre-testing phase, runs from 1 ("very unappealing") to 5 ("very appealing"), and lines show linear fits.

Table 2: Random-effects Tobit Regressions: Online experiments

	Bid														
	Artworks					Haiku mugs					Pooled				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
AI Info	-0.349	-0.731*	-0.758**	-0.757**	-0.816**	-0.893**	-0.911**	-0.906**	-0.908**	-1.101***	-0.353	-0.682*	-0.730*	-0.730*	-0.755**
	(0.381)	(0.382)	(0.382)	(0.382)	(0.377)	(0.363)	(0.364)	(0.364)	(0.364)	(0.359)	(0.373)	(0.373)	(0.373)	(0.373)	(0.367)
Aesthetic rating (pre-test)		1.323***		0.187	0.183		0.635***		0.145	0.149		1.142***		0.109	0.110
		(0.054)		(0.200)	(0.200)		(0.065)		(0.128)	(0.129)		(0.040)		(0.116)	(0.116)
Hypothetical WTP (pre-test)			0.700***	0.607***	0.607***			0.448***	0.367***	0.367***			0.648***	0.592***	0.591***
			(0.028)	(0.103)	(0.103)			(0.042)	(0.083)	(0.083)			(0.021)	(0.062)	(0.063)
25-34 years old					-0.710					-0.446					-0.583
					(0.695)					(0.712)					(0.495)
35-44 years old					-1.303*					-0.400					-0.935*
					(0.682)					(0.745)					(0.502)
45-54 years old					-1.134					-0.300					-0.773
					(0.716)					(0.727)					(0.508)
55+ years old					-0.979					-1.100*					-1.098**
					(0.664)					(0.669)					(0.469)
Female					0.183					-0.308					-0.061
					(0.378)					(0.361)					(0.259)
UK national					0.472					0.248					0.350
					(0.776)					(0.751)					(0.532)
Income above average					0.383					0.535					0.455
					(0.464)					(0.457)					(0.324)
Student or degree holder					-0.504					-1.218***					-0.869***
					(0.430)					(0.401)					(0.293)
Knowledge of art/poetry					0.662***					0.785***					0.701***
					(0.196)					(0.193)					(0.136)
Mug											0.153	0.213	0.495	0.472	0.596
											(0.370)	(0.370)	(0.370)	(0.371)	(0.366)
AI Info × Mug											-0.558	-0.259	-0.197	-0.199	-0.326
											(0.524)	(0.524)	(0.524)	(0.524)	(0.518)
Constant	5.363***	1.700***	3.341***	3.092***	2.119*	5.572***	3.845***	4.510***	4.308***	3.665***	5.402***	2.237***	3.526***	3.384***	2.596***
	(0.270)	(0.310)	(0.282)	(0.389)	(1.141)	(0.254)	(0.310)	(0.273)	(0.327)	(1.099)	(0.264)	(0.287)	(0.272)	(0.311)	(0.796)
Uncensored obs.	4,213	4,213	4,213	4,213	4,199	4,581	4,581	4,581	4,581	4,551	8,794	8,794	8,794	8,794	8,750
Observations	5,250	5,250	5,250	5,250	5,230	5,410	5,410	5,410	5,410	5,380	10,660	10,660	10,660	10,660	10,610
Respondents	525	525	525	525	523	541	541	541	541	538	1,066	1,066	1,066	1,066	1,061

Notes: The dependent variable is the bid amount (£). The table shows estimates from random effects tobit regressions, with bids censored below at £0 and above at £12. Columns (1)-(5) use the artwork experiment, columns (6)-(10) the haiku mug experiment and columns (11)-(15) pool the two; the pooled columns include a mug dummy and its interaction with *AI Info*. Standard errors are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure 5: Distribution of Bids by Treatment: Online experiments, 18-34 year olds



Notes: The figure shows kernel density estimates of the distribution of average bids (£) per individual, for the subsample of participants aged 18 to 34. Dashed vertical lines show mean bids for each treatment condition. The left panel compares bids for artworks when the source is AI vs human (H4) and the right panel compares bids for haiku mugs when the source is AI vs human (H5).

Table 3: Random-effects Tobit Regressions by Age: Online experiments

	Bid					
	18-34 year olds			35+ year olds		
	Artworks (1)	Mugs (2)	Pooled (3)	Artworks (4)	Mugs (5)	Pooled (6)
AI Info	-0.502	-0.068	-0.367	-0.964**	-1.563***	-0.928**
	(0.676)	(0.623)	(0.635)	(0.450)	(0.439)	(0.446)
Aesthetic rating (pre-test)	0.573*	0.017	0.188	0.004	0.207	0.078
	(0.336)	(0.245)	(0.207)	(0.249)	(0.151)	(0.140)
Hypothetical WTP (pre-test)	0.341**	0.195	0.458***	0.729***	0.434***	0.649***
	(0.173)	(0.159)	(0.111)	(0.128)	(0.098)	(0.076)
Female	0.011	-1.206*	-0.653	0.276	0.107	0.200
	(0.690)	(0.618)	(0.462)	(0.452)	(0.440)	(0.312)
UK national	0.358	-0.072	0.123	0.273	0.374	0.383
	(1.168)	(1.000)	(0.765)	(1.038)	(1.082)	(0.725)
Income above average	0.419	0.205	0.342	0.231	0.778	0.474
	(0.738)	(0.752)	(0.527)	(0.581)	(0.563)	(0.401)
Student or degree holder	0.746	-0.165	0.146	-0.845*	-1.447***	-1.161***
	(0.835)	(0.708)	(0.541)	(0.496)	(0.474)	(0.341)
Knowledge of art/poetry	0.865***	0.486	0.687***	0.515**	0.968***	0.739***
	(0.329)	(0.346)	(0.238)	(0.236)	(0.234)	(0.164)
Mug			0.063			0.811*
			(0.642)			(0.441)
AI Info × Mug			0.322			-0.562
			(0.920)			(0.622)
Constant	-0.140	4.582***	2.150*	1.857	2.139*	1.473*
	(1.694)	(1.495)	(1.168)	(1.261)	(1.226)	(0.878)
Uncensored obs.	1,321	1,330	2,651	2,878	3,221	6,099
Observations	1,620	1,510	3,130	3,610	3,870	7,480
Respondents	162	151	313	361	387	748

Notes: The dependent variable is the bid amount (£). The table shows estimates from random effects tobit regressions, with bids censored below at £0 and above at £12, estimated separately for participants aged 18 to 34 and for those aged 35 and over. Within each group, the first column uses the artwork experiment, the second the haiku mug experiment and the third pools the two; the pooled columns include a mug dummy and its interaction with *AI Info*. All specifications include the pre-test aesthetic rating, the pre-test hypothetical WTP, demographic controls and knowledge of the good. Standard errors are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

3.3 Moderating Effects of General Attitude toward AI

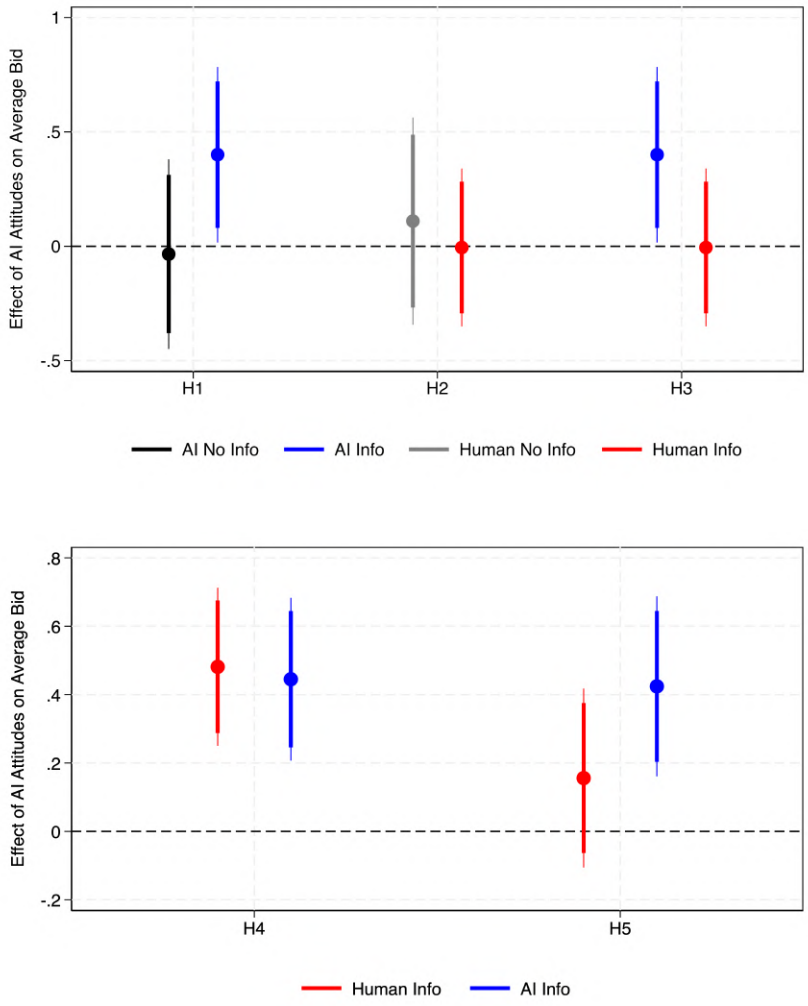
Considering differences in the rate of adoption (Bick et al., 2024) and occupation-level exposure to LLMs (Eloundou et al., 2024), it is likely that these interactions with AI would moderate how one values work done by AI generally. As such, it is plausible that those who view AI as complementary to human capability may value AI-generated work greater than those who view AI as a substitute for human labour. With this in mind, Figure 6 presents estimated effects of attitudes toward AI on average bids within each treatment, for the lab experiment (upper panel) and the online experiment (lower panel). These estimates are obtained from Tobit regressions that control for demographic characteristics.

In the lab experiment, we observe a clearly visible pattern for AI artworks: while more positive AI attitudes have no apparent effect on bids when the AI source is not disclosed, they are associated with significantly higher bids when AI authorship is revealed. Specifically, a one-unit increase in AI attitudes is associated with approximately a £0.40 higher bid in the *AI Info* treatment (p=0.041), whereas in the *AI No Info* treatment the effect is essentially zero. The magnitude of the estimate is economically meaningful also: it implies an increase of approximately 18% of the mean bid (£2.22) or 15% of a standard deviation (£2.66). For human-made artworks, the estimates suggest no significant effect of attitudes toward AI on bids for human artwork, regardless of whether or not the source is revealed (p=0.790 for *Human Info* and p=0.371 for *Human No Info*).[42]

The online experiments allow us to ask whether this relationship survives in a sam-

[42]Our results remain unchanged when using an alternative OLS estimator (Figure E.7) or when conducting bid-level analysis using a random-effects Tobit estimator (Figure E.8). We also examine whether our results regarding the effects of the AI Index reflect attitudes toward AI specifically, rather than mere knowledge of AI capabilities. We replace our AI attitudes index with a measure of participants' knowledge about AI image generation algorithms. The results, presented in Figure E.9, reveal no effect of AI knowledge on bids. Similarly, when we consider participants' art knowledge, Figure E.10 suggests a slight preference for human work among more art-knowledgeable participants, but this difference is not statistically significant at the 5% level.

Figure 6: Effect of AI Attitudes on Bids by Treatment



Notes: The dependent variable is the average bid (£) per individual. Points show the marginal effects of the AI Attitudes Index in each treatment condition, estimated from four separate Tobit regressions. Thick (thin) lines denote 90% (95%) confidence intervals. The upper panel corresponds to the lab experiment, where the comparisons are between the informed and uninformed conditions for AI art (H1) and human art (H2), and between the AI and human informed conditions (H3); these specifications include individual-level controls (age 18-21, male, UK national, above average income, studying economics or finance, AI development awareness, AI use frequency, and art knowledge). The lower panel corresponds to the online experiments, where the comparisons are between the AI and human informed conditions for artworks (H4) and haiku mugs (H5); these specifications include individual-level controls (age band, female, UK national, above average income, student or degree holder, AI development awareness, AI use frequency, and art/poetry knowledge).

ple drawn from the broad UK population, and whether it extends to a second domain. The effect of AI attitudes on bids for AI-generated work replicates closely, and in both domains. A one-unit increase in the attitudes index is associated with a £0.45 higher bid for AI-generated artworks and a £0.42 higher bid for AI-generated poetry mugs, both significant at the 1% level and both remarkably close to the £0.40 estimated in the lab. These magnitudes are again economically meaningful: they correspond to roughly 8% of the mean bid in the *AI Info* treatments, so that moving from an AI-sceptic to an AI-enthusiast – a shift of one standard deviation in the index (2.68 points) – is worth around £1.20, or 20% of the mean bid.

The pattern in the *Human Info* treatments differs between the two domains. For poetry mugs, the online experiment reproduces the pattern observed in the lab: the effect of AI attitudes on bids for human-written work is small and statistically insignificant (a one-unit increase in AI attitudes is associated with a £0.16 higher bid, p=0.240). For artworks, by contrast, AI attitudes predict bids for human-made work just as strongly as for AI-generated work, which suggests that AI attitudes may be correlated with other unobserved characteristics in this treatment.

Taken together, we read these results as follows. Attitudes toward AI robustly predict what consumers are willing to pay for AI-generated work: the estimate is stable in magnitude across two settings and two product domains. Aside from a single anomalous result in the *Human Info* condition for artworks, the findings are consistent. What our data support, then, is that pre-existing views about AI technology shift the valuation of production, and shift it most clearly where the work is known to be AI-generated.

4 Discussion

In this paper, we experimentally investigate whether AI-generated content is inherently devalued relative to human-made content of comparable baseline appeal. In our lab experiment with a student sample, we find no significant premium placed on human-made images. If anything, there is a small premium on AI-generated images when their origin is revealed. This finding contradicted our pre-registered hypotheses. Our follow-up online experiments with a sample from the general UK population puts this finding into perspective: on average, consumers pay a significant premium for human-made items, artworks and poetry mugs alike. Yet this premium increases with age: among the youngest respondents — the demographic similar in age to our student sample — the online experiment reproduces the null result from the lab.

Our results align with existing research, which has mostly relied on hypothetical elicitation methods and has identified a general tendency for a premium to be placed on human-made outputs, though with null results also present in this literature (Qin et al., 2025; Zehnle et al., 2025). We note that our elicitations use a between-subject design, in which each subject was faced with either human-made or AI-generated items, rather than a within-subject design wherein each subject faced both. Prior evidence suggests the human premium may be enhanced by within-subject elicitations (Horton et al., 2023; Neef et al., 2024; Jacobs et al., 2026), suggesting that our findings in this direction may have been strengthened had we implemented such methods. Importantly, both within-subject and between-subject experiments carry external validity to real-world consumption settings. Cases will exist in which consumers are presented with choices between different products, some of which are human-made and others AI-generated; in other cases consumers will encounter purchase decisions over items of one or other production origin, presented to them in isolation. Our results are of greater relevance to cases of the latter type.

Our results also reveal substantial heterogeneity in the valuation of AI-generated work based on pre-existing attitudes toward AI technology. Participants with more positive attitudes toward AI assign significantly higher value to AI-generated items when informed of their origin. This polarisation may plausibly reflect the varying degrees to which individuals perceive AI as complementary to human capabilities.

The age pattern admits two interpretations, which our cross-sectional data cannot separate. The preference for human-made work may come with age: today's young consumers will grow into it, and the premium will renew with each generation. Or the preference for human-made work is tied to cohorts: consumers socialised with generative AI never acquire it, and the premium will fade as older cohorts are replaced. Which interpretation is correct will determine whether consumer preferences can help sustain human labour in the long run.

Three *caveats* merit consideration in relation to the generalisability of the findings. First, our experimental settings capture immediate valuation decisions but cannot speak to how preferences might evolve as AI-generated content becomes more prevalent in everyday life. Historical evidence suggests that adaptation to technological changes occurs only gradually (Comin and Hobijn, 2010).

Second, our experiments considered two product domains: visual artworks and poetry mugs. While it is somewhat reassuring that our online experiments obtain similar findings for the two domains, we must be cautious over the extent to which they can be extrapolated to other domains, art-related or otherwise. We certainly do not argue that our results are informative about markets for works of art (or similar) – which, unlike the copies in our auctions, are new, significant and engaging – and where provenance, historical significance, and investment can drive valuations for collectible items independently of their aesthetic appeal (Ashenfelter and Graddy, 2002; Etro and Stepanova, 2017). At the same time, it is plausible that, since the products in our studies are associated with human

creativity, consumers may feel a stronger sense of attachment to the producer than they would for a purely utilitarian product, such as computer code. Thus our evidence does not necessarily demonstrate that the raw preference for human output we detect would remain present in markets for such products. [43]

Third, our estimates all come from experimental settings, in which subjects know they are being observed and stakes are relatively low. Where we do find a premium on human-made output, we provide some evidence that it expands for more valuable products. We also note that our measurements are elicited in auctions, while in real-world markets a range of different pricing mechanisms are used.

For workers in fields susceptible to AI substitution, our findings offer a qualified silver lining: a consumer premium for human-made work exists, but it is absent among the young. Translating this premium into employment outcomes would require additional knowledge of demand elasticities and of AI's cost advantages in these markets. Whether the demand-side barrier we document fades with cohort replacement or renews with each generation, and whether it is strong enough to protect workers from AI-driven displacement, remain open empirical questions.

[43] We thank the Editor and an anonymous referee for providing clarifications in relation to these important factors.

5 References

Abel, M. and Johnson, R. (2025). AI bias for creative writing: Subjective assessment versus willingness to pay. *IZA Discussion Paper No. 17646*.

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Ashenfelter, O. and Graddy, K. (2002). Art auctions: A survey of empirical studies. *NBER Working Paper*, (w8997).

Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Beckett, L. (2024). World's first AI art museum to explore 'creative potential of machines' in LA. The Guardian (September 25), Available at: https://www.theguardian.com/us-news/2024/sep/25/ai-art-museum-los-angeles-dataland (last accessed February 27, 2025).

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative AI. Technical report, National Bureau of Economic Research.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942.

Burton, J. W., Stein, M.-K., and Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2):220–239.

Castelo, N., Bos, M. W., and Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5):809–825.

Celebi, C., Exley, C., Harrs, S., Kivimaki, H., Serra-Garcia, M., and Yusof, J. (2026). Mission possible: The collection of high-quality data. *CESifo Working Paper No. 12535*.

Charness, G. and Grieco, D. (2024). Creativity and AI. *Available at SSRN 4686415*.

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

Danz, D., Vesterlund, L., and Wilson, A. J. (2022). Belief elicitation and behavioral incentive compatibility. *American Economic Review*, 112(9):2851–2883.

Dargnies, M.-P., Hakimov, R., and Kübler, D. (2026). Aversion to hiring algorithms: Transparency, gender profiling, and self-confidence. *Management Science*, 72(1):285–301.

Dietvorst, B. J., Simmons, J. P., and Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1):114.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.

Ernst, E., Merola, R., and Samaan, D. (2019). Economics of artificial intelligence: Implications for the future of work. *IZA Journal of Labor Policy*, 9(1).

Etro, F. and Stepanova, E. (2017). Art auctions and art investment in the golden age of British painting. *Scottish Journal of Political Economy*, 64(2):191–225.

Feng, W. and Song, X. (2026). Consumers' preference for human-generated versus AI-generated artwork: A social identification account. *Journal of Consumer Behaviour*, 25(1):118–137.

Friedrichsen, J., Schwarz, J., and Clement, M. (2026). When music is made by AI: Effects on preferences and willingness to pay.

Fuchs, C., Schreier, M., and Van Osselaer, S. M. (2015). The handmade effect: What's love got to do with it? *Journal of Marketing*, 79(2):98–110.

Goldstein, J. A., Chao, J., Grossman, S., Stamos, A., and Tomz, M. (2024). How persuasive is AI-generated propaganda? *PNAS Nexus*, 3(2):pgae034.

Gorenz, D. and Schwarz, N. (2024). How funny is ChatGPT? A comparison of human- and A.I.-produced jokes. *PLOS ONE*, 19(7):1–13.

Granulo, A., Fuchs, C., and Puntoni, S. (2021). Preference for human (vs. robotic) labor is stronger in symbolic consumption contexts. *Journal of Consumer Psychology*, 31(1):72–80.

Grassini, S. (2023). Development and validation of the AI attitude scale (AIAS-4): A brief measure of general attitude toward artificial intelligence. *Frontiers in Psychology*, 14.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1):114–125.

Harrison, G. W., Harstad, R. M., and Rutström, E. E. (2004). Experimental methods and elicitation of values. *Experimental Economics*, 7:123–140.

Healy, P. J. and Leo, G. (2025). Belief elicitation: A user's guide. In *Handbook of Experimental Methodology*, volume 1, pages 81–162. North-Holland.

Herbold, S., Hautli-Janisz, A., Heuer, U., Kikteva, Z., and Trautsch, A. (2023). A large-scale comparison of human-written versus ChatGPT-generated essays. *Scientific Reports*, 13(1):18617.

Hill, A. (2025). AI art on trial at Christie's. *Columbia Journal of Law & the Arts*.

Horton, C. B., White, M. W., and Iyengar, S. S. (2023). Bias against AI art can enhance perceptions of human creativity. *Scientific Reports*, 13(1):19001.

Jacobs, O., Pazhoohi, F., Mullen, G., and Kingstone, A. (2026). Comparative designs reveal preferences for human-generated rather than AI-generated art. *Empirical Studies of the Arts*, 44(2):729–743.

Jussupow, E., Benbasat, I., and Heinzl, A. (2020). Why are we averse towards algorithms? A comprehensive literature review on algorithm aversion. *Proceedings of the 28th European Conference on Information Systems (ECIS)*.

Kanazawa, K., Kawaguchi, D., Shigeoka, H., and Watanabe, Y. (2026). AI, skill, and productivity: The case of taxi drivers. *Management Science*, 72(2):1376–1388.

Kirk, U., Skov, M., Hulme, O., Christensen, M. S., and Zeki, S. (2009). Modulation of aesthetic value by semantic context: An fMRI study. *NeuroImage*, 44(3):1125–1132.

Lane, T., Pickard, H., and Walker, M. (2024). Valuing AI art: An experimental study. *AEA RCT Registry, AEARCTR-0014469*.

Lane, T., Pickard, H., Walker, M., and Heinrich, T. (2026). Valuing AI art: An experimental study. *AEA RCT Registry, AEARCTR-0018880*.

Locci, I. and Massoni, S. (2024). Judged by robots: Preferences and perceived fairness of algorithmic versus human punishments. *Review of Law & Economics*.

Lu, Y. and Zhou, Y. (2021). A review on the economics of artificial intelligence. *Journal of Economic Surveys*, 35(4):1045–1072.

Mahmud, H., Islam, A. N., Ahmed, S. I., and Smolander, K. (2022). What influences algorithmic decision-making? A systematic literature review on algorithm aversion. *Technological Forecasting and Social Change*, 175:121390.

Mandel, G. and Imas, A. (2026). Art and the machine: Why people devalue AI-generated creative work. *Available at SSRN 6302659*.

Millet, K., Buehler, F., Du, G., and Kokkoris, M. D. (2023). Defending humankind: Anthropocentric bias in the appreciation of AI art. *Computers in Human Behavior*, 143:107707.

Neef, N. E., Zabel, S., Papoli, M., and Otto, S. (2024). Drawing the full picture on diverging findings: Adjusting the view on the perception of art created by artificial intelligence. *AI & Society*, pages 1–21.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

Pope, A. (2024). AI artwork of Alan Turing sells for $1m. BBC (November 07), Available at: https://www.bbc.co.uk/news/articles/cpqdvz4w45wo (last accessed February 27, 2025).

Porter, B. and Machery, E. (2024). AI-generated poetry is indistinguishable from human-written poetry and is rated more favorably. *Scientific Reports*, 14(1):26133.

Qin, X., Zhou, X., Chen, C., Wu, D., Zhou, H., Dong, X., Cao, L., and Lu, J. G. (2025). AI aversion or appreciation? A capability–personalization framework and a meta-analytic review. *Psychological Bulletin*, 151(5):580.

Rutström, E. E. (1998). Home-grown values and incentive compatible auction design. *International Journal of Game Theory*, 27:427–441.

Schotter, A. and Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, 6(1):103–128.

Stephenson, R. and Armstrong, C. (2026). Student generative artificial intelligence survey 2026. Technical Report HEPI Report 199, Higher Education Policy Institute.

Trautmann, S. T. and van de Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *Economic Journal*, 125(589):2116–2135.

van Waning, J. (2024). Human art vs AI art, a potential danger for artist?: Artistic and economic evaluations across multiple art genres. Master's thesis.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37.

Yee, L., Madgavkar, A., Smit, S., Krivkovich, A., Chui, M., Ramirez, M. J., and Castresana, D. (2025). Agents, robots, and us: Skill partnerships in the age of AI. *McKinsey Global Institute*.

Zehnle, M., Hildebrand, C., and Valenzuela, A. (2025). Not all AI is created equal: A meta-analysis revealing drivers of AI resistance across markets, methods, and time. *International Journal of Research in Marketing*.

Zhang, H., Bai, X., and Ma, Z. (2022). Consumer reactions to AI design: Exploring consumer willingness to pay for AI-designed products. *Psychology & Marketing*, 39(11):2171–2183.

Zhu, F. and Zou, W. (2023). The role of generative AI in human creative processes: Experimental evidence. *Available at SSRN 4676053*.

A Silver Lining? Consumer Preferences over Human-Made and AI-Generated Production in Two Domains

July 17, 2026

Supplementary Information (SI) Appendix

For Online Publication

Tom Lane

Newcastle University

Harry Pickard

Newcastle University

Matthew J. Walker

Newcastle University

Timo Heinrich

Hamburg University of Technology

Contents

A AI Content Generation **3**

A.1 Artwork 3

A.1.1 Lab experiment 3

A.1.2 Online experiment 7

A.2 Poetry 9

B Experiment instructions and protocol **13**

B.1 Lab experiment 13

B.2 Online experiments 23

C Experiment auction program screenshots **35**

C.1 Lab experiment 35

C.2 Online experiments 47

D Post-experiment questionnaire **61**

D.1 Lab experiment 61

D.2 Online experiments 64

E Descriptives and additional data analyses: Lab experiment **77**

F Descriptives and additional data analyses: Online experiments **94**

G References **105**

A AI Content Generation

A.1 Artwork

A.1.1 Lab experiment

To create the AI-generated artworks for our experiment, we used Midjourney version 6.1,[1] a generative artificial intelligence algorithm that creates images from natural language descriptions. We used Midjourney's default parameters as of August 2024 and set the aspect ratio to 4:3.

The first approach leveraged Claude 3.5 Sonnet,[2] a Large Language Model from Anthropic, to generate a description of imaginary artwork. Our initial prompt was as follows: *"Generate some descriptions of artwork in various styles that you have imagined also briefly describe what material is used. Paintings or prints only, no sculptures etc. Give me 10."* We then followed up with some fine tuning of the prompt to remove an imaginary name given to each description, and then repeated a request to produce 10 more examples and keep them varied from one another. We iterated this process until 50 images were created.

The second approach began with randomly selected human artworks. We uploaded an image of each human artwork to Midjourney, which subsequently generated a description of the artwork. These descriptions then

[1] Available at midjourney.com.
[2] Available at claude.ai.

served as prompts for Midjourney to create new images, using the original human artwork as a style reference. We generated 50 images through this method.

For each prompt, Midjourney generates four variations simultaneously. We randomly selected one variation from each set and used the up-scale function to create a high-quality version. Figure A.1 demonstrates the four variations generated from a single prompt using our first approach, with the upper-left image selected for our sample. Figure A.2 shows the four variations generated using our second approach, alongside the original human artwork in Panel (a) that served as the style reference. The bottom-right image in Panel (b) was randomly selected for our sample.

At the time of conducting the lab experiment, Midjourney version 6.1 and Claude Sonnet 3.5 were frontier models.

Figure A.1: Midjourney creation: Example of approach 1



Notes: The four images are created using the following prompt: "An atmospheric watercolor and ink painting of a misty Japanese garden, with delicate brush strokes suggesting the forms of a bridge, lanterns, and maple trees." The upper-left hand panel was randomly selected as the image to enter the sample.

Figure A.2: Midjourney creation: Example of approach 2

Panel (b)

Panel (a)





Notes: Panel (a) shows the original human artwork. Panel (b) shows the four images generated from it using our second approach: Midjourney was given the artwork in Panel (a) as a style reference and produced the following text description of it, which was then used as the prompt: "A painting of an abstract piece with leaves and fruit". The bottom-right image in Panel (b) was randomly selected to enter the sample.

A.1.2 Online experiment

For the online experiment, we followed the same two approaches to artwork generation, again using the frontier models available at the time. The experiment was developed in May 2026, by which point Midjourney had reached version 7.0 and Claude had advanced to Sonnet 4.6. We used these versions to generate the corresponding image sets.

After the same pre-testing phases, we selected a final sample of 10 human and 10 AI images. The final AI images, alongside the final sample of human images, are presented in Figure A.3.

Figure A.3: Artworks auctioned in online experiment

AI Image 1



AI Image 2



AI Image 3



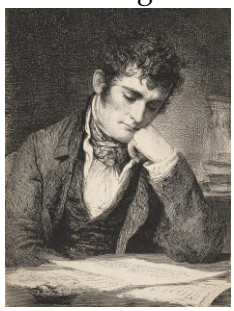
AI Image 4



AI Image 5



AI Image 6



AI Image 7



AI Image 8



AI Image 9



AI Image 10



Human Image 1



Human Image 2



Human Image 3



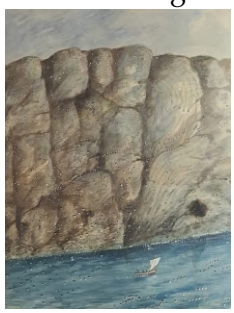
Human Image 4



Human Image 5



Human Image 6



Human Image 7



Human Image 8



Human Image 9



Human Image 10



A.2 Poetry

We generated AI-written haikus using Claude Sonnet 4.6. As with the artwork stream, we used two approaches.

In the first approach, Claude generated original haikus, rather than haikus based on existing work. We asked for 500 haikus, each labelled with one of seven styles (traditional, urban, philosophical, humorous, surreal, melancholic, minimalist), spread roughly evenly across the categories. We asked for a 5-7-5 syllable pattern, allowing small variations (e.g. 4-7-5 or 5-6-5) in around a third of the haikus. We also asked for a wide range of topics, with no two haikus repeating the same theme.

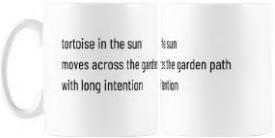
In the second approach, Claude was shown a sample of the Project Gutenberg haikus and asked to write new haikus in the same style. We gave Claude a style guide describing these haikus as fragments of early twentieth-century Western literary poetry, rather than traditional Japanese nature haikus, with a literary, dramatic, often melancholic tone, sometimes wry or comic, mixing formal or archaic language with plain modern speech, and often reading like part of a longer sentence. Syllable counts were left flexible, usually between four and eight per line. We gave Claude eight example lines from the Project Gutenberg corpus to show the target style, and again asked for a wide range of themes, with no repetition.

Together, these two approaches provided 1,000 AI-generated haikus, from

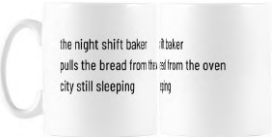
which we randomly selected 70 generated by the first approach and 30 generated by the second approach to take into pre-testing (the 70:30 split was chosen to match the proportion of AI artworks that survived the initial source-detection pre-test for our lab experiment that were generated using the first and second approaches outlined above). After the same pre-testing process, we selected a final sample of 10 human-written and 10 AI-written haikus. The final samples are presented as printed on the mugs in Figure A.4 below. For clarity, we also provide the poetry text in Table A.1.

Figure A.4: Poetry mugs auctioned in online experiment

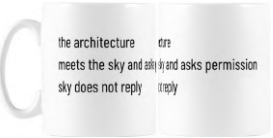
AI Mug 1



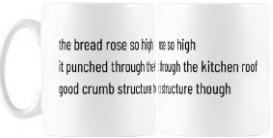
AI Mug 2



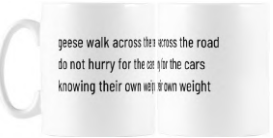
AI Mug 3



AI Mug 4



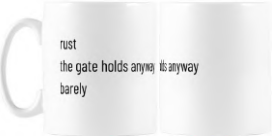
AI Mug 5



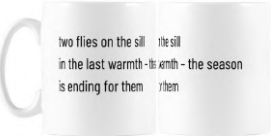
AI Mug 6



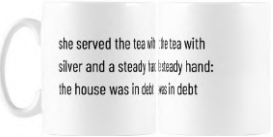
AI Mug 7



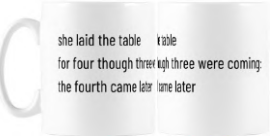
AI Mug 8



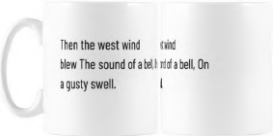
AI Mug 9



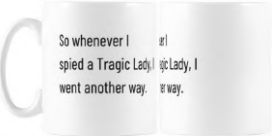
AI Mug 10



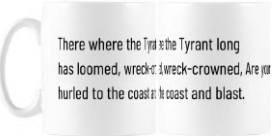
Human Mug 1



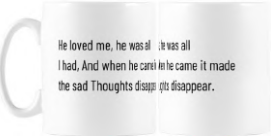
Human Mug 2



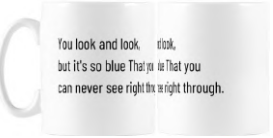
Human Mug 3



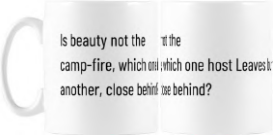
Human Mug 4



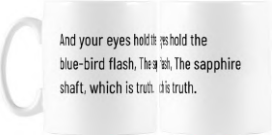
Human Mug 5



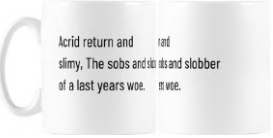
Human Mug 6



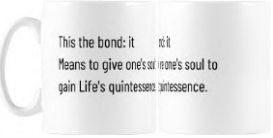
Human Mug 7



Human Mug 8



Human Mug 9



Human Mug 10

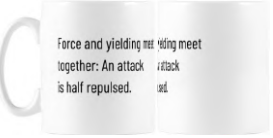


Table A.1: Poetry used in online experiment

AI Mug 1	tortoise in the sun moves across the garden path with long intention	Human Mug 1	Then the west wind blew The sound of a bell, On a gusty swell.
AI Mug 2	the night shift baker pulls the bread from the oven city still sleeping	Human Mug 2	So whenever I spied a Tragic Lady, I went another way.
AI Mug 3	the architecture meets the sky and asks permission sky does not reply	Human Mug 3	There where the Tyrant long has loomed, wreck-crowned, Are young and old hurled to the coast and blast.
AI Mug 4	the bread rose so high it punched through the kitchen roof good crumb structure though	Human Mug 4	He loved me, he was all I had, And when he came it made the sad Thoughts disappear.
AI Mug 5	geese walk across the road do not hurry for the cars knowing their own weight	Human Mug 5	You look and look, but it's so blue That you can never see right through.
AI Mug 6	bike share at midnight the helmets all still damp from the previous heads	Human Mug 6	Is beauty not the camp-fire, which one host Leaves burning for another, close behind?
AI Mug 7	rust the gate holds anyway barely	Human Mug 7	And your eyes hold the blue-bird flash, The sapphire shaft, which is truth.
AI Mug 8	two flies on the sill in the last warmth – the season is ending for them	Human Mug 8	Acrid return and slimy, The sobs and slobber of a last years woe.
AI Mug 9	she served the tea with silver and a steady hand: the house was in debt	Human Mug 9	This the bond: it Means to give one's soul to gain Life's quintessence.
AI Mug 10	she laid the table for four though three were coming: the fourth came later	Human Mug 10	Force and yielding meet together: An attack is half repulsed.

B Experiment instructions and protocol

B.1 Lab experiment

This section presents the three variants of our lab experiment instructions. Each variant is identical except for a specific treatment phrase that appears on the third page of the instructions. These variants correspond to our experimental treatments, whereby the instructions for *Human No Info* and *AI No Info* are the same (and so are only included once below).

The treatment phrases differ as follows: In the *Human Info* treatment, participants read "All artworks being sold are created by human artists." In the *AI Info* treatment, participants read "All artworks being sold are created by an artificial intelligence (AI) program." The *No Info* treatments contain no statement about the artwork's source.

Paper copies of the experiment instructions were distributed to all participants and the instructions were read aloud by the experimenter at the beginning of each session.

Experimental Instructions

Welcome. You are now taking part in a decision-making experiment. You were randomly selected from the Experimental and Behavioural Economics Laboratory's pool of subjects to be invited to participate in this session. There will be several pauses for you to ask questions. During such a pause, please raise your hand if you want to ask a question. Apart from asking questions in this way, you must not communicate with anybody else in this room. If you have a mobile phone, please may we request that you place this in the area designated for personal belongings (at the front of the laboratory) for the duration of the session.

Today's session

In today's session, you can earn money and may purchase a work of art. The amount you earn and whether or not you purchase a work of art depends on your decisions and those of others. You will receive a £12 participation payment for showing up on time and taking part in the experiment. You may choose to use none, some, or all of this amount to offer on an artwork.

Upon leaving the experiment, you will be given your cash earnings and/or artwork in private and you will be asked to sign a receipt. The receipt will be used by administrative personnel solely to ascertain that the funds allocated for the research project are properly utilized. Cash earnings will be rounded up to the nearest 10 pence (e.g., £8.40 if it is £8.33).

The session is made up of two parts. In the first part you will take part in a series of auctions for different works of art as described below. In the second part, you will complete a short questionnaire (your responses to the questionnaire will not affect your earnings in the experiment).

Instructions for Part I

I will now describe the task which makes up the first part of the session.

At the beginning of this part, the computer will randomly assign you to a group of 4 participants (including you). You will not know the identity of any other person in your group and no-one else in your group will know your identity.

You and the other three people in your group will take part in a type of auction known as the "second-price auction". In an auction procedure, everyone in the group submits a bid for an item for purchase. The person submitting the highest bid in the group buys the item. Unlike in other types of auctions that you may have encountered before, in the second-price auction, the person submitting the highest bid **purchases the item at a price equal to the next highest bid, which has been submitted by another participant in the group**.

Why are we using a "second-price auction" procedure?

The best strategy in this type of auction is to always bid exactly the maximum amount that you would be willing to pay for the item.

To understand why, let us assume that you have decided that you would be willing to pay up to £10 for the item. Let us call this your true value of the item. Which strategy is best for you: to bid higher than £10, lower than £10, or to bid exactly £10? The answer is: to bid exactly £10.

If you were to bid less than £10 (say £8) the auction might be won by a higher bidder whose bid is also less than £10 (say £9), and you will not be able to buy the item even though you were truly willing to pay more than the winner's bid (which is the amount you would have paid if you had bid higher). Therefore, bidding less than your true valuation is not a good strategy.

If you were to bid more than £10 (say £12) you might end up winning the auction but having to pay a price higher than what you really value the item at (say if the second highest bid was £11). Therefore, bidding more than your true valuation is not a good strategy.

Is there a point in trying to guess what the next highest bid after yours is, in order to bid just a penny more than that person? No, because the way the purchase price is determined in this auction you will only have to pay the price of the next highest bid anyway, independent of how much higher your bid was.

The best strategy in this type of auction is therefore to **always bid exactly the maximum amount that you decided that you would be willing to pay**. You do not risk paying more for the item than what you want to, and you will never pay more than the second highest bid anyway.

Now let us do a numerical example to illustrate.

Bidder	True value	Bid
1	£10	?
2	£7	£7
3	£4	£4
4	£1	£1

Suppose you are bidder 1. This is a situation where you would win the auction if you bid your true value. Would it be better for you to bid less than your true value, say £8? No. It does not matter to your purchasing price whether you bid £10 or £8, since you will still only pay £7.

Can you improve your situation by bidding more than your true value if you think other bidders are not bidding their true value? No. If bidder 2 overbids at £11, and you manage to outbid him or her, then you would have to pay £11 for the item, which is more than you are really willing to pay. If bidder 2 overbids at £9, a £10 bid from you will still win. If bidder 2 underbids at £6, a £10 bid from you will still win. There is no reason for you to overbid, independent of what other participants might decide to do.

Is there ever a reason for you to underbid? No. If bidder 2 overbids at £9, underbidding by you

could result in bidder 2 winning the auction even though you were really willing to pay £10. If bidder 2 underbids at £6, you will only pay £6 whether or not you bid £10 or £7. So why risk losing the auction unnecessarily by underbidding?

Therefore, the best strategy for you is always to bid exactly your true value for the item.

If you are not willing to pay any amount to purchase the item, then you should offer a zero bid.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** works of art for sale. For each work of art, a physical copy of it has been created, as a small print, on high quality card of size A5 (14.8 x 21 cm). Each physical copy is available for sale in the lab today.

All artworks being sold are created by an artificial intelligence (AI) program.

Each work of art will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The artwork will be presented on the computer screen.
2. Each bidder is invited to observe the physical copy of the artwork in the laboratory.
3. Each participant submits a bid for the artwork via the computer.
4. The computer ranks bids from highest to lowest (separately for each group of 4 participants). Ties will be broken randomly.

After all the rounds are completed, the computer will randomly select one of the 10 auction rounds as binding (separately for each group of 4 participants). Therefore, **one and only one artwork will be sold in each group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the artwork** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation payment of £12** to determine your final payment. Bids may not exceed £12.

Practice rounds

Before the main auction rounds, two practice rounds will be performed for some hypothetical items that are not available in the laboratory. In the practice rounds (only), you will be assigned a "true value" for the item (in other words, we will ask you to imagine that this amount is your maximum willingness to pay for the item) and you will bid against hypothetical opponents. The practice rounds are solely for familiarizing yourself with the auction procedure and will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

We will now move on to the practice rounds to check your understanding. Please raise your hand if you are stuck or have a question. After everyone in today's session has finished the practice rounds, the main auction rounds will begin.

Experimental Instructions

Welcome. You are now taking part in a decision-making experiment. You were randomly selected from the Experimental and Behavioural Economics Laboratory's pool of subjects to be invited to participate in this session. There will be several pauses for you to ask questions. During such a pause, please raise your hand if you want to ask a question. Apart from asking questions in this way, you must not communicate with anybody else in this room. If you have a mobile phone, please may we request that you place this in the area designated for personal belongings (at the front of the laboratory) for the duration of the session.

Today's session

In today's session, you can earn money and may purchase a work of art. The amount you earn and whether or not you purchase a work of art depends on your decisions and those of others. You will receive a £12 participation payment for showing up on time and taking part in the experiment. You may choose to use none, some, or all of this amount to offer on an artwork.

Upon leaving the experiment, you will be given your cash earnings and/or artwork in private and you will be asked to sign a receipt. The receipt will be used by administrative personnel solely to ascertain that the funds allocated for the research project are properly utilized. Cash earnings will be rounded up to the nearest 10 pence (e.g., £8.40 if it is £8.33).

The session is made up of two parts. In the first part you will take part in a series of auctions for different works of art as described below. In the second part, you will complete a short questionnaire (your responses to the questionnaire will not affect your earnings in the experiment).

Instructions for Part I

I will now describe the task which makes up the first part of the session.

At the beginning of this part, the computer will randomly assign you to a group of 4 participants (including you). You will not know the identity of any other person in your group and no-one else in your group will know your identity.

You and the other three people in your group will take part in a type of auction known as the "second-price auction". In an auction procedure, everyone in the group submits a bid for an item for purchase. The person submitting the highest bid in the group buys the item. Unlike in other types of auctions that you may have encountered before, in the second-price auction, the person submitting the highest bid **purchases the item at a price equal to the next highest bid, which has been submitted by another participant in the group**.

Why are we using a "second-price auction" procedure?

The best strategy in this type of auction is to always bid exactly the maximum amount that you would be willing to pay for the item.

To understand why, let us assume that you have decided that you would be willing to pay up to £10 for the item. Let us call this your true value of the item. Which strategy is best for you: to bid higher than £10, lower than £10, or to bid exactly £10? The answer is: to bid exactly £10.

If you were to bid less than £10 (say £8) the auction might be won by a higher bidder whose bid is also less than £10 (say £9), and you will not be able to buy the item even though you were truly willing to pay more than the winner's bid (which is the amount you would have paid if you had bid higher). Therefore, bidding less than your true valuation is not a good strategy.

If you were to bid more than £10 (say £12) you might end up winning the auction but having to pay a price higher than what you really value the item at (say if the second highest bid was £11). Therefore, bidding more than your true valuation is not a good strategy.

Is there a point in trying to guess what the next highest bid after yours is, in order to bid just a penny more than that person? No, because the way the purchase price is determined in this auction you will only have to pay the price of the next highest bid anyway, independent of how much higher your bid was.

The best strategy in this type of auction is therefore to **always bid exactly the maximum amount that you decided that you would be willing to pay**. You do not risk paying more for the item than what you want to, and you will never pay more than the second highest bid anyway.

Now let us do a numerical example to illustrate.

Bidder	True value	Bid
1	£10	?
2	£7	£7
3	£4	£4
4	£1	£1

Suppose you are bidder 1. This is a situation where you would win the auction if you bid your true value. Would it be better for you to bid less than your true value, say £8? No. It does not matter to your purchasing price whether you bid £10 or £8, since you will still only pay £7.

Can you improve your situation by bidding more than your true value if you think other bidders are not bidding their true value? No. If bidder 2 overbids at £11, and you manage to outbid him or her, then you would have to pay £11 for the item, which is more than you are really willing to pay. If bidder 2 overbids at £9, a £10 bid from you will still win. If bidder 2 underbids at £6, a £10 bid from you will still win. There is no reason for you to overbid, independent of what other participants might decide to do.

Is there ever a reason for you to underbid? No. If bidder 2 overbids at £9, underbidding by you

could result in bidder 2 winning the auction even though you were really willing to pay £10. If bidder 2 underbids at £6, you will only pay £6 whether or not you bid £10 or £7. So why risk losing the auction unnecessarily by underbidding?

Therefore, the best strategy for you is always to bid exactly your true value for the item.

If you are not willing to pay any amount to purchase the item, then you should offer a zero bid.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** works of art for sale. For each work of art, a physical copy of it has been created, as a small print, on high quality card of size A5 (14.8 x 21 cm). Each physical copy is available for sale in the lab today.

All artworks being sold are created by human artists.

Each work of art will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The artwork will be presented on the computer screen.
2. Each bidder is invited to observe the physical copy of the artwork in the laboratory.
3. Each participant submits a bid for the artwork via the computer.
4. The computer ranks bids from highest to lowest (separately for each group of 4 participants). Ties will be broken randomly.

After all the rounds are completed, the computer will randomly select one of the 10 auction rounds as binding (separately for each group of 4 participants). Therefore, **one and only one artwork will be sold in each group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the artwork** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation payment of £12** to determine your final payment. Bids may not exceed £12.

Practice rounds

Before the main auction rounds, two practice rounds will be performed for some hypothetical items that are not available in the laboratory. In the practice rounds (only), you will be assigned a "true value" for the item (in other words, we will ask you to imagine that this amount is your maximum willingness to pay for the item) and you will bid against hypothetical opponents. The practice rounds are solely for familiarizing yourself with the auction procedure and will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

We will now move on to the practice rounds to check your understanding. Please raise your hand if you are stuck or have a question. After everyone in today's session has finished the practice rounds, the main auction rounds will begin.

Experimental Instructions

Welcome. You are now taking part in a decision-making experiment. You were randomly selected from the Experimental and Behavioural Economics Laboratory's pool of subjects to be invited to participate in this session. There will be several pauses for you to ask questions. During such a pause, please raise your hand if you want to ask a question. Apart from asking questions in this way, you must not communicate with anybody else in this room. If you have a mobile phone, please may we request that you place this in the area designated for personal belongings (at the front of the laboratory) for the duration of the session.

Today's session

In today's session, you can earn money and may purchase a work of art. The amount you earn and whether or not you purchase a work of art depends on your decisions and those of others. You will receive a £12 participation payment for showing up on time and taking part in the experiment. You may choose to use none, some, or all of this amount to offer on an artwork.

Upon leaving the experiment, you will be given your cash earnings and/or artwork in private and you will be asked to sign a receipt. The receipt will be used by administrative personnel solely to ascertain that the funds allocated for the research project are properly utilized. Cash earnings will be rounded up to the nearest 10 pence (e.g., £8.40 if it is £8.33).

The session is made up of two parts. In the first part you will take part in a series of auctions for different works of art as described below. In the second part, you will complete a short questionnaire (your responses to the questionnaire will not affect your earnings in the experiment).

Instructions for Part I

I will now describe the task which makes up the first part of the session.

At the beginning of this part, the computer will randomly assign you to a group of 4 participants (including you). You will not know the identity of any other person in your group and no-one else in your group will know your identity.

You and the other three people in your group will take part in a type of auction known as the "second-price auction". In an auction procedure, everyone in the group submits a bid for an item for purchase. The person submitting the highest bid in the group buys the item. Unlike in other types of auctions that you may have encountered before, in the second-price auction, the person submitting the highest bid **purchases the item at a price equal to the next highest bid, which has been submitted by another participant in the group**.

Why are we using a "second-price auction" procedure?

The best strategy in this type of auction is to always bid exactly the maximum amount that you would be willing to pay for the item.

To understand why, let us assume that you have decided that you would be willing to pay up to £10 for the item. Let us call this your true value of the item. Which strategy is best for you: to bid higher than £10, lower than £10, or to bid exactly £10? The answer is: to bid exactly £10.

If you were to bid less than £10 (say £8) the auction might be won by a higher bidder whose bid is also less than £10 (say £9), and you will not be able to buy the item even though you were truly willing to pay more than the winner's bid (which is the amount you would have paid if you had bid higher). Therefore, bidding less than your true valuation is not a good strategy.

If you were to bid more than £10 (say £12) you might end up winning the auction but having to pay a price higher than what you really value the item at (say if the second highest bid was £11). Therefore, bidding more than your true valuation is not a good strategy.

Is there a point in trying to guess what the next highest bid after yours is, in order to bid just a penny more than that person? No, because the way the purchase price is determined in this auction you will only have to pay the price of the next highest bid anyway, independent of how much higher your bid was.

The best strategy in this type of auction is therefore to **always bid exactly the maximum amount that you decided that you would be willing to pay**. You do not risk paying more for the item than what you want to, and you will never pay more than the second highest bid anyway.

Now let us do a numerical example to illustrate.

Bidder	True value	Bid
1	£10	?
2	£7	£7
3	£4	£4
4	£1	£1

Suppose you are bidder 1. This is a situation where you would win the auction if you bid your true value. Would it be better for you to bid less than your true value, say £8? No. It does not matter to your purchasing price whether you bid £10 or £8, since you will still only pay £7.

Can you improve your situation by bidding more than your true value if you think other bidders are not bidding their true value? No. If bidder 2 overbids at £11, and you manage to outbid him or her, then you would have to pay £11 for the item, which is more than you are really willing to pay. If bidder 2 overbids at £9, a £10 bid from you will still win. If bidder 2 underbids at £6, a £10 bid from you will still win. There is no reason for you to overbid, independent of what other participants might decide to do.

Is there ever a reason for you to underbid? No. If bidder 2 overbids at £9, underbidding by you

could result in bidder 2 winning the auction even though you were really willing to pay £10. If bidder 2 underbids at £6, you will only pay £6 whether or not you bid £10 or £7. So why risk losing the auction unnecessarily by underbidding?

Therefore, the best strategy for you is always to bid exactly your true value for the item.

If you are not willing to pay any amount to purchase the item, then you should offer a zero bid.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** works of art for sale. For each work of art, a physical copy of it has been created, as a small print, on high quality card of size A5 (14.8 x 21 cm). Each physical copy is available for sale in the lab today.

Each work of art will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The artwork will be presented on the computer screen.
2. Each bidder is invited to observe the physical copy of the artwork in the laboratory.
3. Each participant submits a bid for the artwork via the computer.
4. The computer ranks bids from highest to lowest (separately for each group of 4 participants). Ties will be broken randomly.

After all the rounds are completed, the computer will randomly select one of the 10 auction rounds as binding (separately for each group of 4 participants). Therefore, **one and only one artwork will be sold in each group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the artwork** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation payment of £12** to determine your final payment. Bids may not exceed £12.

Practice rounds

Before the main auction rounds, two practice rounds will be performed for some hypothetical items that are not available in the laboratory. In the practice rounds (only), you will be assigned a "true value" for the item (in other words, we will ask you to imagine that this amount is your maximum willingness to pay for the item) and you will bid against hypothetical opponents. The practice rounds are solely for familiarizing yourself with the auction procedure and will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

We will now move on to the practice rounds to check your understanding. Please raise your hand if you are stuck or have a question. After everyone in today's session has finished the practice rounds, the main auction rounds will begin.

B.2 Online experiments

This section presents the instructions for the online artwork and poetry mug experiments. To the extent possible, we retained the same instructions for the online studies as were used in the lab experiment, with changes made only where necessary to adapt the protocol to the online setting and to reflect the different item types being auctioned (visual artworks or poetry mugs). Between studies, the instructions were identical except for description specific to the item type being auctioned (see pages 1 and 5 of the instructions below).

For a given item type, the instructions were identical except for the specific treatment phrase that appears on page 5. For the visual artwork experiment, the phrase differs as follows (the language used is identical to our lab experiment): In the *Human Info* treatment, participants read "All artworks being sold are created by human artists." In the *AI Info* treatment, participants read "All artworks being sold are created by an artificial intelligence (AI) program."

For the poetry mug experiment, the treatment phrases differ as follows: In the *Human Info* treatment, participants read "The poetry on all mugs being sold is written by human poets." In the *AI Info* treatment, participants read "The poetry on all mugs being sold is written by an artificial intelligence (AI) program."

The combined sample for the online experiments was provided by Pureprofile [3] and the experiments were implemented using Qualtrics survey software (Qualtrics, Provo, UT). Participants were recruited using a single link which randomly assigned them to one of the four treatments. On entering the survey, participants were asked a set of three questions used for stratification of the sample to match the composition of the sample collected for pre-testing of the items based on age, gender and region.

Considering the online implementation, we employed a number of procedures to acquire high-quality data. First, each participant had to successfully complete a Captcha task before proceeding to the consent form. Second, after providing their consent and before the instructions phase, all participants completed a video attention check, which is particularly effective at screening out AI agents (Celebi et al., 2026). Third, participants completed a classic multiple-choice attention check question, designed to capture human inattention, at a random position in the sequence of auction rounds. Data from any participant who failed either the video or classic attention check questions is not included in the analysis. Fourth, we used Qualtric's built-in feature for bot detection.[4] Based on consultation with the survey provider, data from any participant with a "Q_RecaptchaScore" below 0.5 was excluded. Fifth, we exclude from the analysis the fastest 5%

[3] business.pureprofile.com/.

[4] qualtrics.com/support/survey-platform/survey-module/survey-checker/fraud-detection/.

of responses within each treatment, on the grounds that such respondents may not be paying sufficient attention for their provided data to be reliable. Finally, Pureprofile independently excludes participants they identify as duplicate respondents, and those showing suspicious device behaviour, VPN/proxy usage or automated traffic patterns. These exclusion criteria were pre-registered.

Figure B.1: Instructions page 1 - Artwork experiment.

Study Details

In this study, you can earn money and may purchase an artwork. The amount you earn and whether or not you purchase an artwork depends on your decisions and those of others. You will receive a £12 participation bonus for taking part in the study. You may choose to use none, some, or all of this amount to offer on an artwork. Since the outcome of the study depends also on the decisions of other participants, we will not be able to reveal your earnings, and whether you have purchased an artwork, at the point that you complete the survey. We will need to wait to collect the responses of others and will contact you later to inform you of the outcome (normally, this will be within 48 hours). Your monetary payment will then be processed by PureProfile and any artwork you purchase will be mailed to you.

The study is made up of two parts. In the first part you will take part in a series of auctions for different artworks as described below. In the second part, you will complete a short questionnaire (your responses to the questionnaire will not affect your earnings in the study).

At the end of the survey, due to the binding nature of the auctions and the possibility that you may purchase an artwork, we will collect your name and mailing address. These details will only be used if you are an auction winner and will receive an artwork. If you do not purchase an artwork, your name and address details will not be used and will be securely deleted once study outcomes have been determined.

→

Figure B.2: Instructions page 1 - Poetry mug experiment.

Study Details

In this study, you can earn money and may purchase a poetry mug. The amount you earn and whether or not you purchase a poetry mug depends on your decisions and those of others. You will receive a £12 participation bonus for taking part in the study. You may choose to use none, some, or all of this amount to offer on a poetry mug. Since the outcome of the study depends also on the decisions of other participants, we will not be able to reveal your earnings, and whether you have purchased a poetry mug, at the point that you complete the survey. We will need to wait to collect the responses of others and will contact you later to inform you of the outcome (normally, this will be within 48 hours). Your monetary payment will then be processed by PureProfile and any poetry mug you purchase will be mailed to you.

The study is made up of two parts. In the first part you will take part in a series of auctions for different poetry mugs as described below. In the second part, you will complete a short questionnaire (your responses to the questionnaire will not affect your earnings in the study).

At the end of the survey, due to the binding nature of the auctions and the possibility that you may purchase a poetry mug, we will collect your name and mailing address. These details will only be used if you are an auction winner and will receive a poetry mug. If you do not purchase a poetry mug, your name and address details will not be used and will be securely deleted once study outcomes have been determined.

→

Figure B.3: Instructions page 2 (both experiments).

Instructions for Part I

We will now describe the task which makes up the first part of the study.

You will be randomly assigned to a group of 4 participants (including you). You will not know the identity of any other person in your group and no-one else in your group will know your identity.

You and the other three people in your group will take part in a type of auction known as the "second-price auction". In an auction procedure, everyone in the group submits a bid for an item for purchase. The person submitting the highest bid in the group buys the item. Unlike in other types of auctions that you may have encountered before, in the second-price auction, the person submitting the highest bid **purchases the item at a price equal to the next highest bid, which has been submitted by another participant in the group**.

→

Figure B.4: Instructions page 3 (both experiments).

Why are we using a "second-price auction" procedure?

The best strategy in this type of auction is to always bid exactly the maximum amount that you would be willing to pay for the item.

To understand why, let us assume that you have decided that you would be willing to pay up to £10 for the item. Let us call this your true value of the item. Which strategy is best for you: to bid higher than £10, lower than £10, or to bid exactly £10? The answer is: to bid exactly £10.

If you were to bid less than £10 (say £8) the auction might be won by a higher bidder whose bid is also less than £10 (say £9), and you will not be able to buy the item even though you were truly willing to pay more than the winner's bid (which is the amount you would have paid if you had bid higher). Therefore, bidding less than your true valuation is not a good strategy.

If you were to bid more than £10 (say £12) you might end up winning the auction but having to pay a price higher than what you really value the item at (say if the second highest bid was £11). Therefore, bidding more than your true valuation is not a good strategy.

Is there a point in trying to guess what the next highest bid after yours is, in order to bid just a penny more than that person? No, because the way the purchase price is determined in this auction you will only have to pay the price of the next highest bid anyway, independent of how much higher your bid was.

The best strategy in this type of auction is therefore to **always bid exactly the maximum amount that you decided that you would be willing to pay**. You do not risk paying more for the item than what you want to, and you will never pay more than the second highest bid anyway.

Figure B.5: Instructions page 4 (both experiments).

Now let us do a numerical example to illustrate.

Bidder	True value	Bid
1	£10	?
2	£7	£7
3	£4	£4
4	£1	£1

Suppose you are bidder 1. This is a situation where you would win the auction if you bid your true value. Would it be better for you to bid less than your true value, say £8? No. It does not matter to your purchasing price whether you bid £10 or £8, since you will still only pay £7.

Can you improve your situation by bidding more than your true value if you think other bidders are not bidding their true value? No. If bidder 2 overbids at £11, and you manage to outbid him or her, then you would have to pay £11 for the item, which is more than you are really willing to pay. If bidder 2 overbids at £9, a £10 bid from you will still win. If bidder 2 underbids at £6, a £10 bid from you will still win. There is no reason for you to overbid, independent of what other participants might decide to do.

Is there ever a reason for you to underbid? No. If bidder 2 overbids at £9, underbidding by you could result in bidder 2 winning the auction even though you were really willing to pay £10. If bidder 2 underbids at £6, you will only pay £6 whether or not you bid £10 or £7. So why risk losing the auction unnecessarily by underbidding?

Therefore, the best strategy for you is always to bid exactly your true value for the item.

If you are not willing to pay any amount to purchase the item, then you should offer a zero bid.

Figure B.6: Instructions page 5 - Artwork experiment - Human Info treatment.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** artworks. For each artwork, a physical copy of it has been made, as a small print, on high quality card of size A5 (14.8 x 21 cm). Each physical copy is available for sale.

All artworks being sold are created by human artists.

Each artwork will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The artwork will be presented on the computer screen.
2. Each participant submits a bid for the artwork via the computer.
3. The computer ranks bids from highest to lowest in your group. Ties will be broken randomly.

Note that the other bidders are real participants in this study. Since they may not be entering their responses at exactly the same time as you, the auctions are not taking place in real time. However, no information will be revealed to other participants about your bids, nor will you receive any information about other participants' bids, until all participants have made their bids in all 10 auction rounds.

After all the rounds are completed by all participants, the computer will randomly select one of the 10 auction rounds as binding. Therefore, **one and only one artwork will be sold in your group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the artwork** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation bonus of £12** to determine your final payment. Bids may not exceed £12.

Figure B.7: Instructions page 5 - Artwork experiment - AI Info treatment.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** artworks. For each artwork, a physical copy of it has been made, as a small print, on high quality card of size A5 (14.8 x 21 cm). Each physical copy is available for sale.

All artworks being sold are created by an artificial intelligence (AI) program.

Each artwork will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The artwork will be presented on the computer screen.
2. Each participant submits a bid for the artwork via the computer.
3. The computer ranks bids from highest to lowest in your group. Ties will be broken randomly.

Note that the other bidders are real participants in this study. Since they may not be entering their responses at exactly the same time as you, the auctions are not taking place in real time. However, no information will be revealed to other participants about your bids, nor will you receive any information about other participants' bids, until all participants have made their bids in all 10 auction rounds.

After all the rounds are completed by all participants, the computer will randomly select one of the 10 auction rounds as binding. Therefore, **one and only one artwork will be sold in your group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the artwork** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation bonus of £12** to determine your final payment. Bids may not exceed £12.

Figure B.8: Instructions page 5 - Poetry mug experiment - Human Info treatment.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** pictures of poetry mugs. Each is a standard-size white ceramic mug, with poetry printed on it. Each poetry mug is a physical item available for sale.

The poetry on all mugs being sold is written by human poets.

Each poetry mug will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The poetry mug will be presented on the computer screen.
2. Each participant submits a bid for the poetry mug via the computer.
3. The computer ranks bids from highest to lowest in your group. Ties will be broken randomly.

Note that the other bidders are real participants in this study. Since they may not be entering their responses at exactly the same time as you, the auctions are not taking place in real time. However, no information will be revealed to other participants about your bids, nor will you receive any information about other participants' bids, until all participants have made their bids in all 10 auction rounds.

After all the rounds are completed by all participants, the computer will randomly select one of the 10 auction rounds as binding. Therefore, **one and only one poetry mug will be sold in your group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the poetry mug** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation bonus of £12** to determine your final payment. Bids may not exceed £12.

Figure B.9: Instructions page 5 - Poetry mug experiment - AI Info treatment.

What are the items on which you are bidding?

You and the other three participants in your group will be presented with a sequence of **10** pictures of poetry mugs. Each is a standard-size white ceramic mug, with a poetry printed on it. Each poetry mug is a physical item available for sale.

The poetry on all mugs being sold is written by an artificial intelligence (AI) program.

Each poetry mug will be sold in a separate auction and you will participate in a sequence of 10 auction rounds. Each of the 10 auction rounds will proceed in the following steps:

1. The poetry mug will be presented on the computer screen.
2. Each participant submits a bid for the poetry mug via the computer.
3. The computer ranks bids from highest to lowest in your group. Ties will be broken randomly.

Note that the other bidders are real participants in this study. Since they may not be entering their responses at exactly the same time as you, the auctions are not taking place in real time. However, no information will be revealed to other participants about your bids, nor will you receive any information about other participants' bids, until all participants have made their bids in all 10 auction rounds.

After all the rounds are completed by all participants, the computer will randomly select one of the 10 auction rounds as binding. Therefore, **one and only one poetry mug will be sold in your group (to the highest bidder in the randomly selected round)**. This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round (as you will only purchase at most one item).

Note! Depending on yours and the other participants' bids, you may or may not be the highest bidder in your group in the selected round. If you are the highest bidder, then you will **buy the poetry mug** at a price equal to the second highest bid in the group. This price will be **subtracted from your participation bonus of £12** to determine your final payment. Bids may not exceed £12.

C Experiment auction program screenshots

C.1 Lab experiment

Figure C.1: Human Info treatment - Example Lot Information Screen.

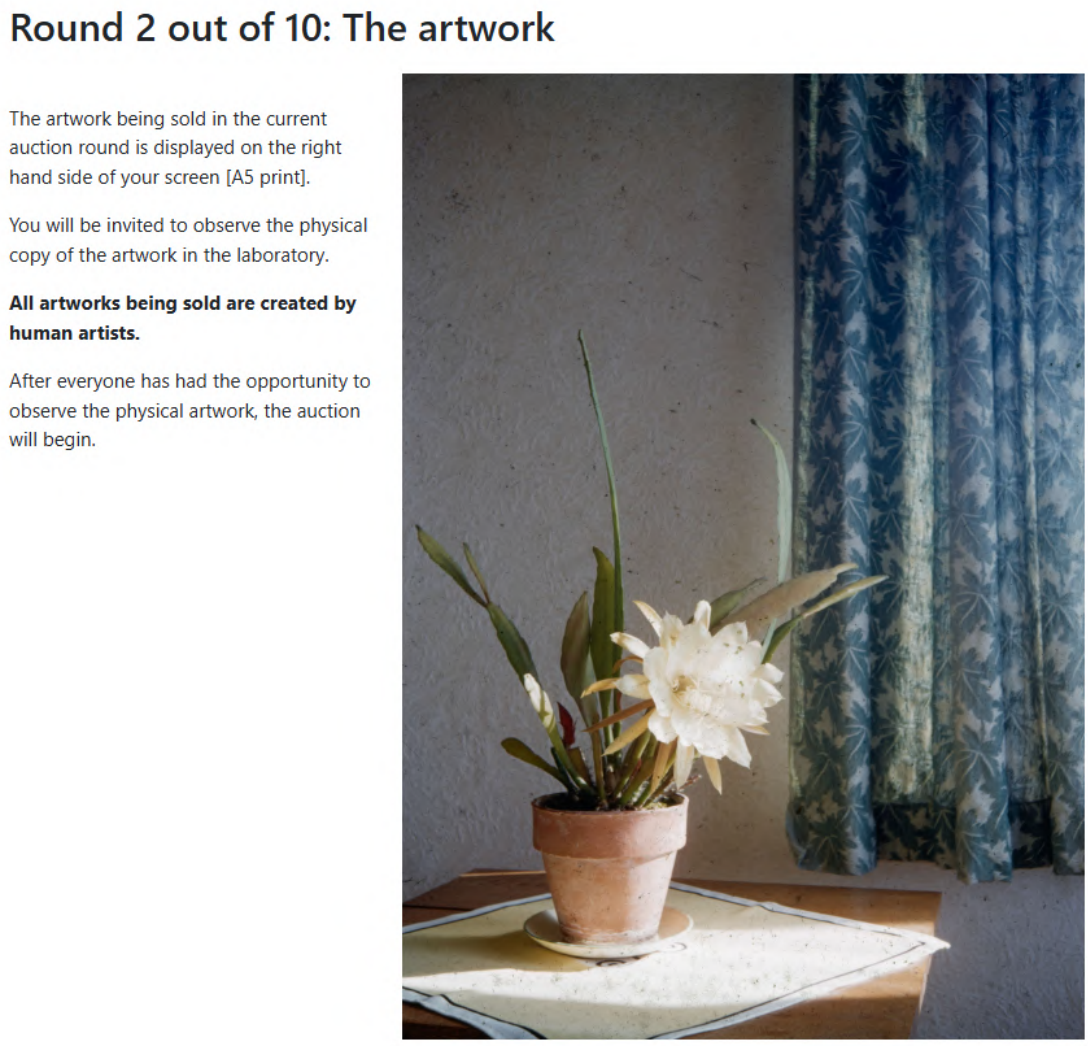


Figure C.2: Human Info treatment - Example Bid Screen.

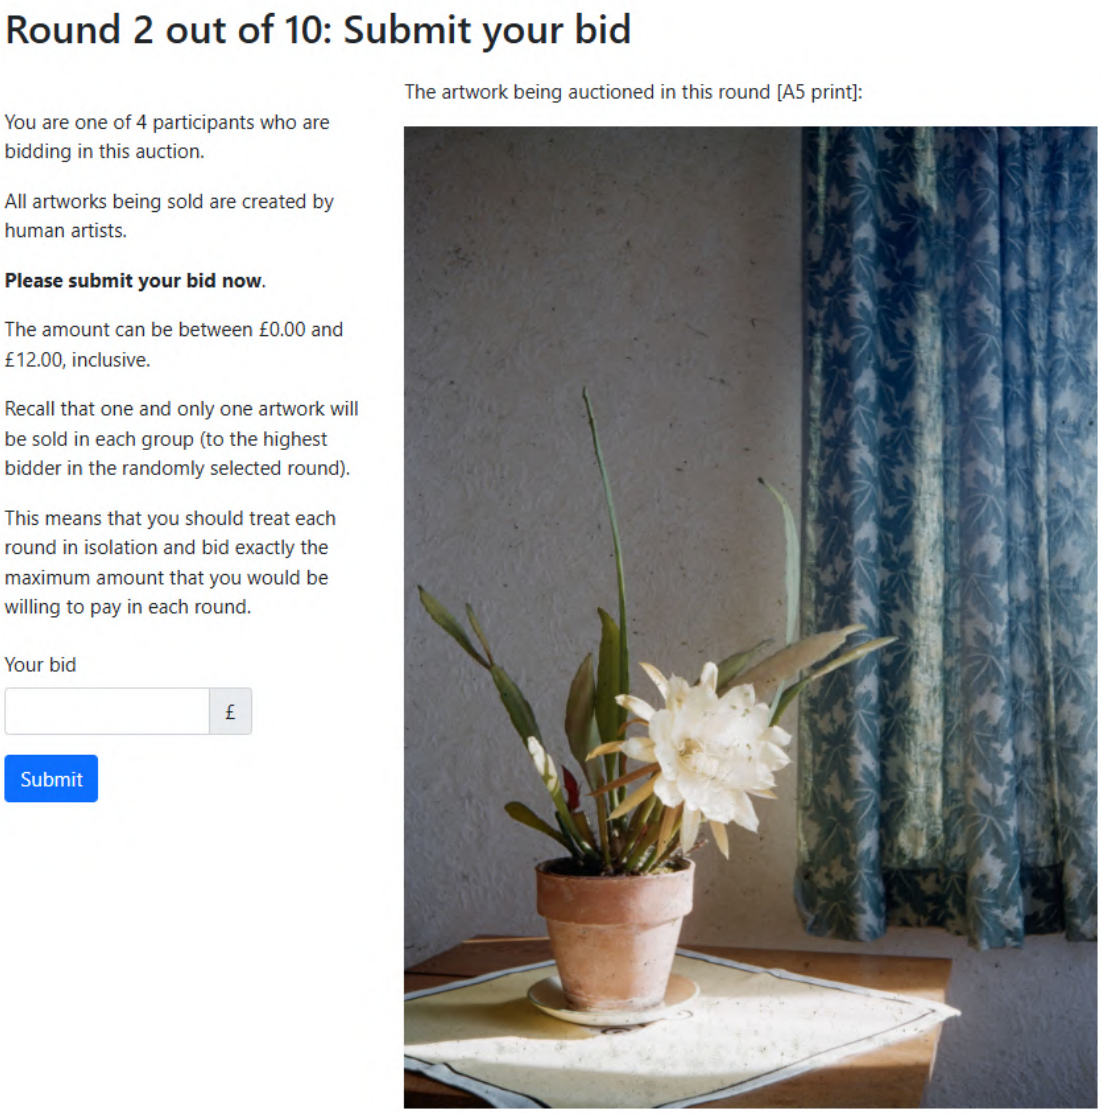


Figure C.3: AI Info treatment - Example Lot Information Screen.

Round 2 out of 10: The artwork

The artwork being sold in the current auction round is displayed on the right hand side of your screen [A5 print].

You will be invited to observe the physical copy of the artwork in the laboratory.

All artworks being sold are created by an artificial intelligence (AI) program.

After everyone has had the opportunity to observe the physical artwork, the auction will begin.



Figure C.4: AI Info treatment - Example Bid Screen.

Round 2 out of 10: Submit your bid

You are one of 4 participants who are bidding in this auction.

All artworks being sold are created by an artificial intelligence (AI) program.

Please submit your bid now.

The amount can be between £0.00 and £12.00, inclusive.

Recall that one and only one artwork will be sold in each group (to the highest bidder in the randomly selected round).

This means that you should treat each round in isolation and bid exactly the maximum amount that you would be willing to pay in each round.

Your bid

£

Submit

The artwork being auctioned in this round [A5 print]:



Figure C.5: Human No Info treatment - Example Lot Information Screen.

Round 2 out of 10: The artwork

The artwork being sold in the current auction round is displayed on the right hand side of your screen [A5 print].

You will be invited to observe the physical copy of the artwork in the laboratory.

After everyone has had the opportunity to observe the physical artwork, the auction will begin.



Figure C.6: Human No Info treatment - Example Bid Screen.

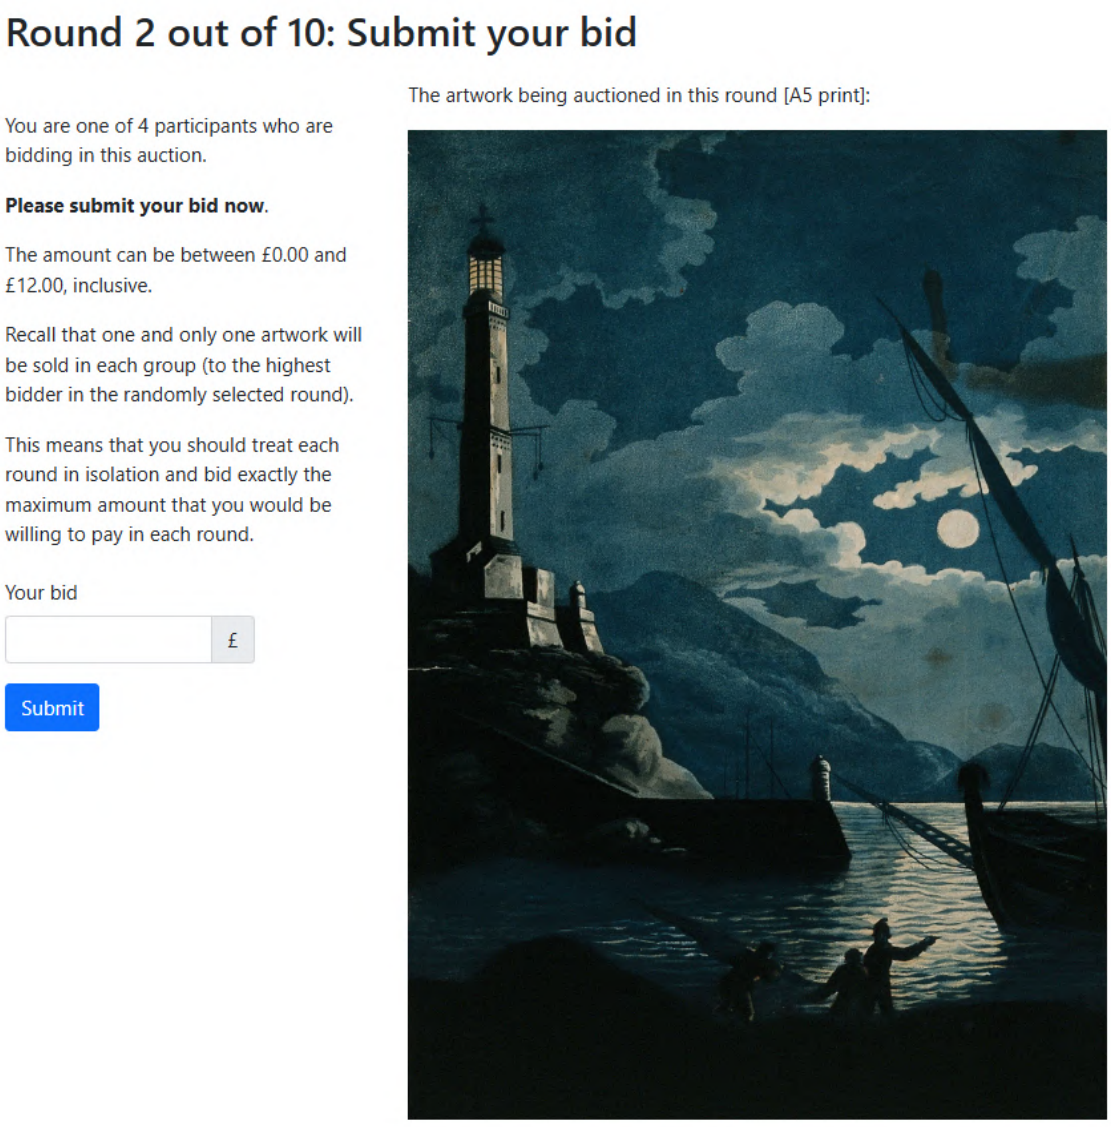


Figure C.7: AI No Info treatment - Example Lot Information Screen.

Round 2 out of 10: The artwork

The artwork being sold in the current auction round is displayed on the right hand side of your screen [A5 print].

You will be invited to observe the physical copy of the artwork in the laboratory.

After everyone has had the opportunity to observe the physical artwork, the auction will begin.



Figure C.8: AI No Info treatment - Example Bid Screen.

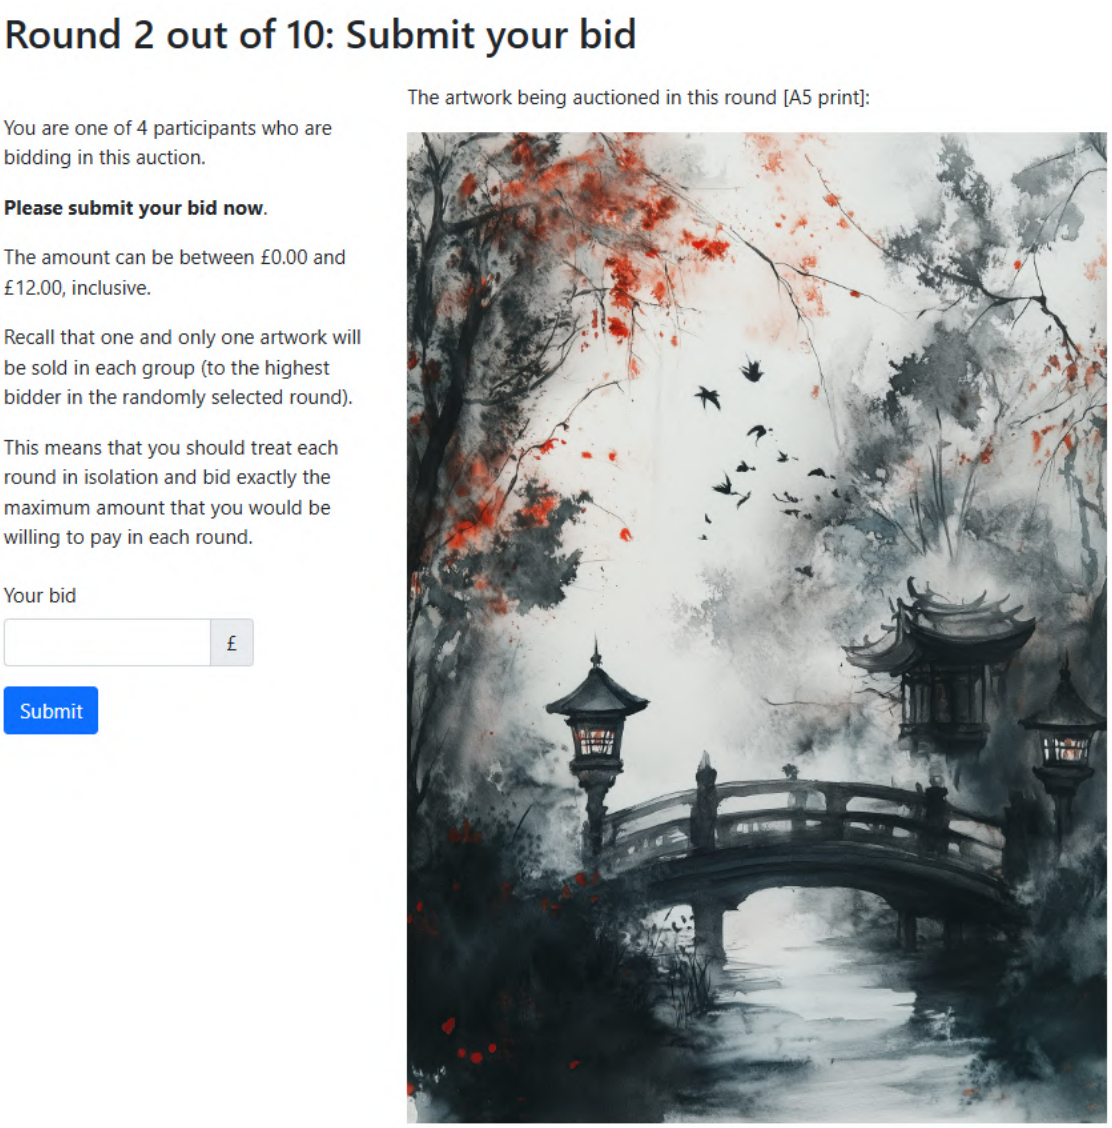


Figure C.9: Practice Round 1 - Information Screen.

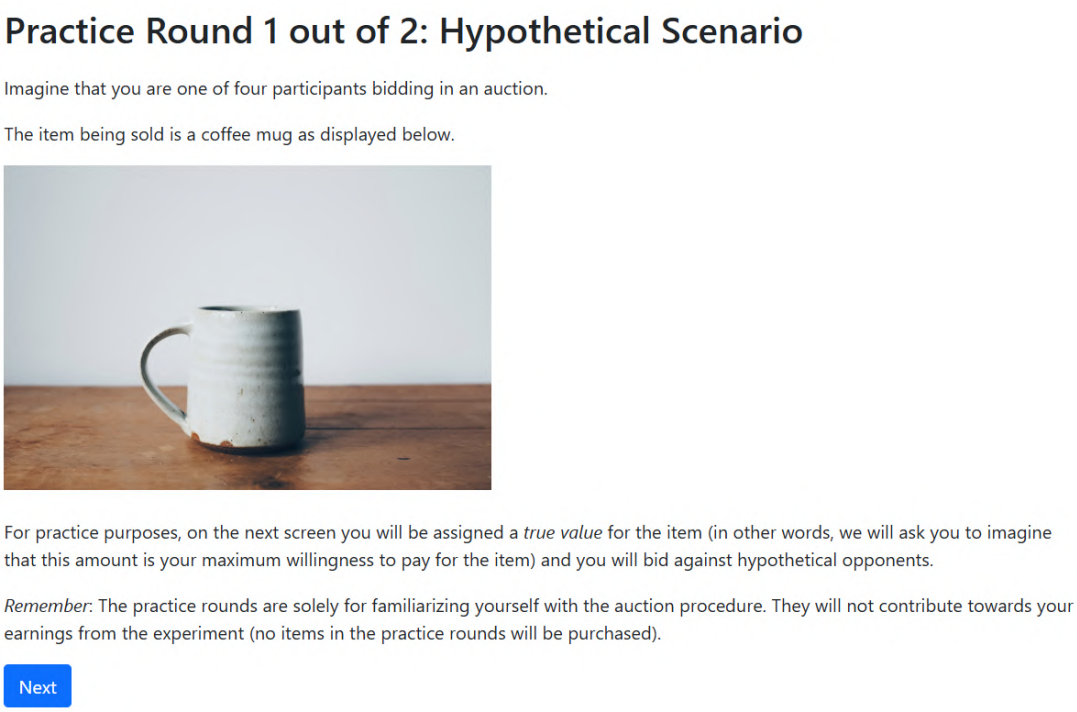
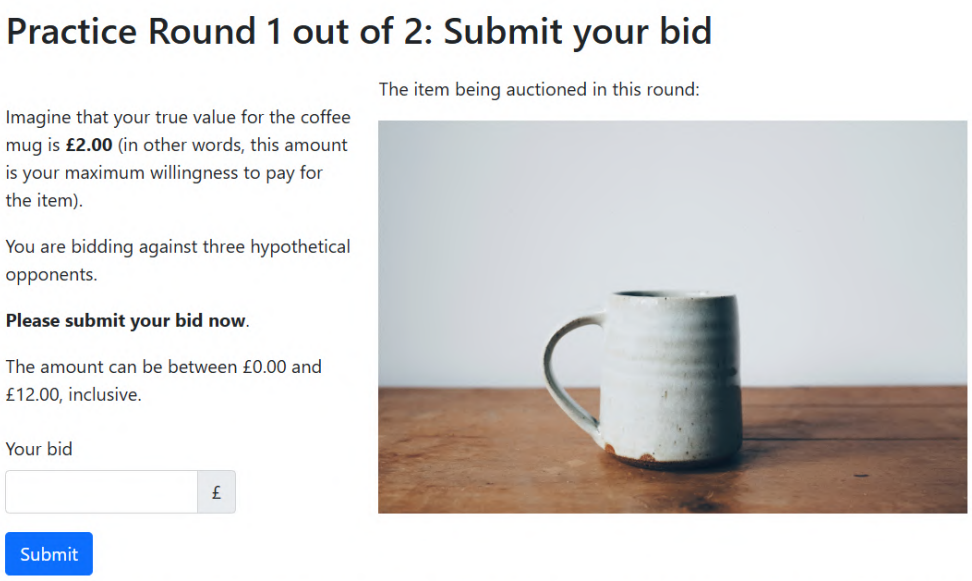
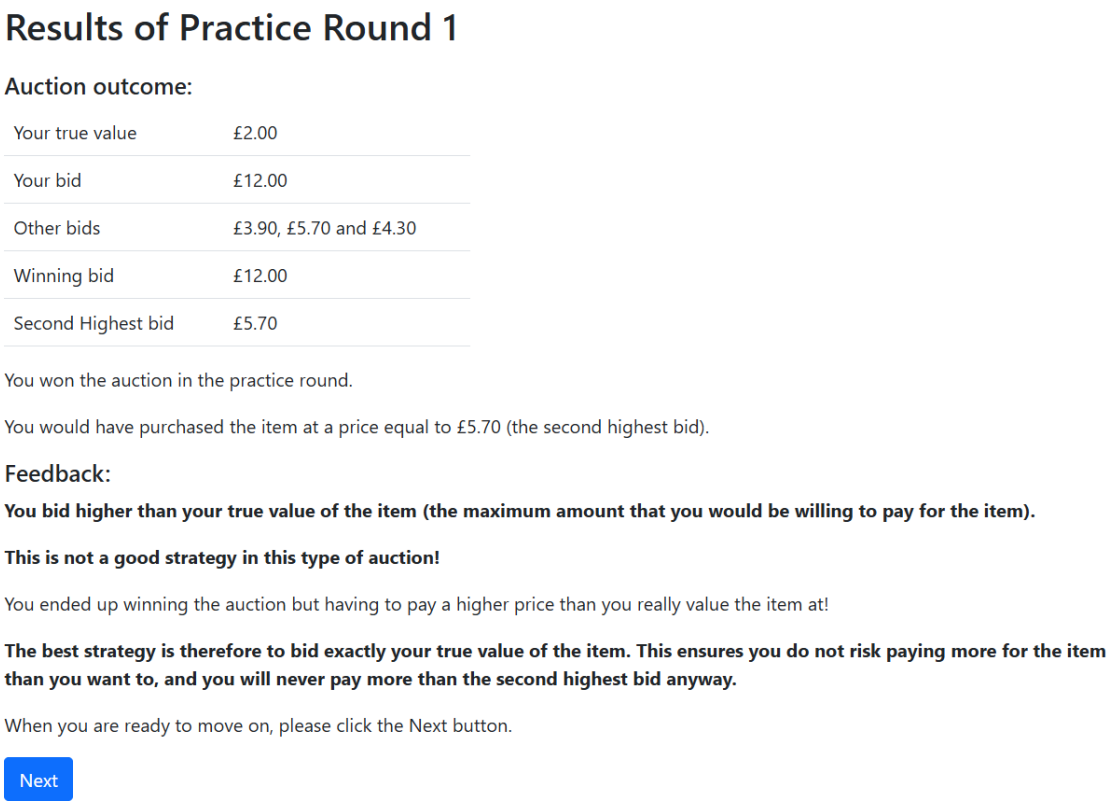


Figure C.10: Practice Round 1 - Example Bid Screen.



Notes: Induced "true value" is randomized between subjects.

Figure C.11: Practice Round 1 - Example Feedback Screen.



Results of Practice Round 1

Auction outcome:

Your true value	£2.00
Your bid	£12.00
Other bids	£3.90, £5.70 and £4.30
Winning bid	£12.00
Second Highest bid	£5.70

You won the auction in the practice round.

You would have purchased the item at a price equal to £5.70 (the second highest bid).

Feedback:

You bid higher than your true value of the item (the maximum amount that you would be willing to pay for the item).

This is not a good strategy in this type of auction!

You ended up winning the auction but having to pay a higher price than you really value the item at!

The best strategy is therefore to bid exactly your true value of the item. This ensures you do not risk paying more for the item than you want to, and you will never pay more than the second highest bid anyway.

When you are ready to move on, please click the Next button.

Next

Figure C.12: Practice Round 2 - Information Screen.

Practice Round 2 out of 2: Hypothetical Scenario

Imagine again that you are one of four participants bidding in an auction.

The item being sold is a tape measure as displayed below.



For practice purposes, on the next screen you will be assigned a *true value* for the item (in other words, we will ask you to imagine that this amount is your maximum willingness to pay for the item) and you will bid against hypothetical opponents.

Remember: The practice rounds are solely for familiarizing yourself with the auction procedure. They will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

Next

Figure C.13: Practice Round 2 - Example Bid Screen.

Practice Round 2 out of 2: Submit your bid

Imagine that your true value for the tape measure is **£11.00** (in other words, this amount is your maximum willingness to pay for the item).

You are bidding against three hypothetical opponents.

Please submit your bid now.

The amount can be between £0.00 and £12.00, inclusive.

The item being auctioned in this round:



Your bid

£

Submit

Remember: The practice rounds are solely for familiarizing yourself with the auction procedure. They will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

Notes: Induced "true value" is randomized between subjects.

Figure C.14: Practice Round 2 - Example Feedback Screen.

Results of Practice Round 2

Auction outcome:

Your true value	£11.00
Your bid	£1.00
Other bids	£7.40, £1.40 and £10.00
Winning bid	£10.00
Second Highest bid	£7.40

You did not win the auction in the practice round.

The item was purchased by the highest bidder at a price equal to £7.40 (the second highest bid).

Feedback:

You bid lower than your true value of the item (the maximum amount that you would be willing to pay for the item).

This is not a good strategy in this type of auction!

The auction was won by a higher bidder whose bid was also less than your true value (but higher than your bid).

This left you unable to buy the item even though you were truly willing to pay more than the winner's bid (the amount you would have paid if you had bid higher).

The best strategy is therefore to bid exactly your true value of the item. This ensures you do not risk paying more for the item than you want to, and you will never pay more than the second highest bid anyway.

When you are ready to move on, please click the Next button.

Next

C.2 Online experiments

Figure C.15: Artwork experiment - Human Info treatment - Example Lot Information Screen.



Figure C.16: Artwork experiment - Human Info treatment - Example Bid Screen.

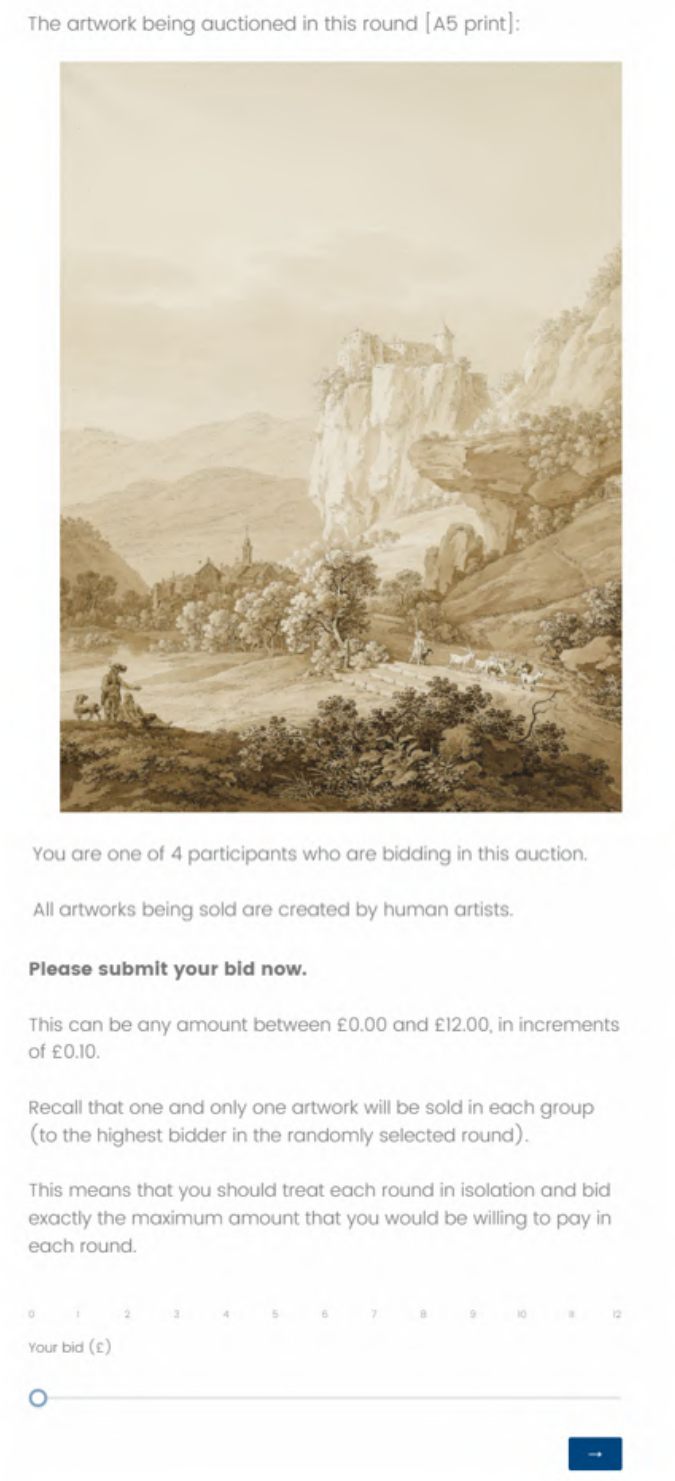


Figure C.17: Artwork experiment - AI Info treatment - Example Lot Information Screen.



Figure C.18: Artwork experiment - AI Info treatment - Example Bid Screen.

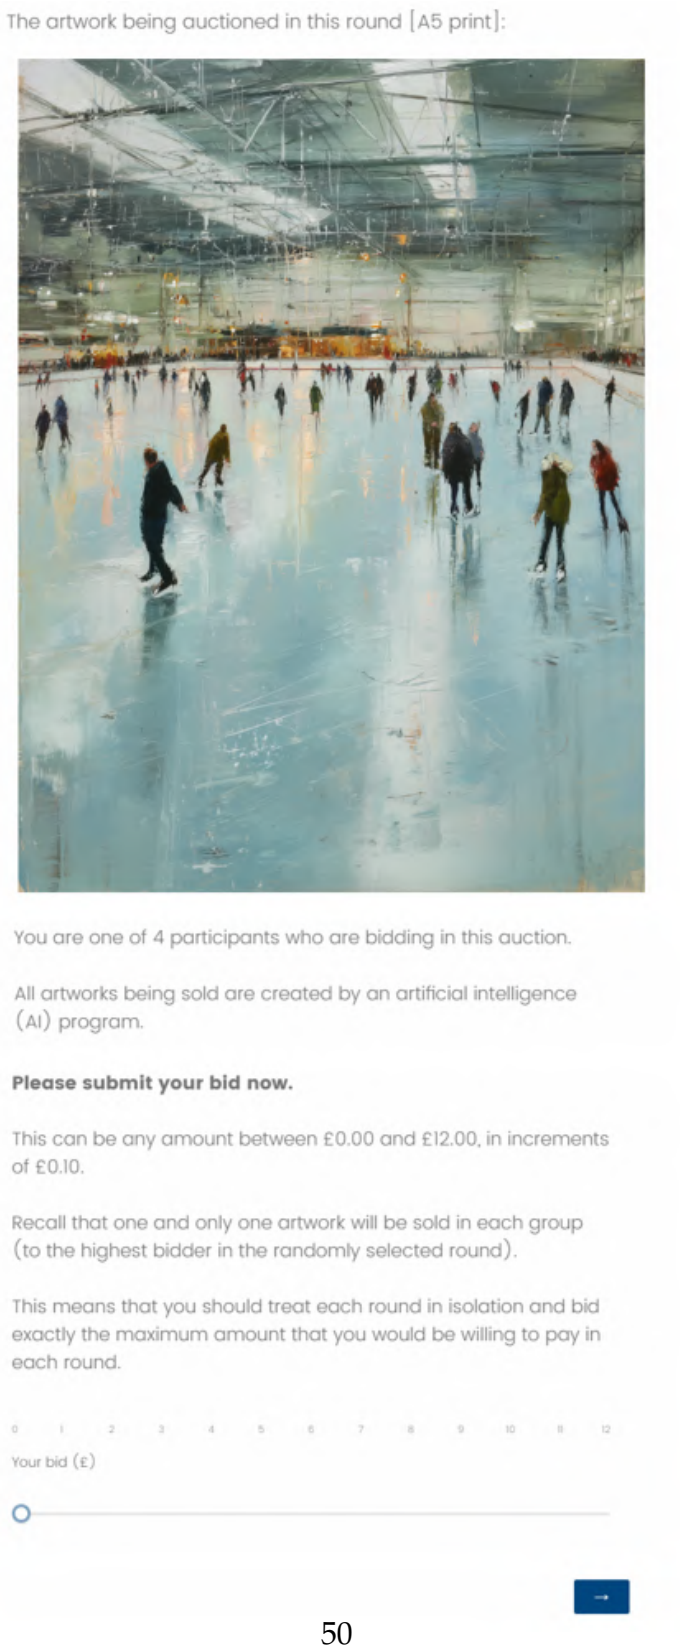


Figure C.19: Poetry mug experiment - Human Info treatment - Example Lot Information Screen.

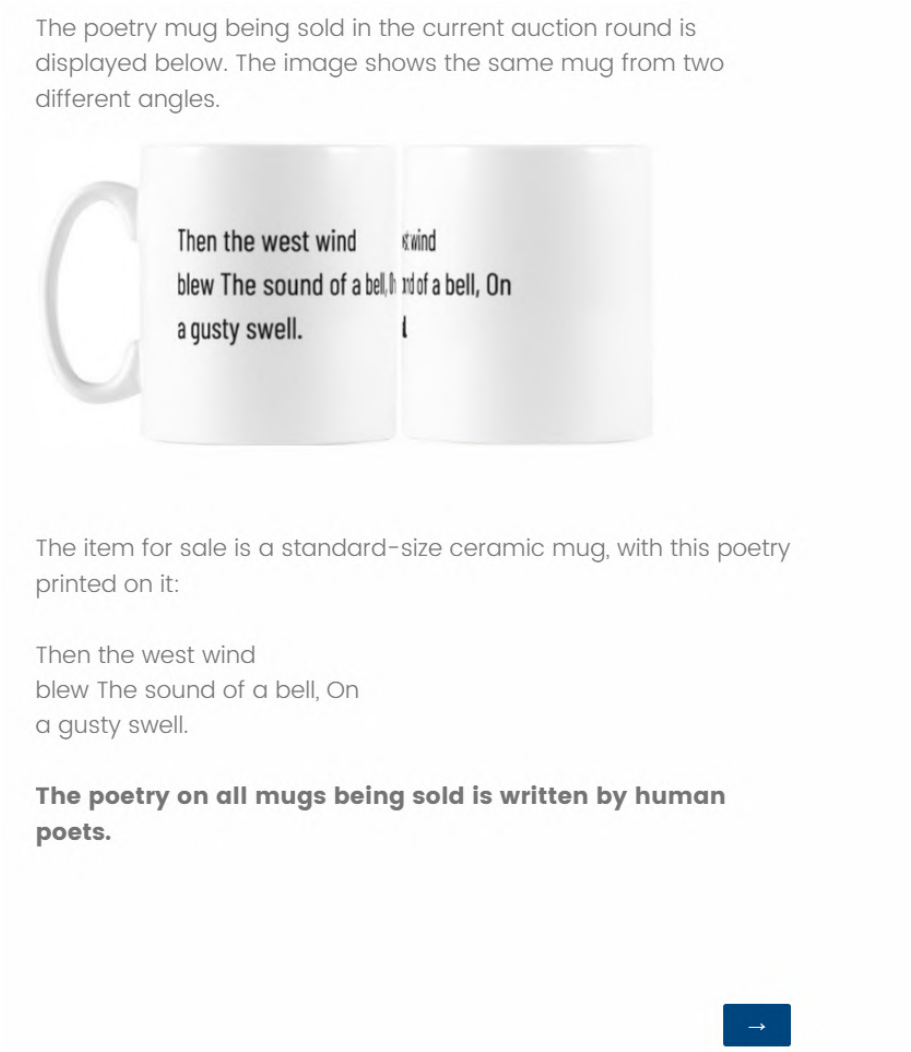


Figure C.20: Poetry mug experiment - Human Info treatment - Example Bid Screen.

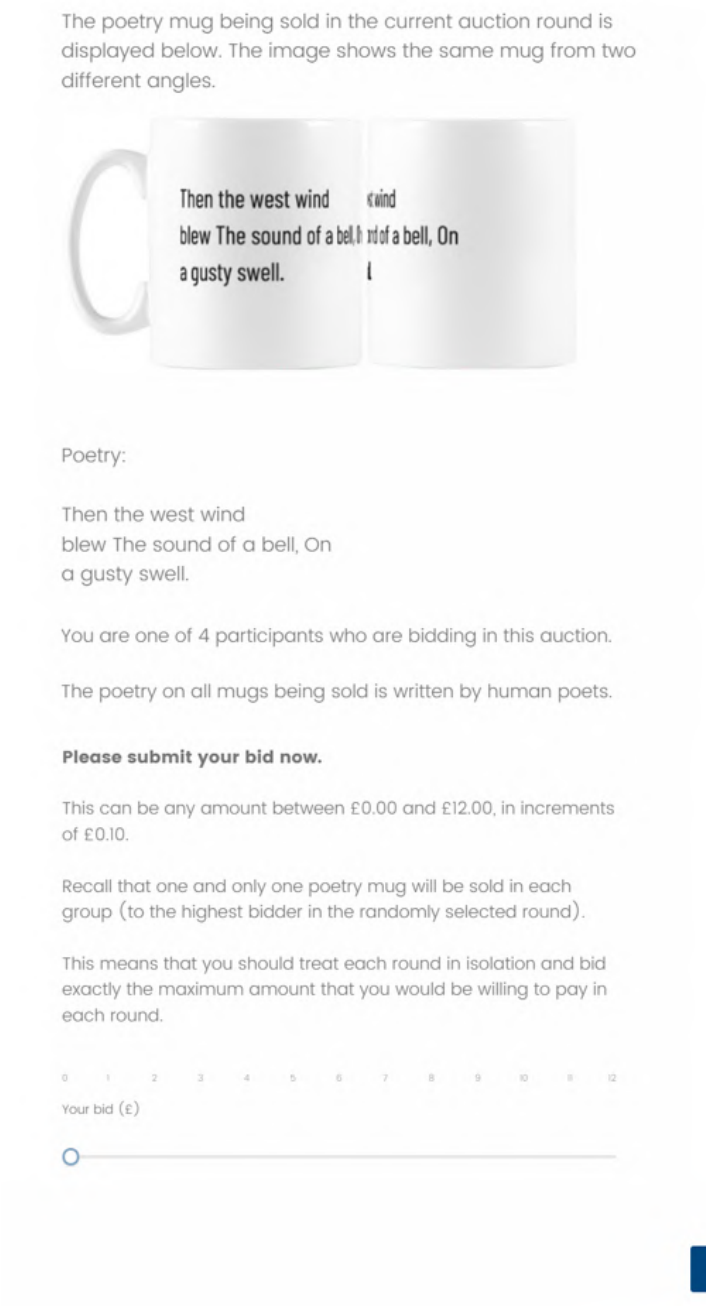


Figure C.21: Poetry mug experiment - AI Info treatment - Example Lot Information Screen.

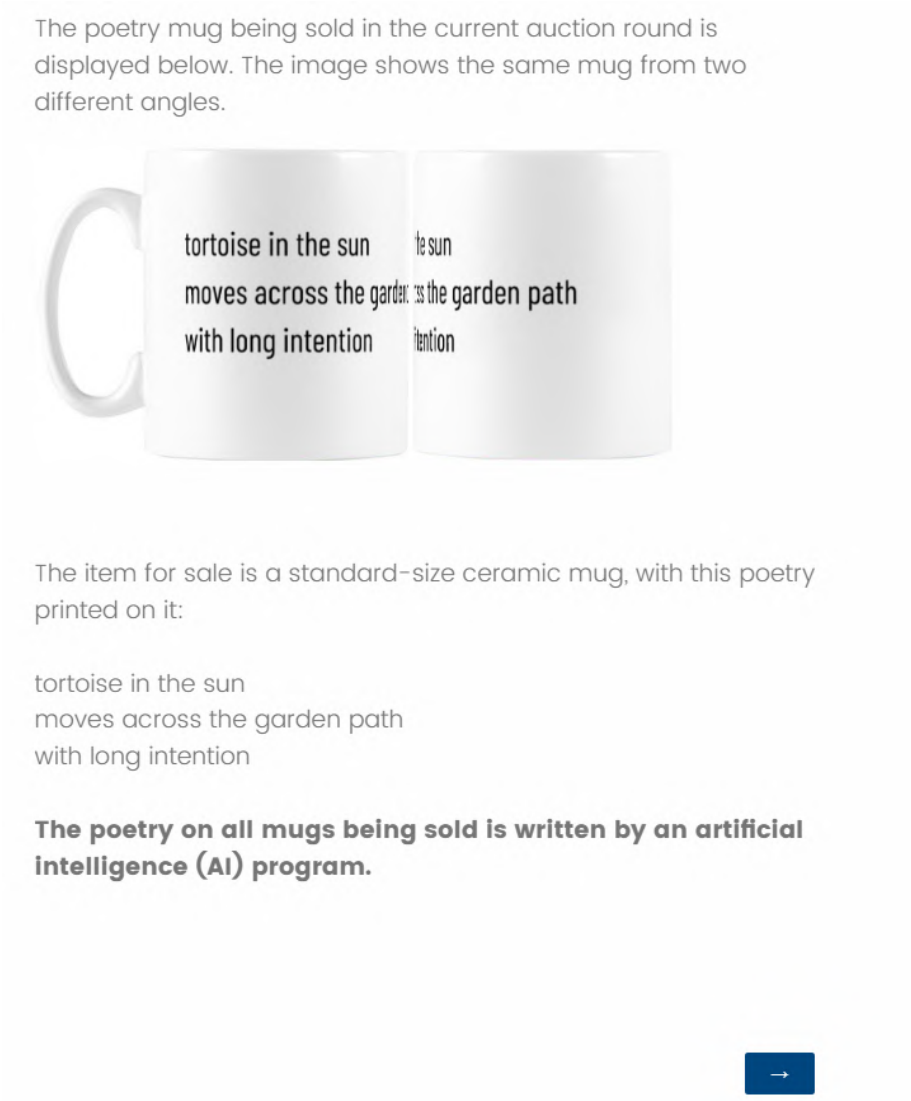


Figure C.22: Poetry mug experiment - AI Info treatment - Example Bid Screen.

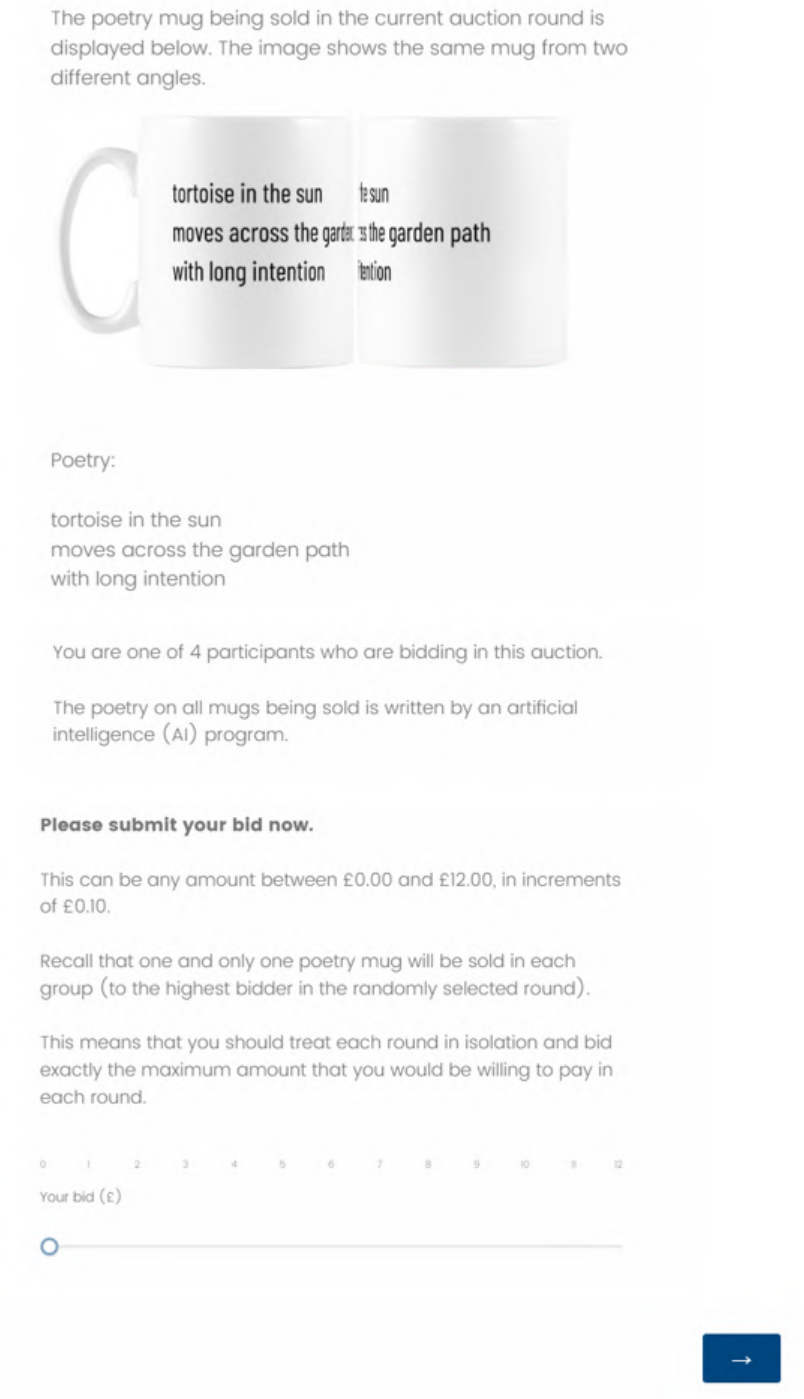


Figure C.23: Practice Round 1 - Information Screen.

Practice Round 1 out of 2: Hypothetical Scenario

Imagine that you are one of four participants bidding in an auction.

The item being sold is a pen as displayed below.



For practice purposes, on the next screen you will be assigned a *true value* for the item (in other words, we will ask you to imagine that this amount is your maximum willingness to pay for the item) and you will bid against hypothetical opponents.

Remember: The practice rounds are solely for familiarizing yourself with the auction procedure. They will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

→

Figure C.24: Practice Round 1 - Example Bid Screen.

Practice Round 1 out of 2: Submit your bid

Imagine that your true value for the pen is **£4** (in other words, this amount is your maximum willingness to pay for the item).

You are bidding against three hypothetical opponents.

The item being auctioned in this round:



Please submit your bid now.

This can be any amount between £0.00 and £12.00, in increments of £0.10.

0 1 2 3 4 5 6 7 8 9 10 11 12

Your bid (£)

Remember: The practice rounds are solely for familiarizing yourself with the auction procedure. They will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

→

Figure C.25: Practice Round 1 - Example Feedback Screen.

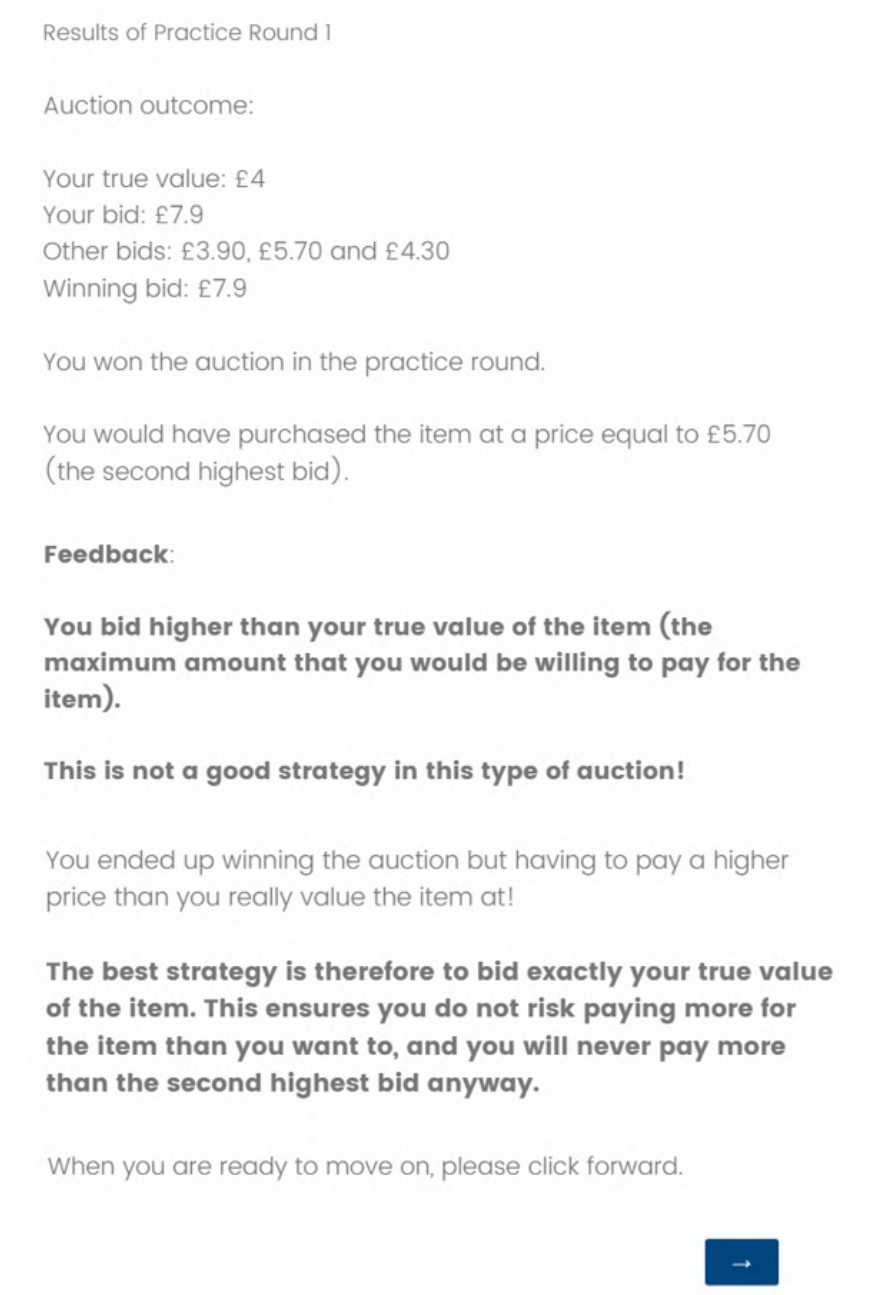


Figure C.26: Practice Round 2 - Information Screen.

Practice Round 2 out of 2: Hypothetical Scenario

Imagine again that you are one of four participants bidding in an auction.

The item being sold is a tape measure as displayed below.



For practice purposes, on the next screen you will be assigned a *true value* for the item (in other words, we will ask you to imagine that this amount is your maximum willingness to pay for the item) and you will bid against hypothetical opponents.

Remember: The practice rounds are solely for familiarizing yourself with the auction procedure. They will not contribute towards your earnings from the experiment (no items in the practice rounds will be purchased).

→

Figure C.27: Practice Round 2 - Example Bid Screen.

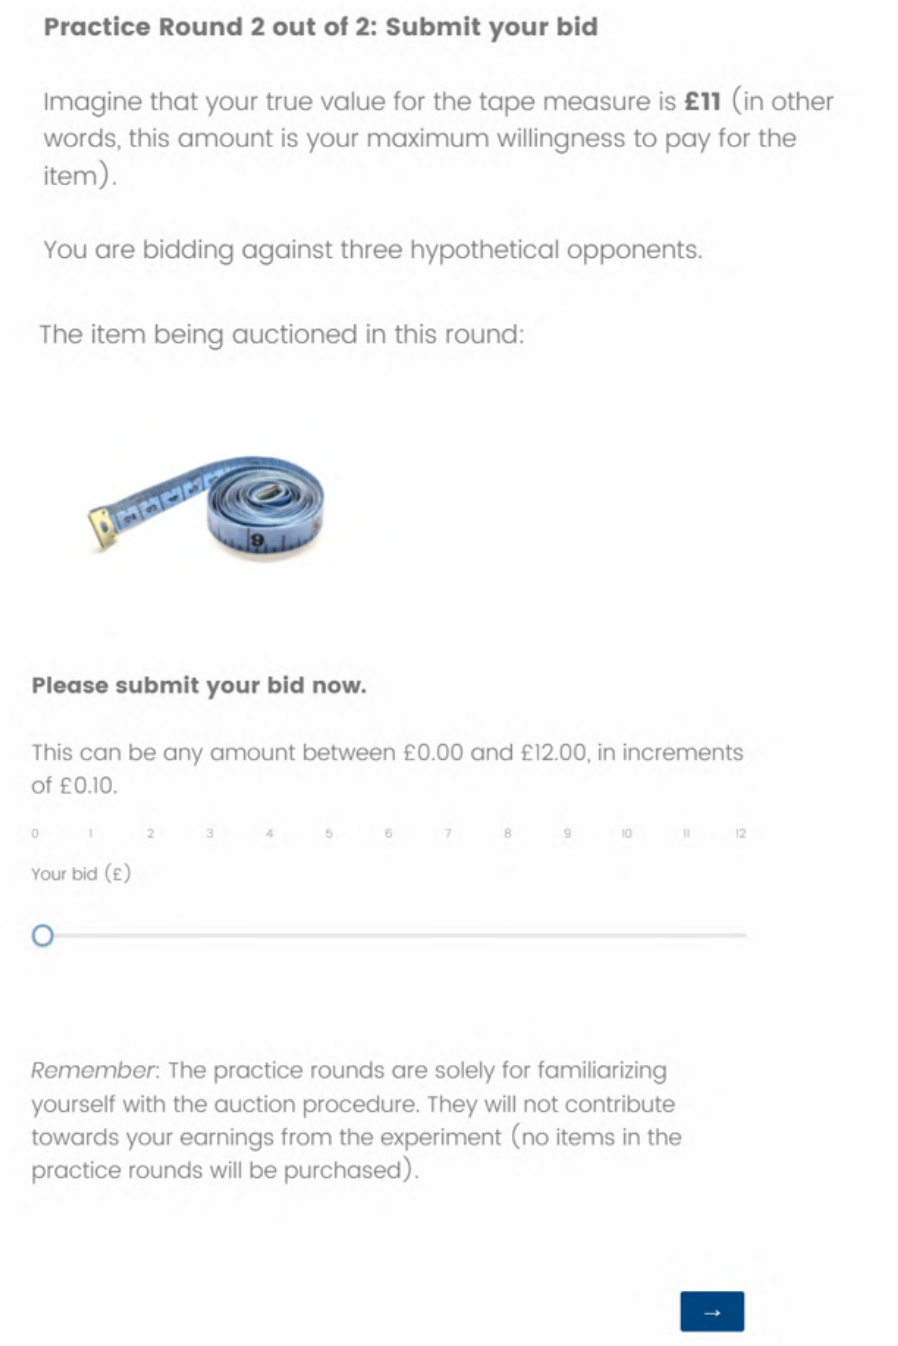
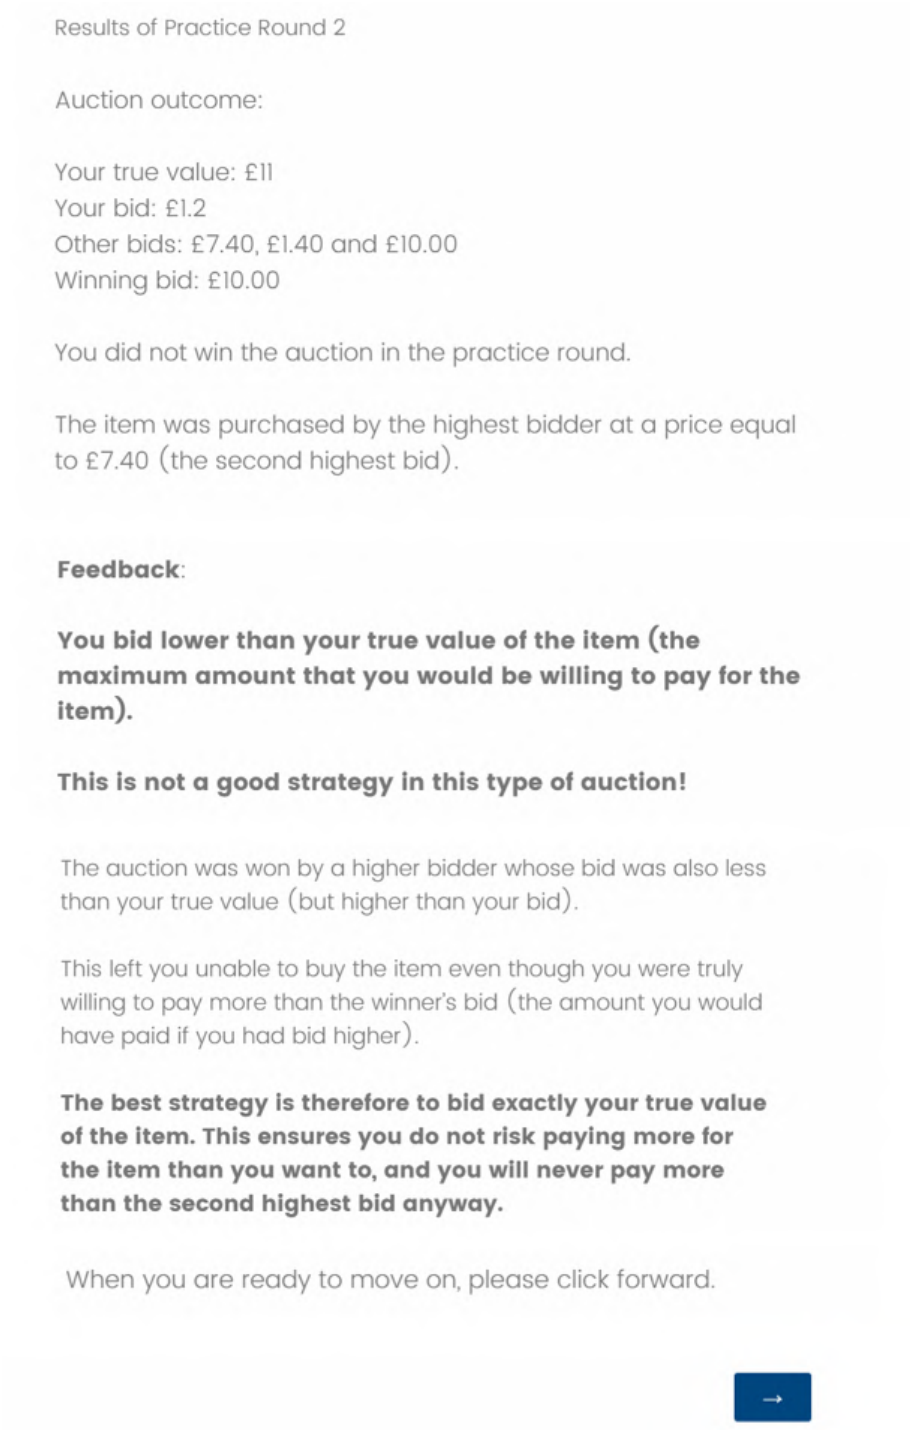


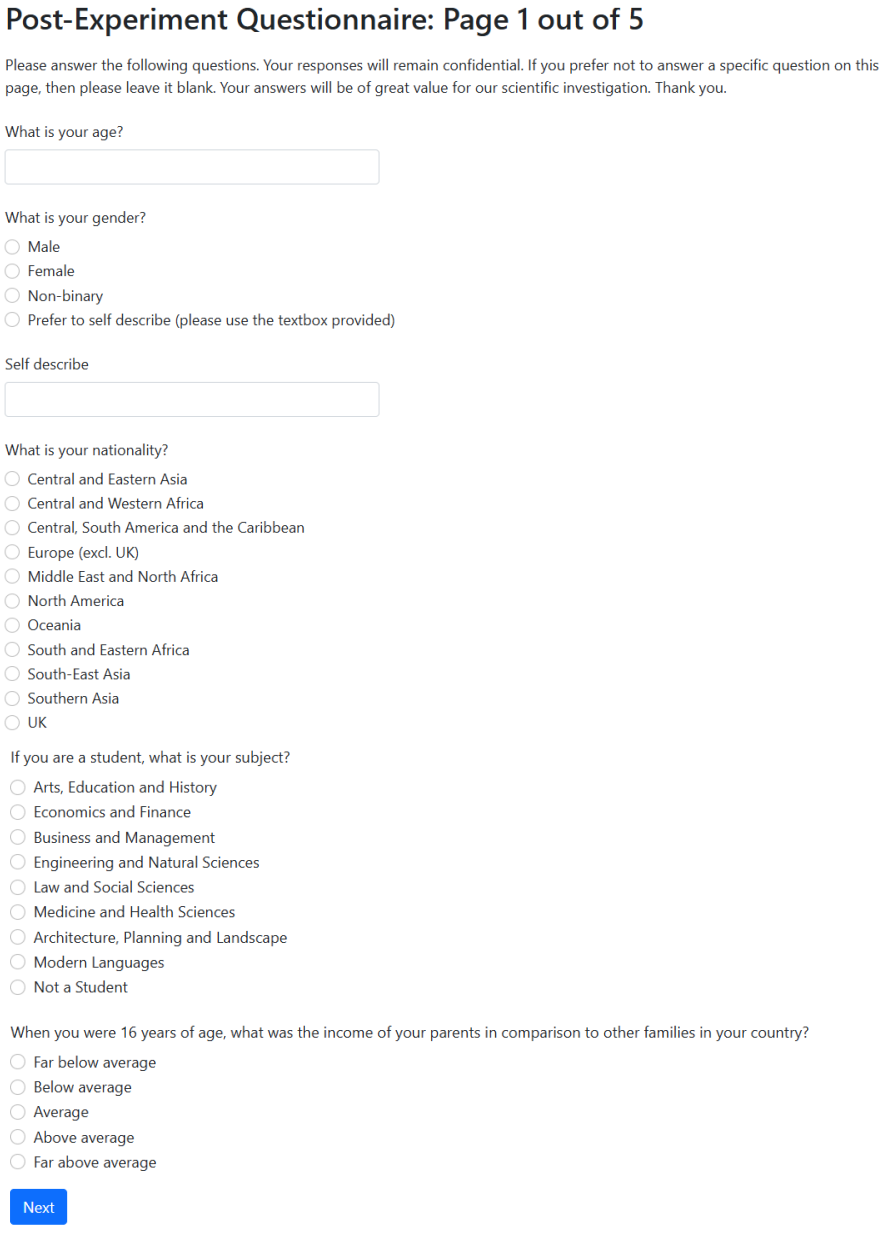
Figure C.28: Practice Round 2 - Example Feedback Screen.



D Post-experiment questionnaire

D.1 Lab experiment

Figure D.1: Socio-demographic characteristics.



Post-Experiment Questionnaire: Page 1 out of 5

Please answer the following questions. Your responses will remain confidential. If you prefer not to answer a specific question on this page, then please leave it blank. Your answers will be of great value for our scientific investigation. Thank you.

What is your age?

What is your gender?

- Male
- Female
- Non-binary
- Prefer to self describe (please use the textbox provided)

Self describe

What is your nationality?

- Central and Eastern Asia
- Central and Western Africa
- Central, South America and the Caribbean
- Europe (excl. UK)
- Middle East and North Africa
- North America
- Oceania
- South and Eastern Africa
- South-East Asia
- Southern Asia
- UK

If you are a student, what is your subject?

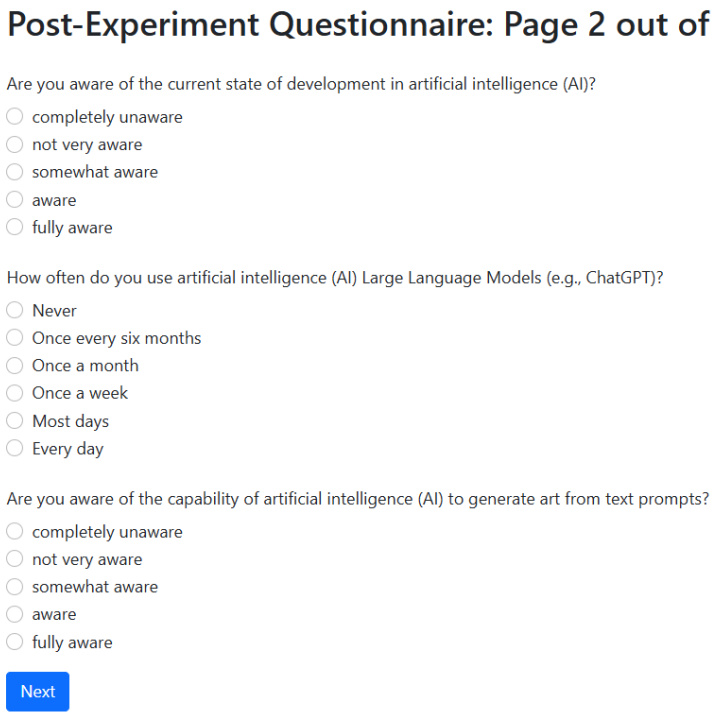
- Arts, Education and History
- Economics and Finance
- Business and Management
- Engineering and Natural Sciences
- Law and Social Sciences
- Medicine and Health Sciences
- Architecture, Planning and Landscape
- Modern Languages
- Not a Student

When you were 16 years of age, what was the income of your parents in comparison to other families in your country?

- Far below average
- Below average
- Average
- Above average
- Far above average

Next

Figure D.2: Awareness, use and knowledge of AI.



Post-Experiment Questionnaire: Page 2 out of 5

Are you aware of the current state of development in artificial intelligence (AI)?

- ○ completely unaware
- ○ not very aware
- ○ somewhat aware
- ○ aware
- ○ fully aware

How often do you use artificial intelligence (AI) Large Language Models (e.g., ChatGPT)?

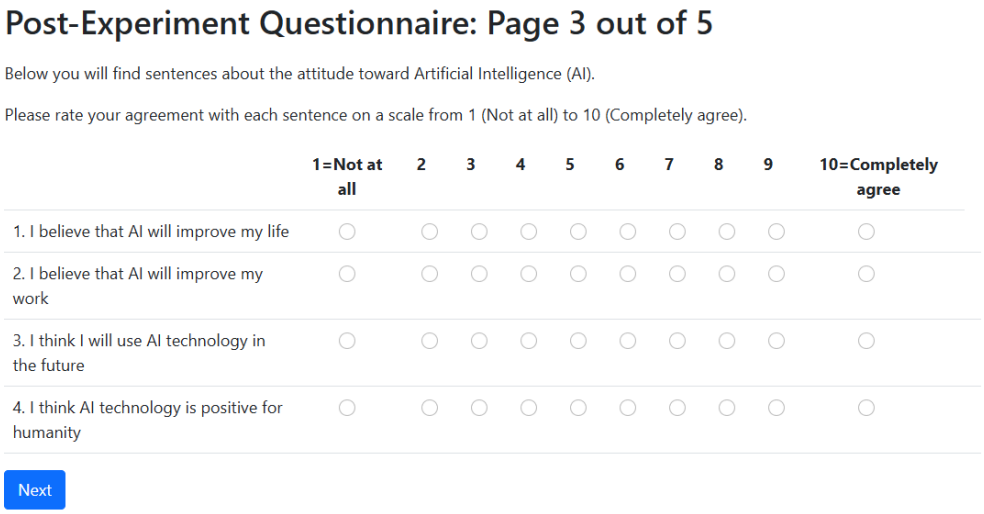
- ○ Never
- ○ Once every six months
- ○ Once a month
- ○ Once a week
- ○ Most days
- ○ Every day

Are you aware of the capability of artificial intelligence (AI) to generate art from text prompts?

- ○ completely unaware
- ○ not very aware
- ○ somewhat aware
- ○ aware
- ○ fully aware

Next

Figure D.3: General attitude toward AI.



Post-Experiment Questionnaire: Page 3 out of 5

Below you will find sentences about the attitude toward Artificial Intelligence (AI).

Please rate your agreement with each sentence on a scale from 1 (Not at all) to 10 (Completely agree).

	1=Not at all	2	3	4	5	6	7	8	9	10=Completely agree
1. I believe that AI will improve my life	○	○	○	○	○	○	○	○	○	○
2. I believe that AI will improve my work	○	○	○	○	○	○	○	○	○	○
3. I think I will use AI technology in the future	○	○	○	○	○	○	○	○	○	○
4. I think AI technology is positive for humanity	○	○	○	○	○	○	○	○	○	○

Next

Notes: Based on validated AI attitude scale (AIAS-4) due to Grassini (2023).

Figure D.4: Artwork debrief (*Human* treatments) and knowledge.

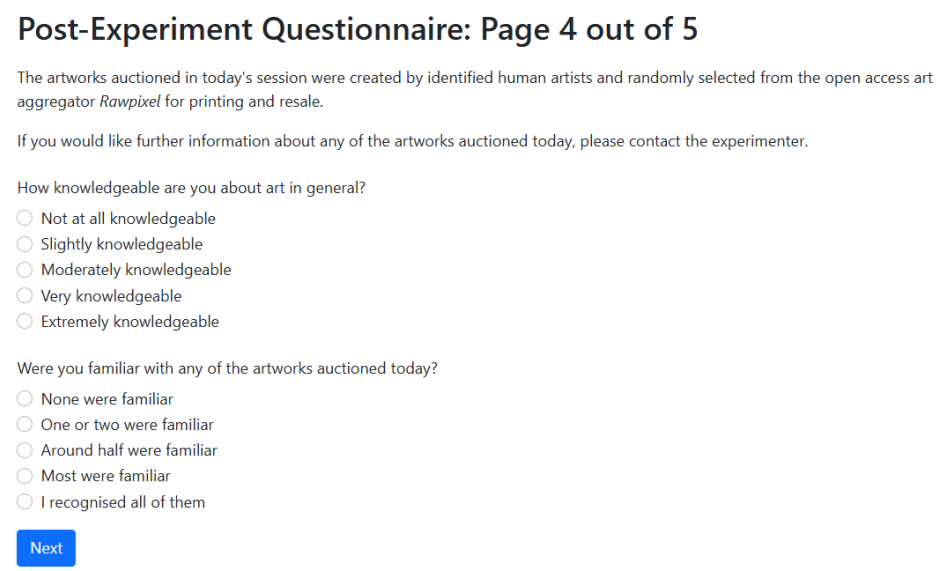


Figure D.5: Artwork debrief (*AI* treatments) and knowledge.

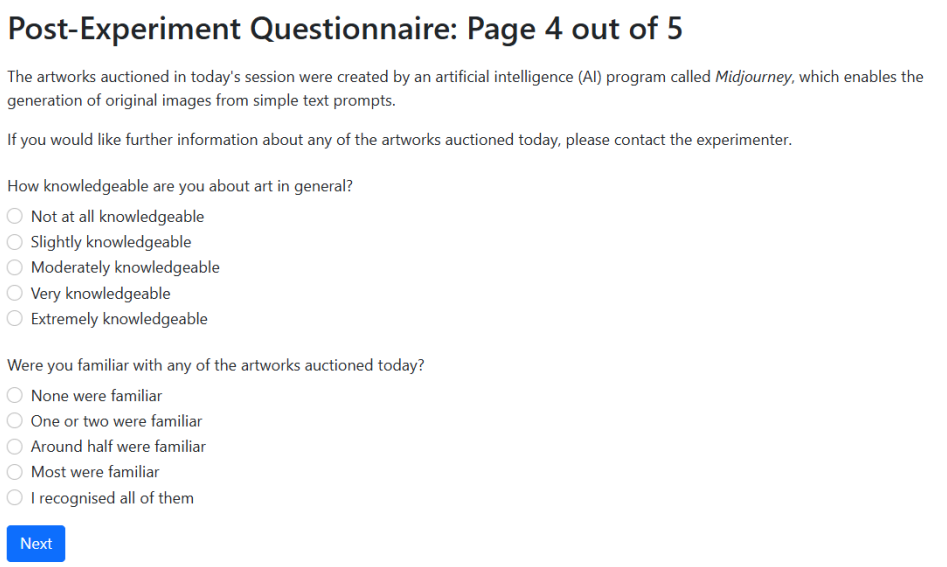
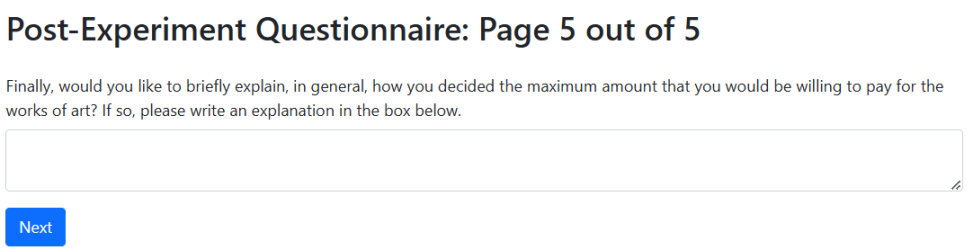


Figure D.6: Open text comment: Willingness to pay rationale.



D.2 Online experiments

Figure D.7: Page 1 - Socio-demographic characteristics.

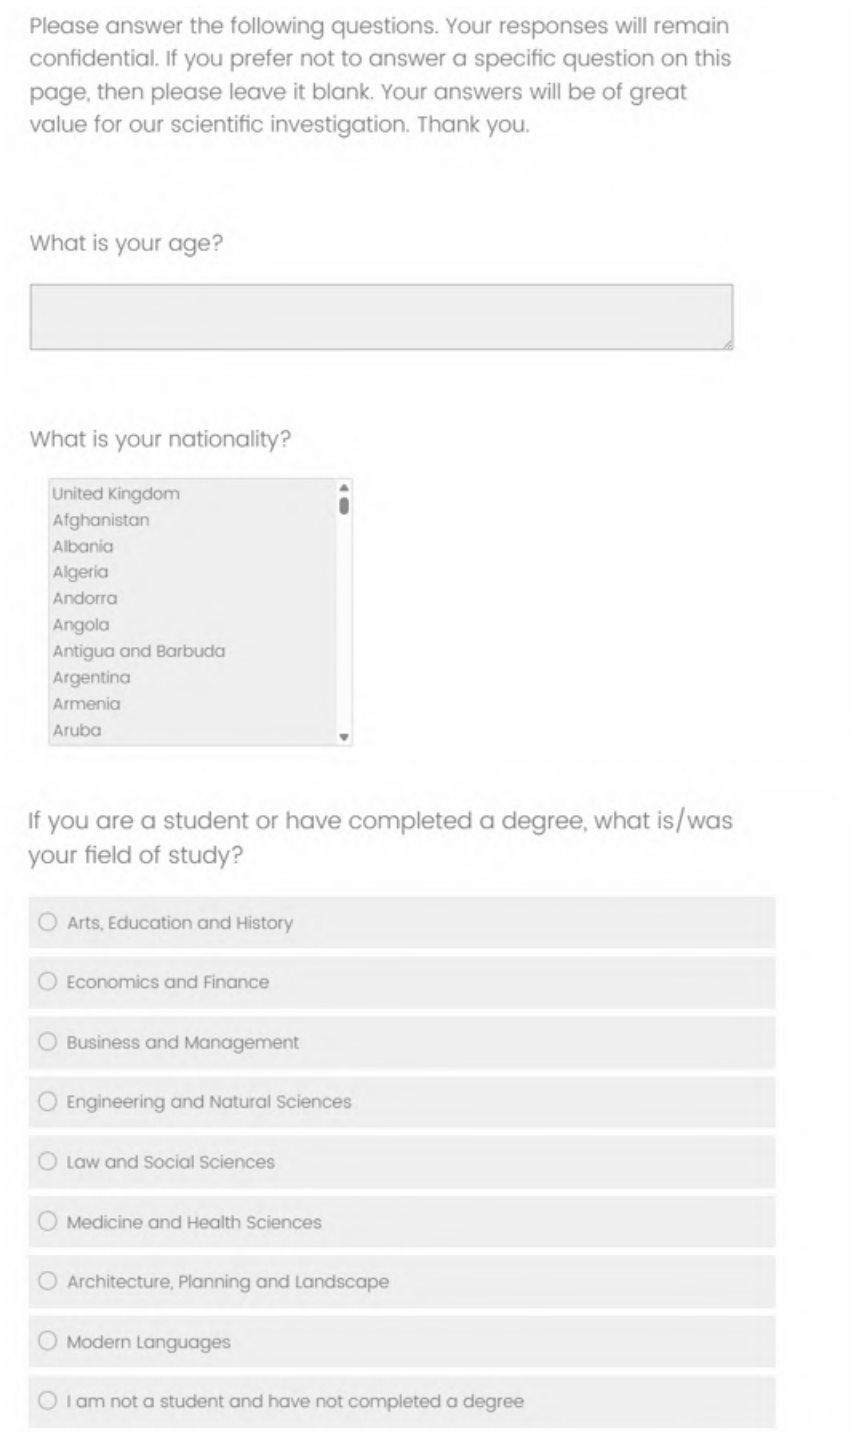


Figure D.8: Page 1 continued - Socio-demographic characteristics.

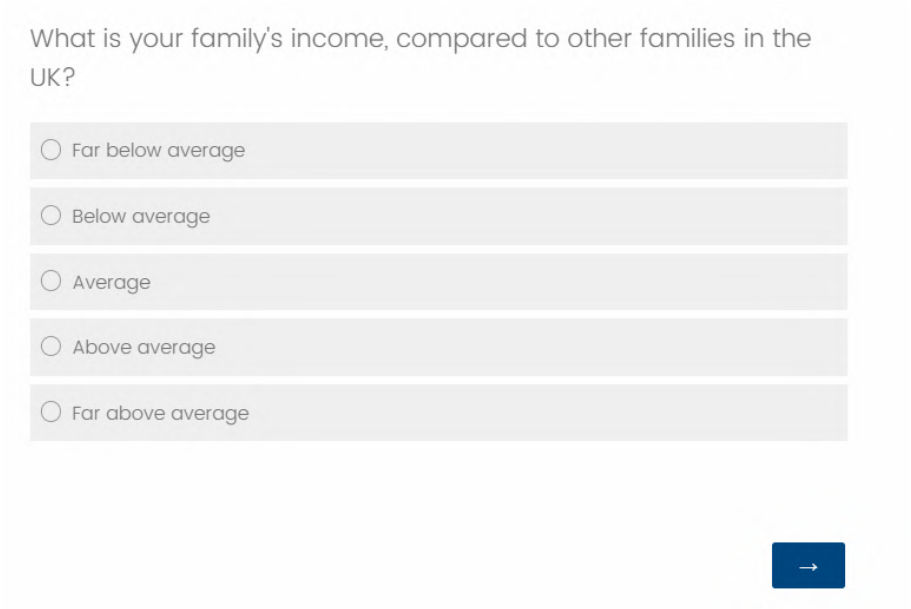


Figure D.9: Page 2 - Awareness of AI.

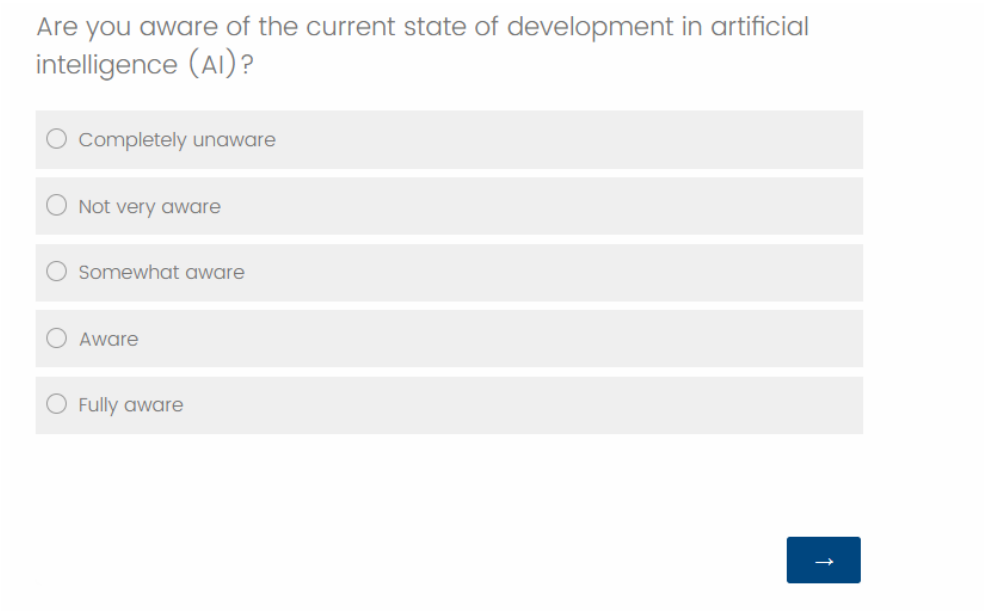


Figure D.10: Page 3 - Use of AI.

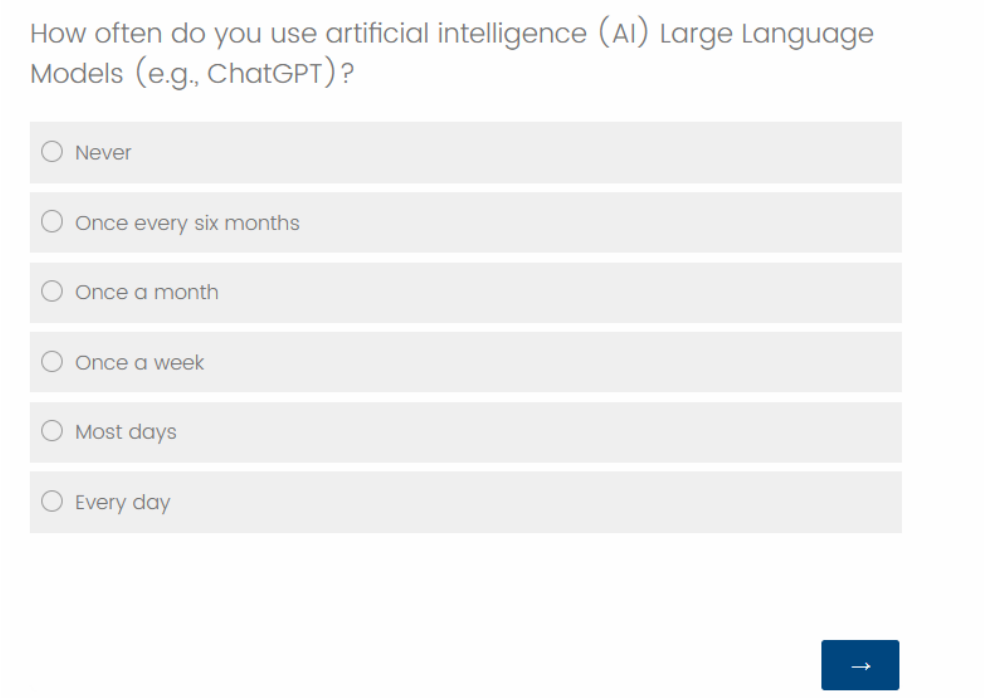


Figure D.11: Page 4 - Knowledge of AI - Artwork experiment.

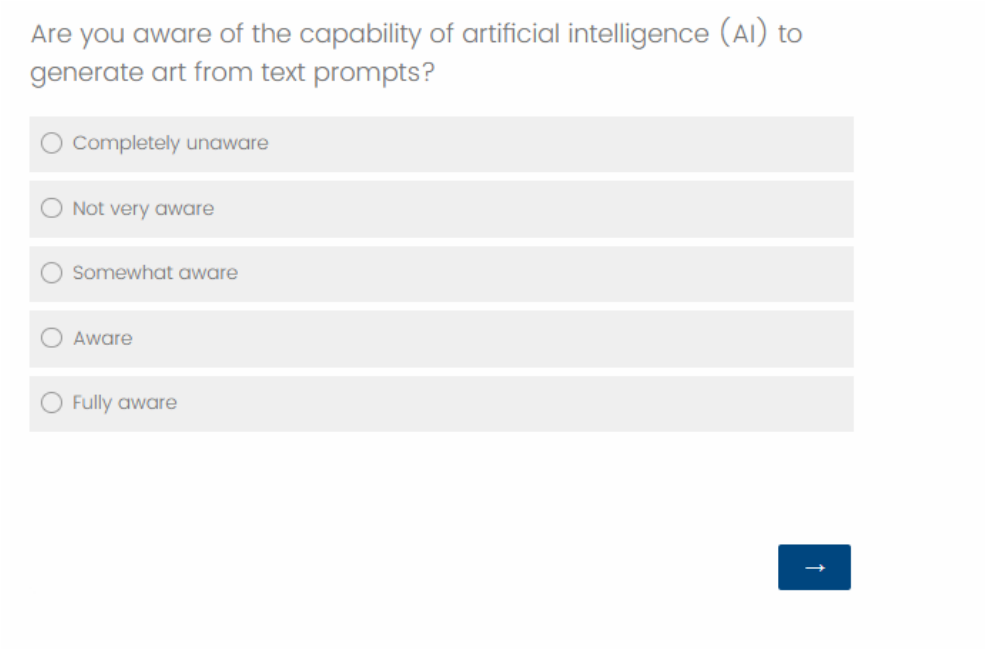


Figure D.12: Page 4 - Knowledge of AI - Poetry mug experiment.

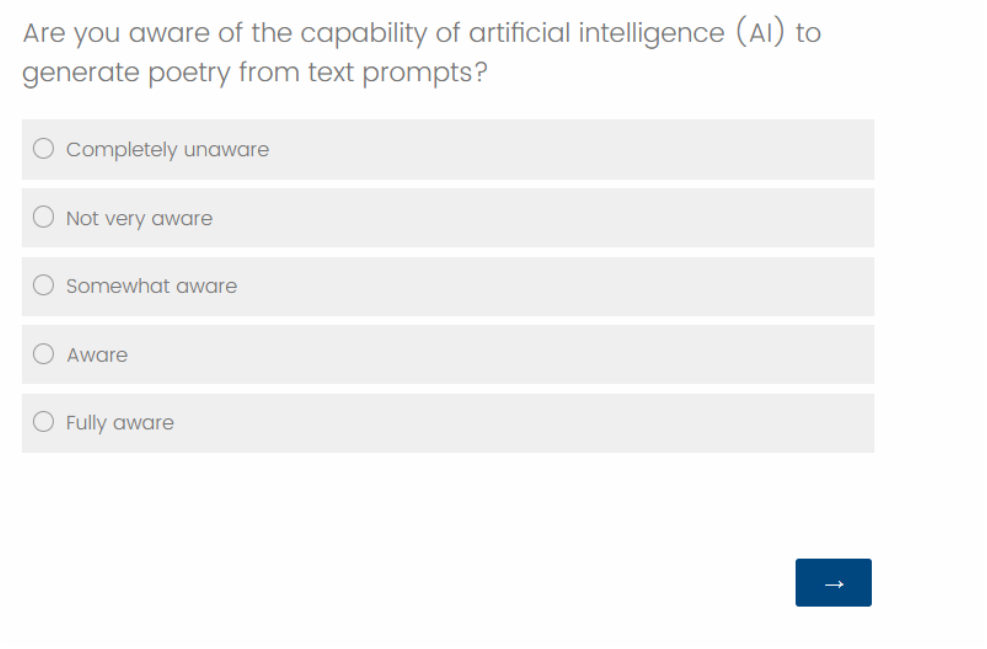
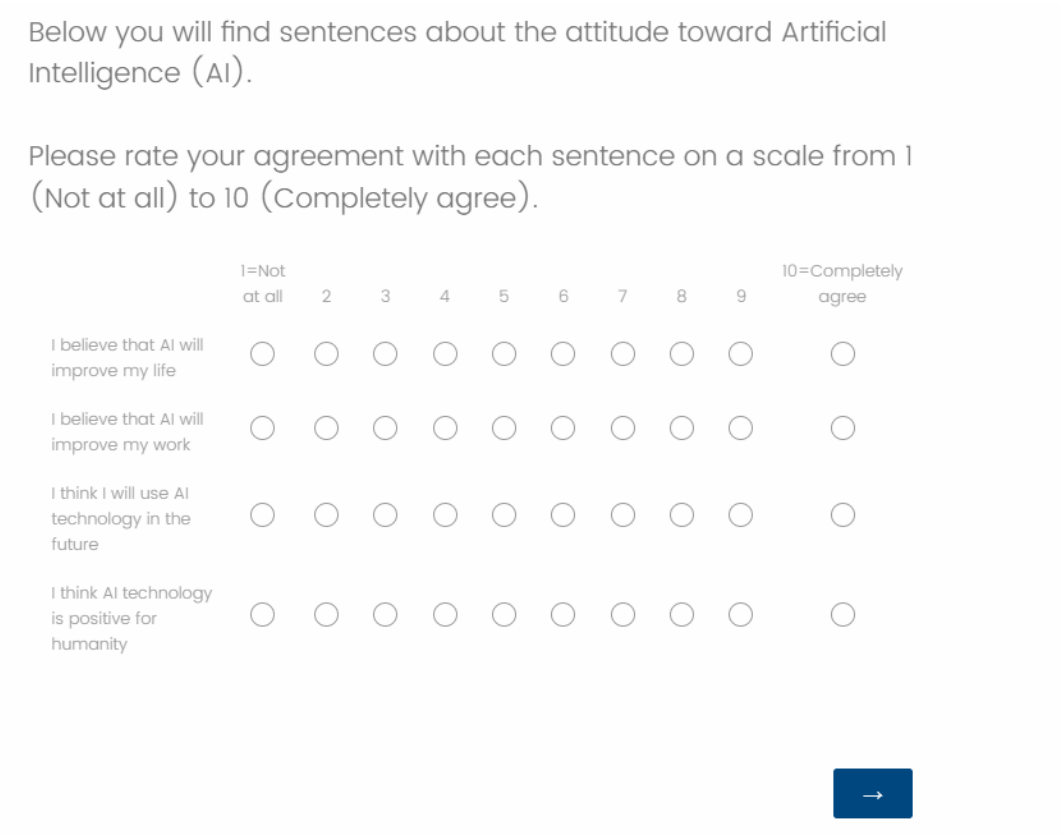


Figure D.13: Page 5 - General attitude toward AI.



Notes: Based on validated AI attitude scale (AIAS-4) due to Grassini (2023).

Figure D.14: Page 6 - Bid determinants - Artwork experiment.

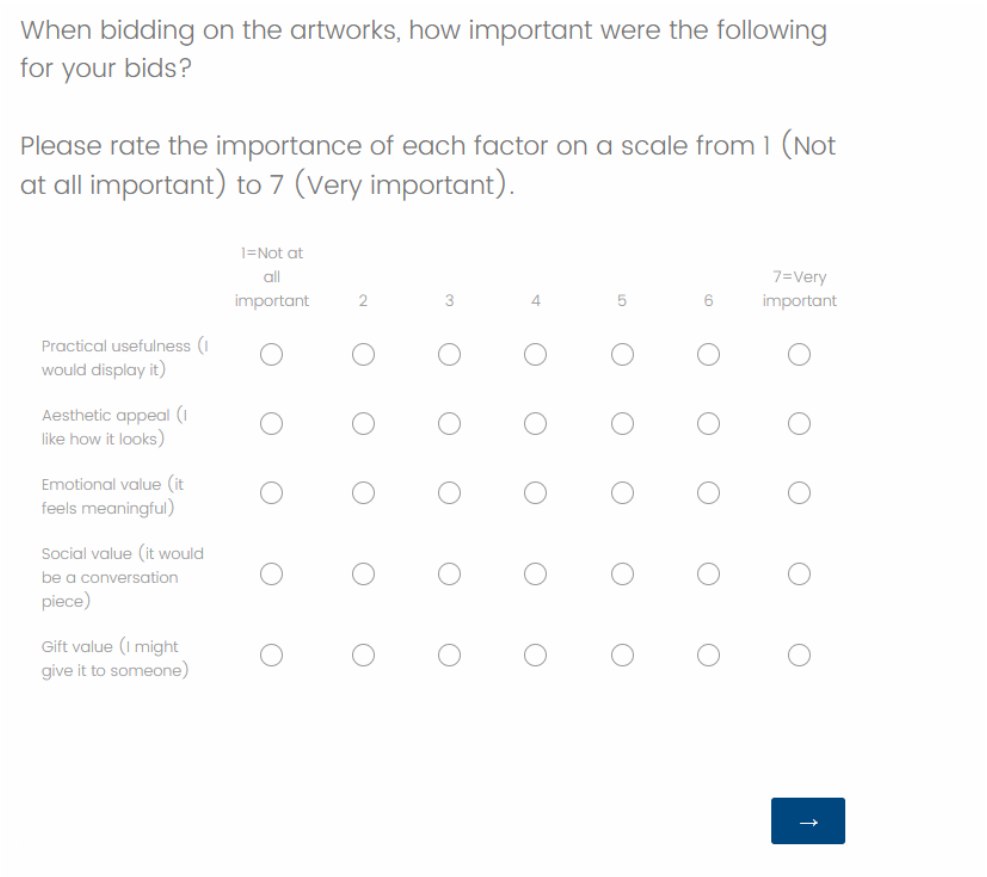


Figure D.15: Page 6 - Bid determinants - Poetry mug experiment.

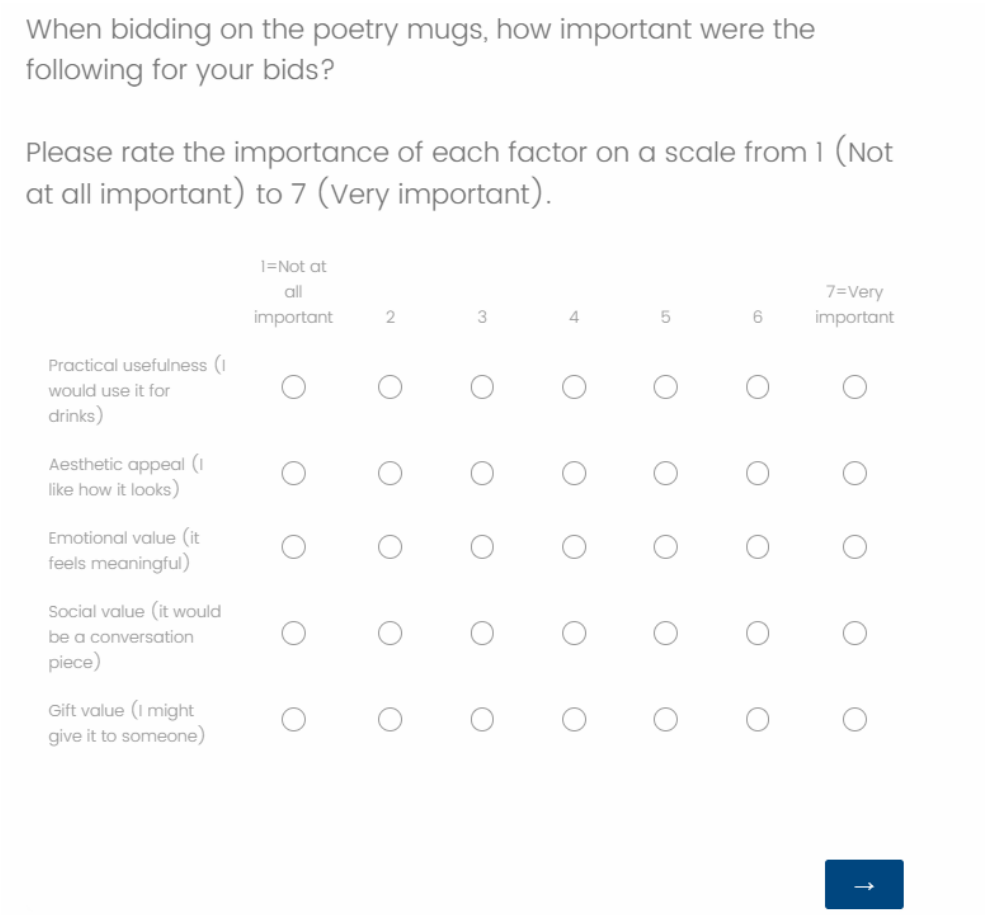


Figure D.16: Page 7 - Debrief - Artwork experiment - Human Info treatment.

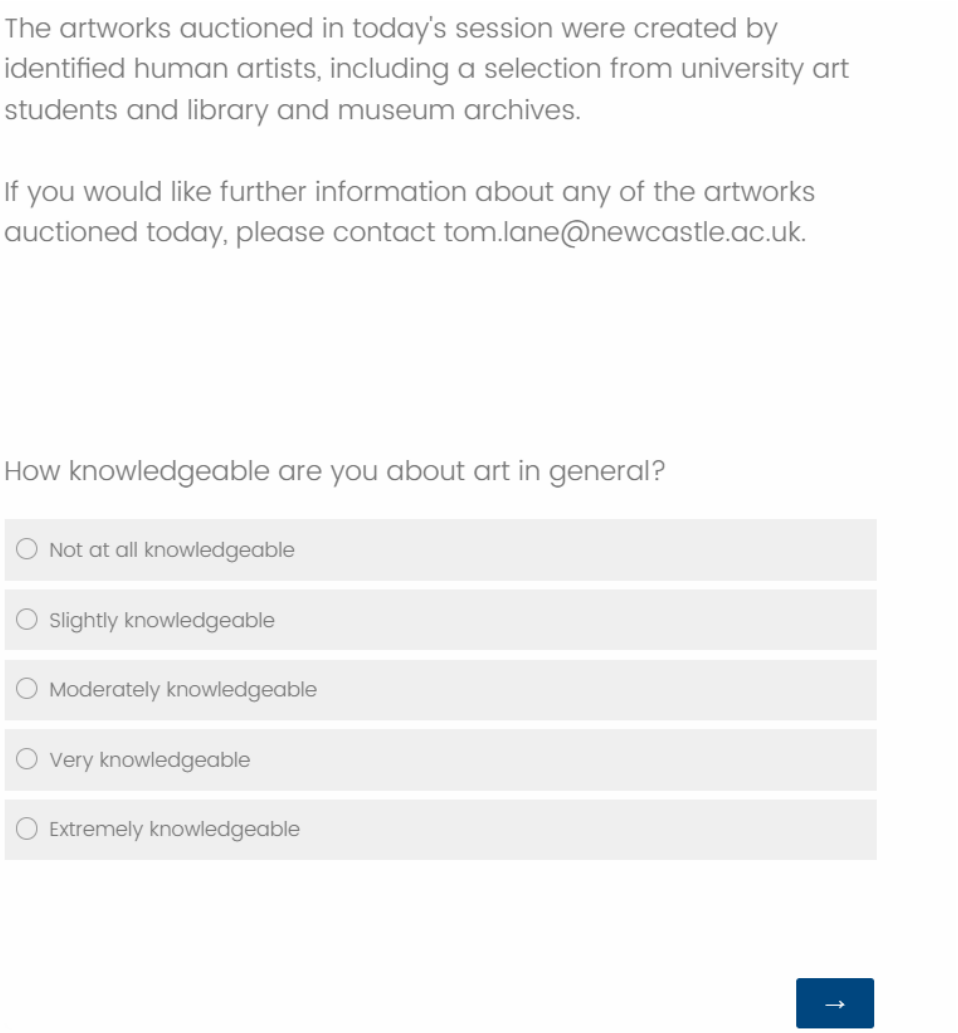


Figure D.17: Page 7 - Debrief - Artwork experiment - AI Info treatment.

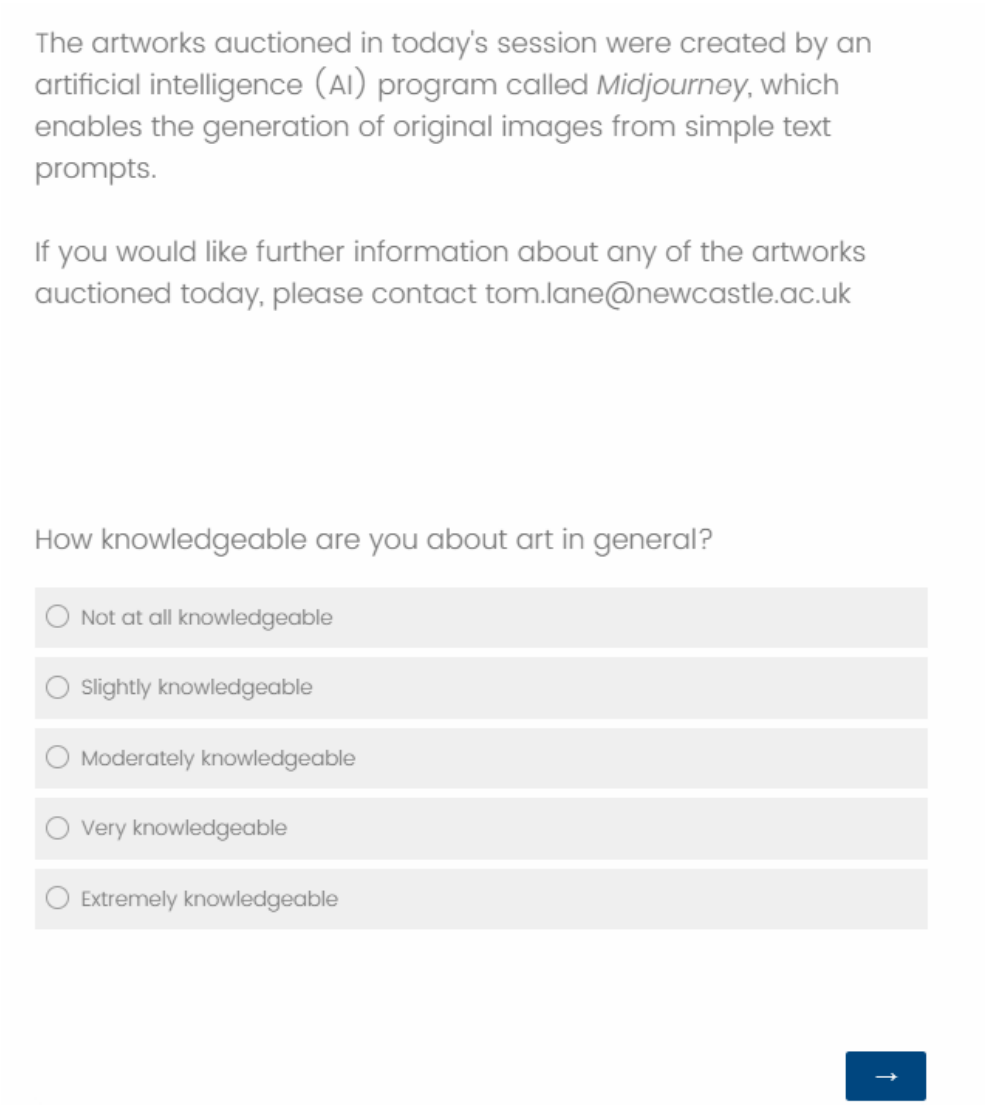


Figure D.18: Page 7 - Debrief - Poetry mug experiment - Human Info treatment.

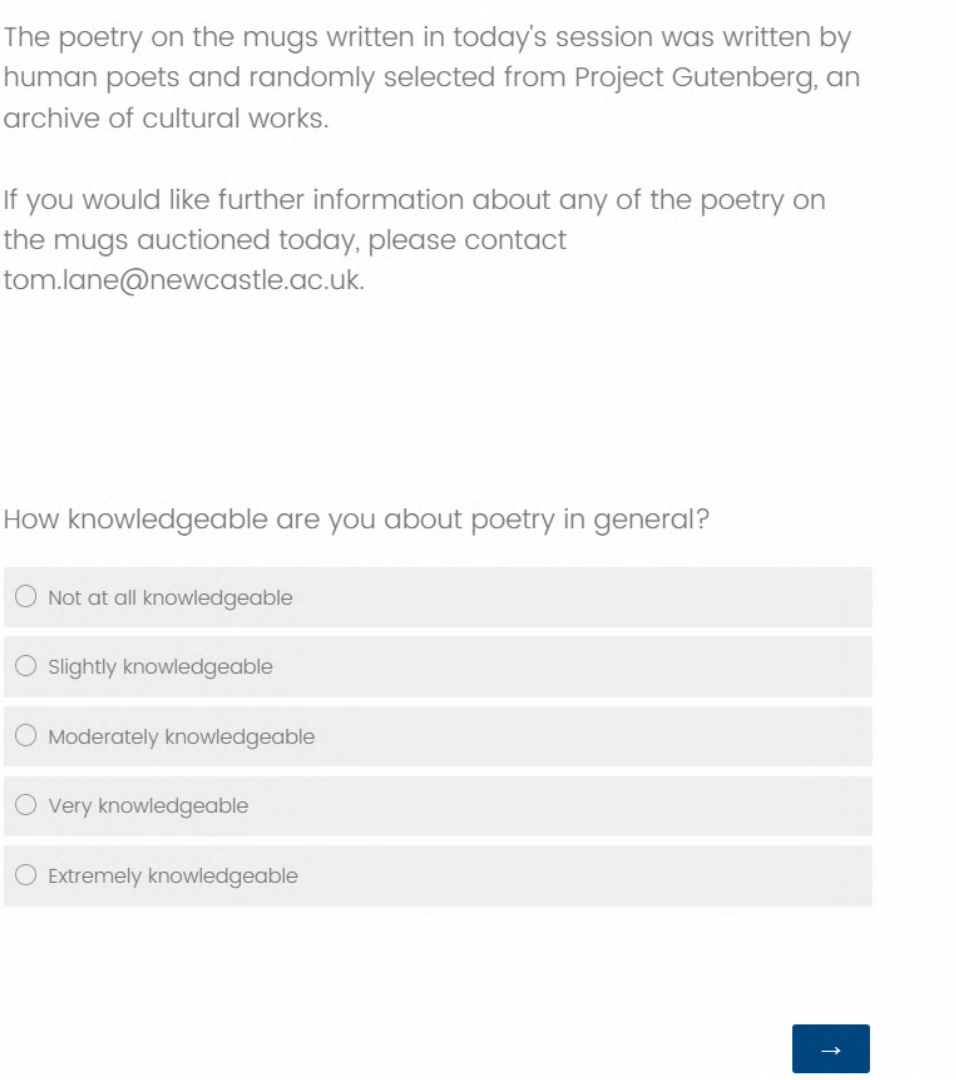


Figure D.19: Page 7 - Debrief - Poetry mug experiment - AI Info treatment.

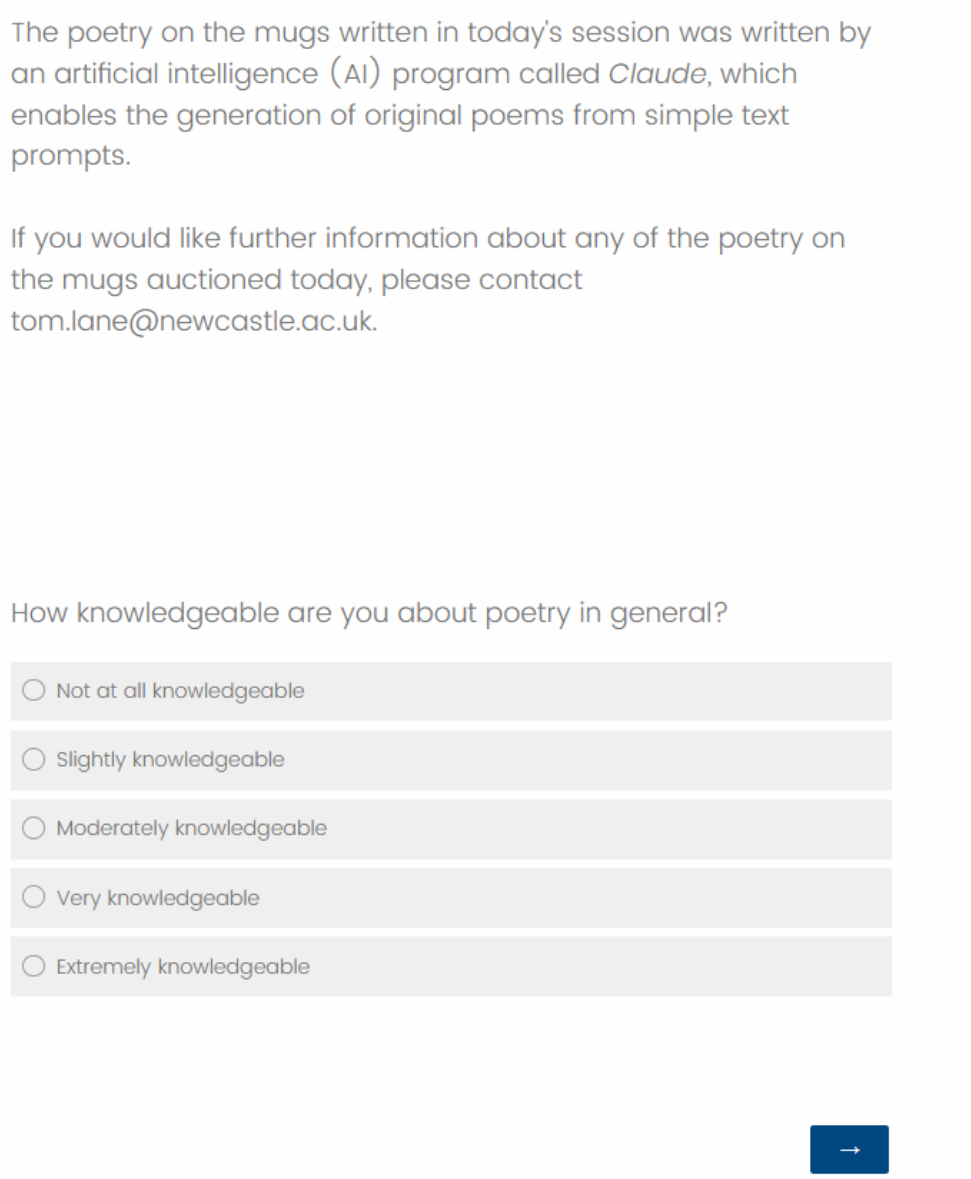


Figure D.20: Page 8 - Familiarity - Artwork experiment.

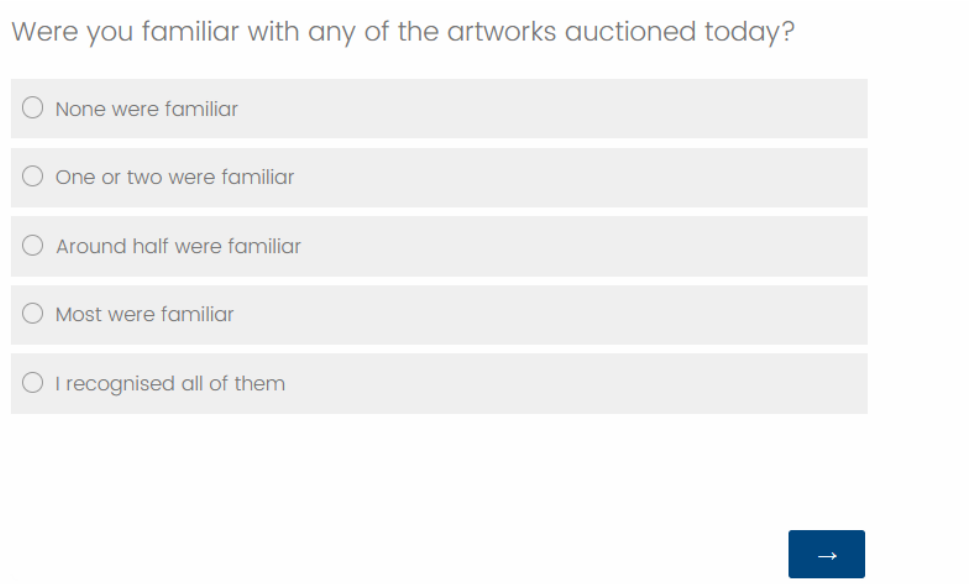


Figure D.21: Page 8 - Familiarity - Poetry mug experiment.

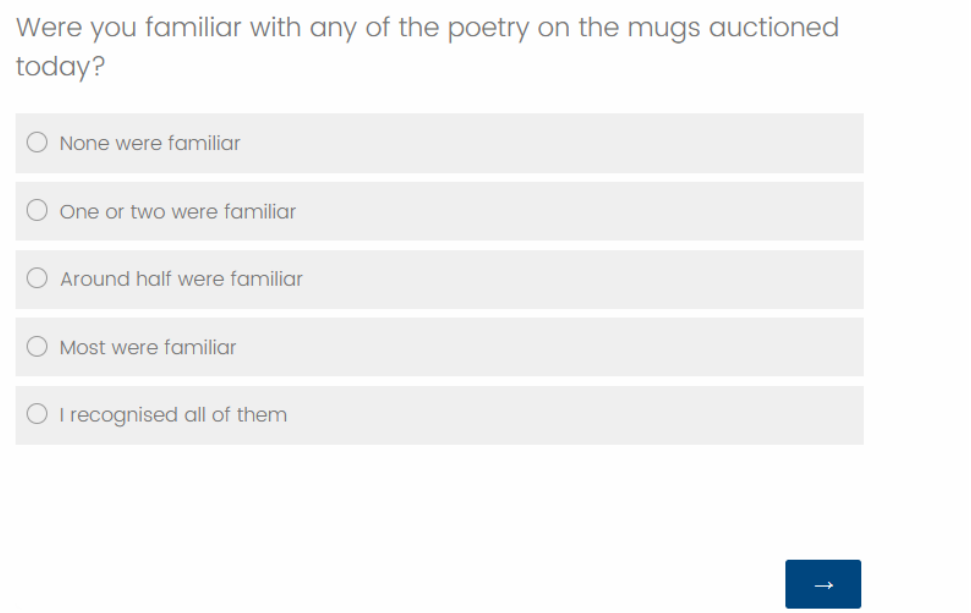


Figure D.22: Page 9 - Open text comment: Willingness to pay rationale - Artwork experiment.

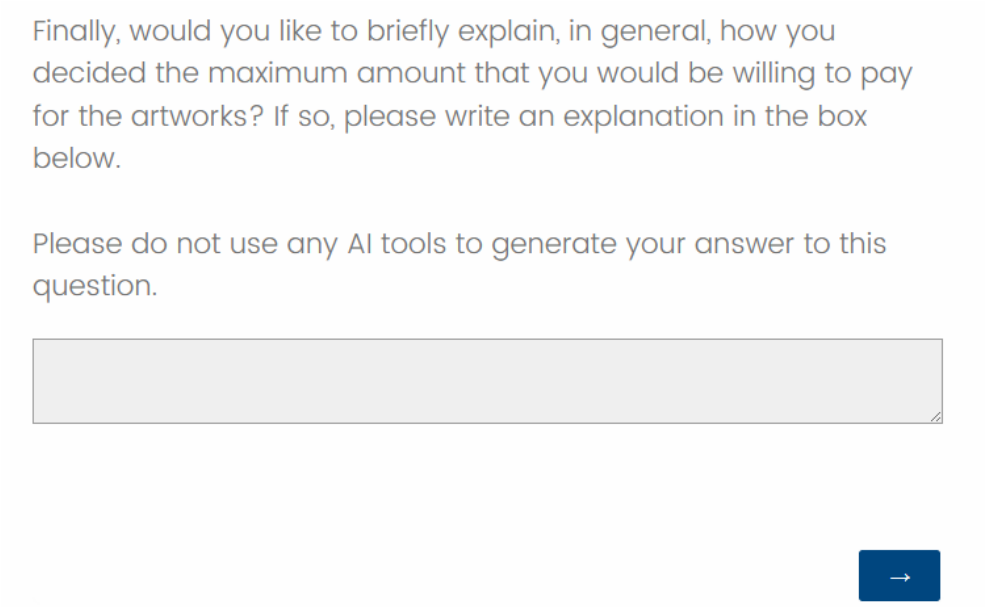
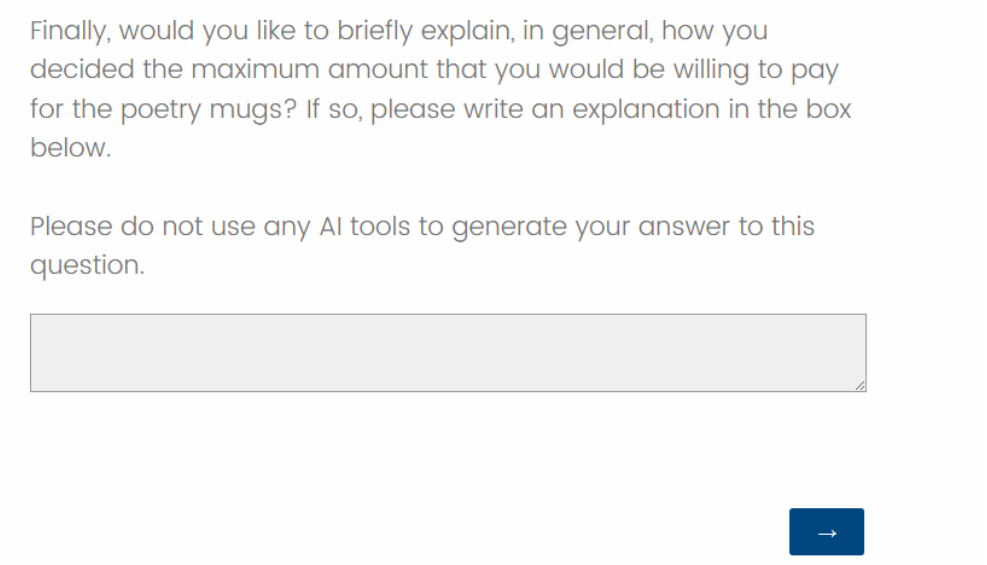


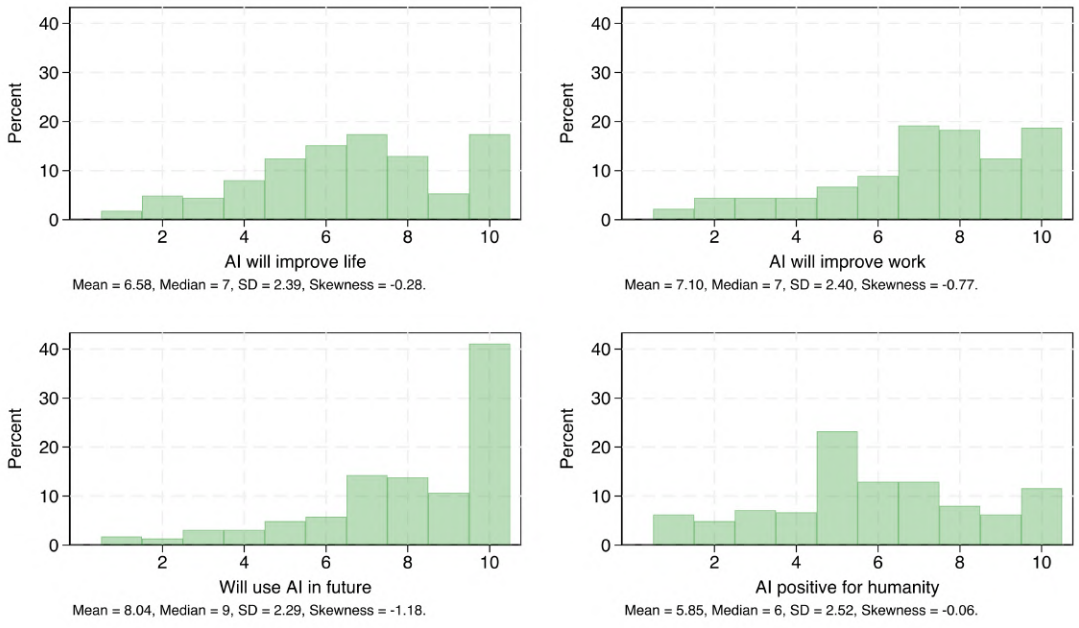
Figure D.23: Page 9 - Open text comment: Willingness to pay rationale - Poetry mug experiment.



E Descriptives and additional data analyses: Lab experiment

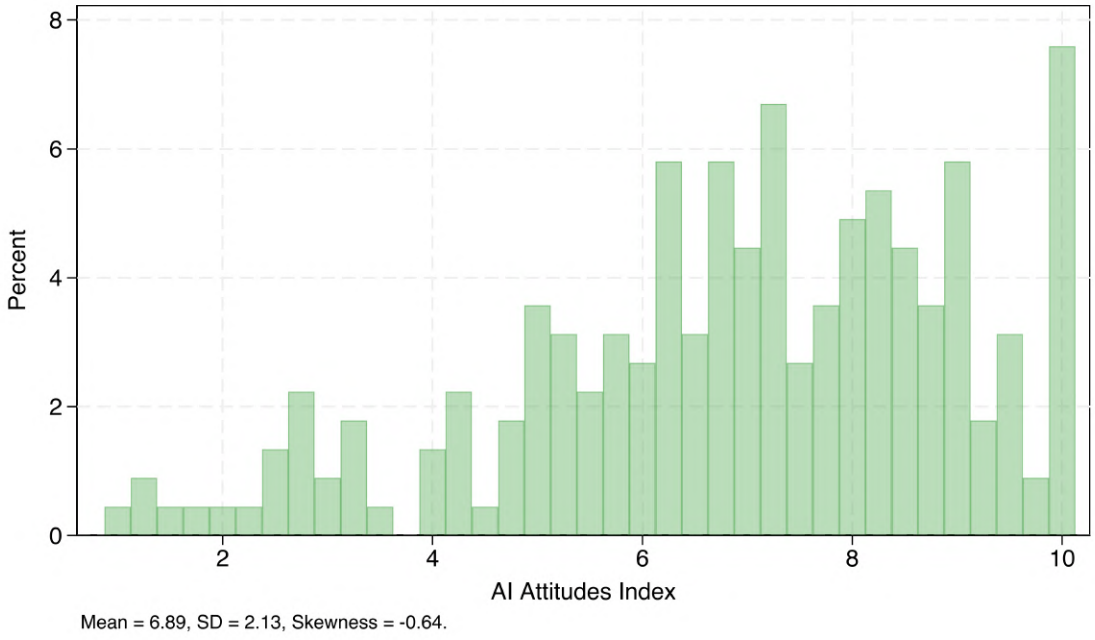
This section presents descriptive statistics and supplementary analyses that complement the findings from our lab experiment. We begin by examining the distributions of participants' AI attitudes, both for individual components and the composite index. We then provide summary statistics and balance tests across experimental conditions, followed by statistical tests of our hypotheses at the participant level. Next, we validate our pretesting procedure by comparing bids between uninformed conditions. We re-estimate our bid-level regressions using alternative estimators and test whether bidding behaviour changes as participants gain experience across the ten auction rounds. We then examine correlations between hypothetical willingness-to-pay measures and actual bids in the experiment. We also provide hypothesis tests at the individual image level, along with detailed bid distributions for each image across treatment conditions. Finally, we explore alternative specifications and heterogeneity in treatment effects based on participants' attitudes toward AI, knowledge of AI image generation capabilities, and knowledge about art.

Figure E.1: The Distribution of the AI Attitudes Index Components



Notes: The figure shows the distribution of the four components of the AI Attitudes index. All responses range from 1 to 10, where higher values indicate a more positive view, or frequent future use, of AI.

Figure E.2: The Distribution of the AI Attitudes Index



Notes: The figure shows the AI Attitudes index, which is the average of four components.

Table E.1: Summary statistics

	Count	Mean	Std. Dev.	Min.	25th pctl.	50th pctl.	75th pctl.	Max.
Average bid	224	2.01	2.23	0	0.29	1.14	3.07	9.36
18-21 years old	224	0.50	–	0	–	–	–	1
Male	224	0.49	–	0	–	–	–	1
UK national	224	0.40	–	0	–	–	–	1
Income above average	223	0.41	–	0	–	–	–	1
Studying Economics or Finance	222	0.16	–	0	–	–	–	1
Awareness of AI development	224	3.83	0.92	1.00	3.00	4.00	4.50	5.00
Frequency of AI/LLM use	224	4.33	1.37	1.00	4.00	5.00	5.00	6.00
Awareness of AI art generation	224	4.11	0.99	1.00	3.00	4.00	5.00	5.00
Art knowledge	224	2.36	0.96	1.00	2.00	2.00	3.00	5.00
AI Attitudes	224	6.89	2.13	1.00	5.75	7.25	8.50	10.00
AI will improve life	224	6.58	2.39	1.00	5.00	7.00	8.00	10.00
AI will improve work	224	7.10	2.40	1.00	6.00	7.00	9.00	10.00
Will use AI in future	224	8.04	2.29	1.00	7.00	9.00	10.00	10.00
AI positive for humanity	224	5.85	2.52	1.00	4.50	6.00	8.00	10.00

Table E.2: Covariate Balance by Treatment

	H1: AI Info vs AI No Info			H2: Human Info vs Human No Info			H3: Human Info vs AI Info		
	β	S.E.	p-value	β	S.E.	p-value	β	S.E.	p-value
18-21 years old	-0.107	0.117	0.365	0.229	0.109	0.038	0.264	0.115	0.024
Male	0.066	0.101	0.515	-0.029	0.096	0.765	-0.182	0.095	0.059
UK national	-0.064	0.118	0.590	-0.051	0.127	0.688	-0.053	0.130	0.684
Income above average	-0.119	0.107	0.269	0.007	0.094	0.937	0.097	0.102	0.345
Studying Economics or Finance	0.110	0.147	0.456	-0.121	0.126	0.339	-0.107	0.139	0.445
Awareness of AI development	0.026	0.058	0.655	-0.031	0.054	0.567	0.053	0.050	0.293
Frequency of AI/LLM use	0.002	0.037	0.947	-0.037	0.045	0.408	0.011	0.037	0.766
Art knowledge	0.038	0.052	0.470	0.116	0.053	0.029	-0.028	0.048	0.565
Constant	0.384	0.257	0.138	0.437	0.329	0.187	0.280	0.250	0.266
F-test p-value	0.611			0.081			0.138		

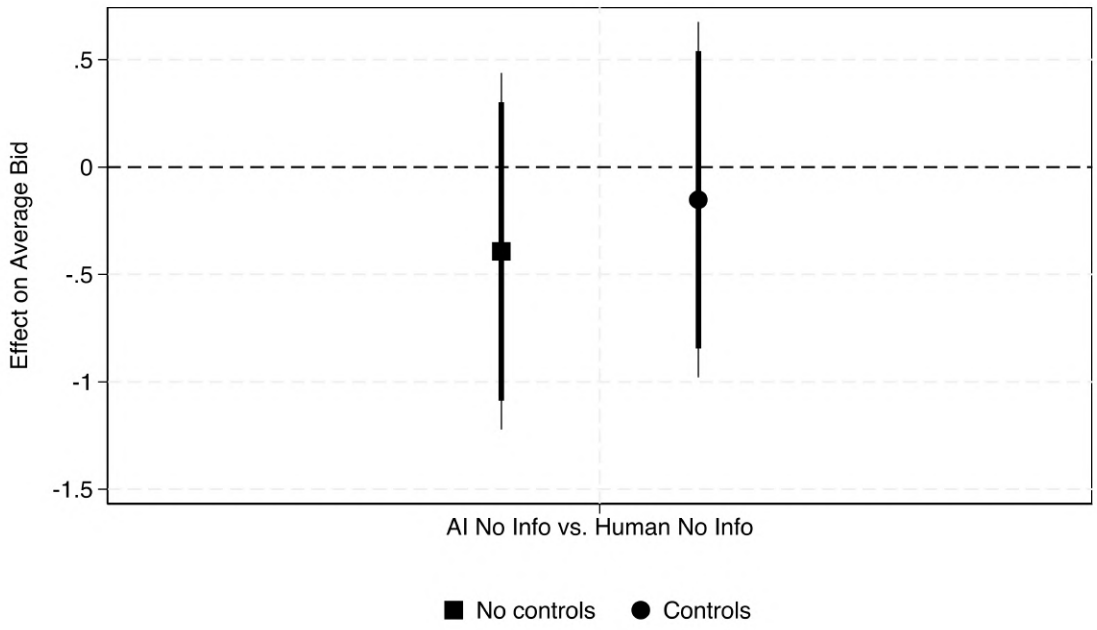
Notes: The dependent variable is a dummy for treatment status. The table shows estimates from a linear probability model, standard errors and p-values. The F-test is a test of the joint significance of all controls.

Table E.3: Main Hypothesis Tests

	Treatment 1 mean	Treatment 2 mean	Diff. in means	T-test p-value	K.S. statistic	K.S. p-value
H1: AI Info vs AI No Info	2.22	1.64	0.58	0.185	0.143	0.641
H2: Human Info vs Human No Info	2.14	2.01	0.12	0.77	0.235	0.083
H3: Human Info vs AI Info	2.14	2.22	-0.09	0.841	0.196	0.23

Notes: The table shows the mean bid under each treatment condition, the difference in the means, the p-value from a two-tailed t-test, the Kolmogorov-Smirnov statistic and p-value.

Figure E.3: Comparing AI No Info and Human No Info bids



Notes: The dependent variable is the average bid amount (£) per individual. Points show the marginal effects of the AI No Info vs Human No Info treatment dummy. All specifications include individual-level controls (age 18-21, male, UK national, above average income, studying economics or finance, AI development awareness, AI use frequency, and art knowledge). Thick (thin) lines denote 90% (95%) confidence intervals.

Table E.4: Alternative Estimators: OLS Regressions, Lab Experiment

	Bid														
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Panel A: Random effects															
AI Info vs AI No Info	0.582	0.582	0.582	0.582	0.307										
	(0.436)	(0.436)	(0.436)	(0.436)	(0.446)										
Human Info vs Human No Info						0.120	0.120	0.120	0.120	0.002					
						(0.410)	(0.410)	(0.410)	(0.410)	(0.431)					
Human Info vs AI Info											-0.089	-0.047	-0.067	-0.080	0.225
											(0.443)	(0.443)	(0.443)	(0.443)	(0.464)
Panel B: Random effects, standard errors clustered on the subject															
AI Info vs AI No Info	0.582	0.582	0.582	0.582	0.307										
	(0.428)	(0.428)	(0.428)	(0.429)	(0.407)										
Human Info vs Human No Info						0.120	0.120	0.120	0.120	0.002					
						(0.406)	(0.406)	(0.406)	(0.406)	(0.449)					
Human Info vs AI Info											-0.089	-0.047	-0.067	-0.080	0.225
											(0.441)	(0.441)	(0.441)	(0.444)	(0.430)
Aesthetic rating (pre-test)		✓		✓	✓		✓		✓	✓		✓		✓	✓
Hypothetical WTP (pre-test)			✓	✓	✓			✓	✓	✓			✓	✓	✓
Individual controls					✓					✓					✓
Observations	1,080	1,080	1,080	1,080	1,060	1,160	1,160	1,160	1,160	1,160	1,120	1,120	1,120	1,120	1,100
Subjects	108	108	108	108	106	116	116	116	116	116	112	112	112	112	110

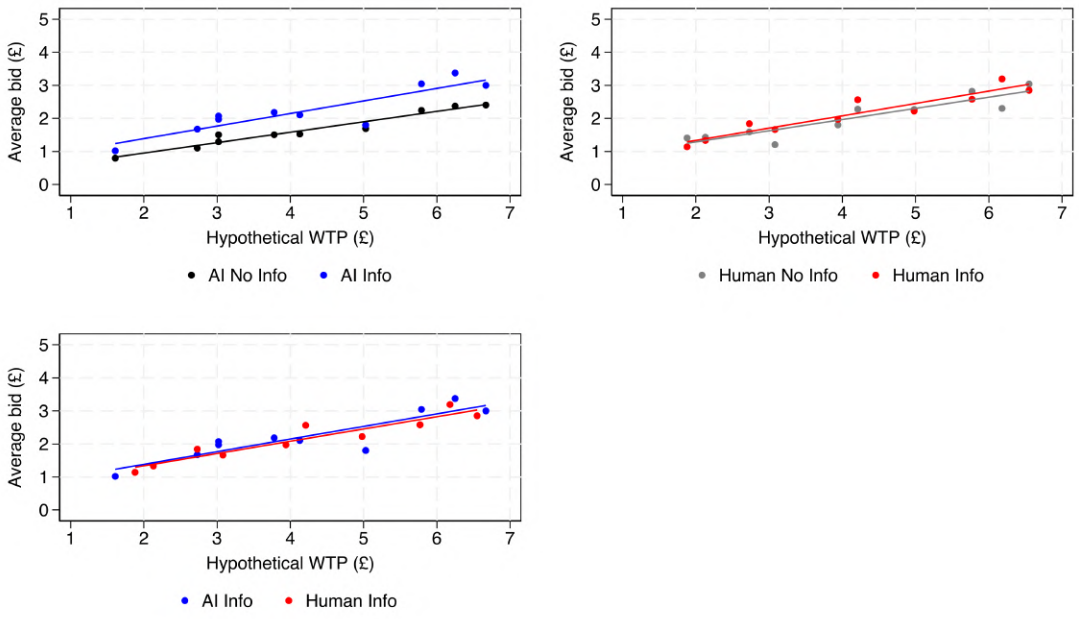
Notes: The dependent variable is the bid amount (£). The table re-estimates the specifications of Table 1 by OLS. Panel A reports random-effects estimates; Panel B reports the same estimates with standard errors clustered on the subject. Only the treatment indicators are shown; the rows beneath record which pre-test item scores and individual-level controls each column includes, and the samples and control sets are as in Table 1. Standard errors are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table E.5: Learning Across the Ten Rounds

	Bid		
	H1: AI Info vs AI No Info (1)	H2: Human Info vs Human No Info (2)	H3: Human Info vs AI Info (3)
Auction round	0.031	0.032	0.011
	(0.024)	(0.027)	(0.027)
Treatment × Round	-0.019	-0.045	-0.022
	(0.038)	(0.041)	(0.041)
Subject fixed effects	✓	✓	✓
Aesthetic rating (pre-test)	✓	✓	✓
Observations	1,080	1,160	1,120
Subjects	108	116	112

Notes: The dependent variable is the bid amount (£). The table shows estimates from linear regressions with subject fixed effects, estimated separately for each hypothesis. All specifications control for the artwork's aesthetic rating, elicited in the Prolific pre-testing phase. Standard errors, clustered at the subject level, are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure E.4: Relationship Between Hypothetical WTP and Auction Bids by Treatment



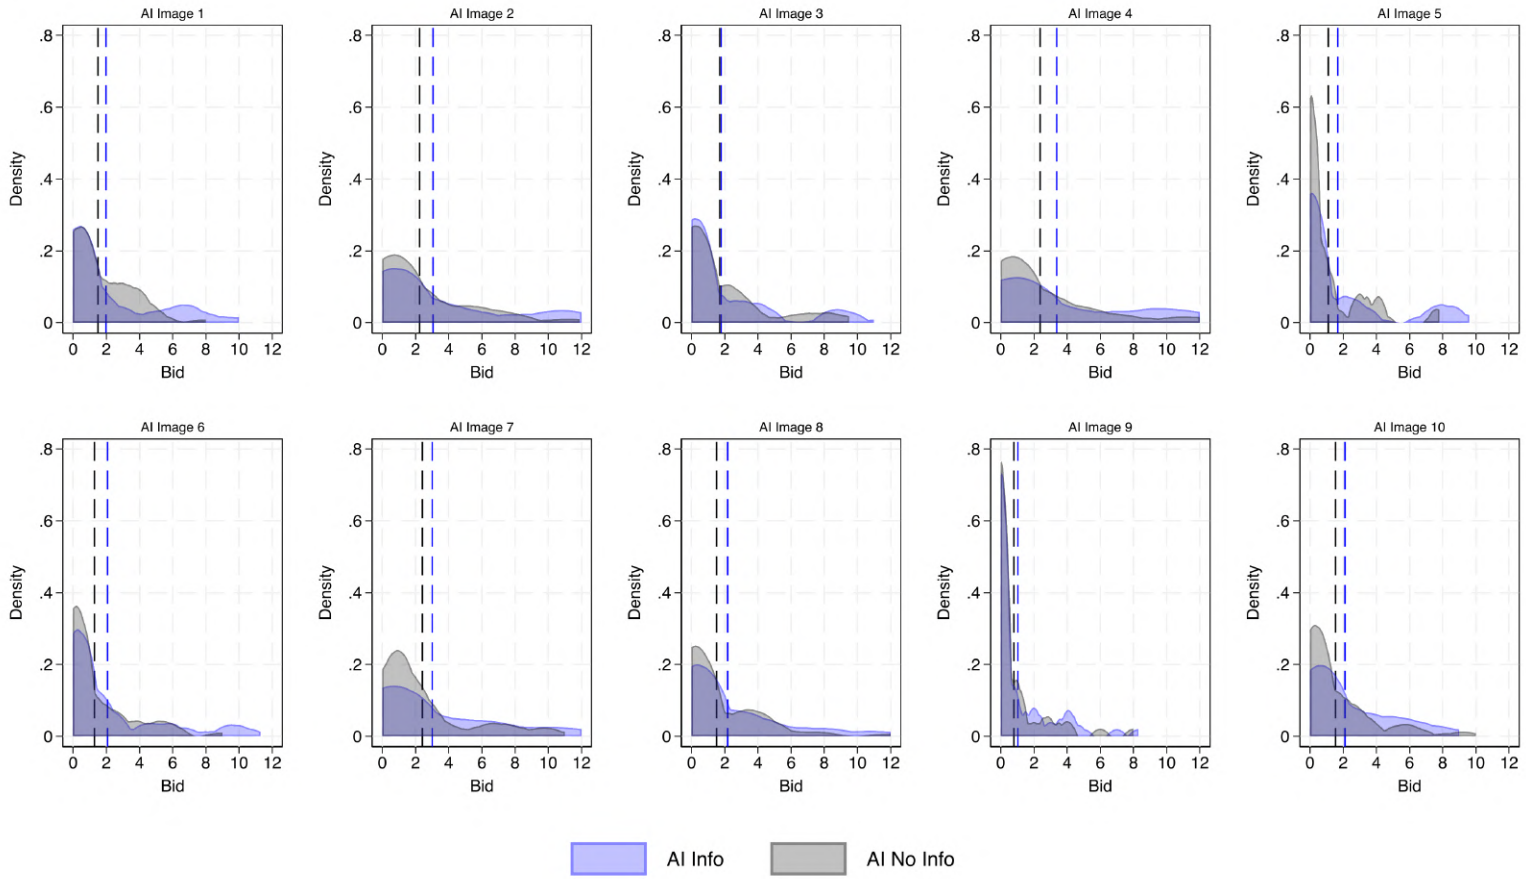
Notes: Each point shows the mean bid (£) and hypothetical WTP for one of the 10 artworks per treatment condition. The hypothetical WTP measure is the average willingness to pay (£) for each artwork elicited in the Prolific pre-testing phase. Lines show linear fits. The upper panel compares bids for AI artworks when their source is informed vs not informed (H1), the middle panel compares bids for human artworks when their source is informed vs not informed (H2), and the lower panel compares bids when participants are informed about AI vs human sources (H3).

Table E.6: Hypothesis Tests by Image

	Treatment 1 mean	Treatment 2 mean	Diff. in means	T-test p-value	K.S. statistic	K.S. p-value
H1: AI Info vs AI No Info						
AI image 1	1.97	1.50	0.47	0.316	0.159	0.5
AI image 2	3.04	2.24	0.80	0.232	0.122	0.815
AI image 3	1.80	1.69	0.11	0.835	0.08	0.996
AI image 4	3.37	2.37	1.00	0.166	0.137	0.689
AI image 5	1.67	1.10	0.57	0.222	0.107	0.916
AI image 6	2.07	1.29	0.78	0.13	0.135	0.713
AI image 7	3.00	2.40	0.60	0.359	0.148	0.593
AI image 8	2.18	1.50	0.68	0.225	0.118	0.846
AI image 9	1.02	0.79	0.22	0.509	0.078	0.996
AI image 10	2.11	1.52	0.58	0.224	0.135	0.713
H2: Human Info vs Human No Info						
Human image 1	3.19	2.30	0.89	0.122	0.252	0.05
Human image 2	1.66	1.21	0.45	0.218	0.213	0.144
Human image 3	2.56	2.28	0.28	0.603	0.14	0.617
Human image 4	2.58	2.82	-0.24	0.689	0.104	0.915
Human image 5	2.22	2.27	-0.05	0.926	0.162	0.433
Human image 6	1.33	1.43	-0.10	0.822	0.069	0.999
Human image 7	2.85	3.04	-0.19	0.761	0.077	0.995
Human image 8	1.14	1.41	-0.27	0.544	0.094	0.96
Human image 9	1.84	1.59	0.25	0.571	0.121	0.787
Human image 10	1.97	1.80	0.17	0.716	0.2	0.197

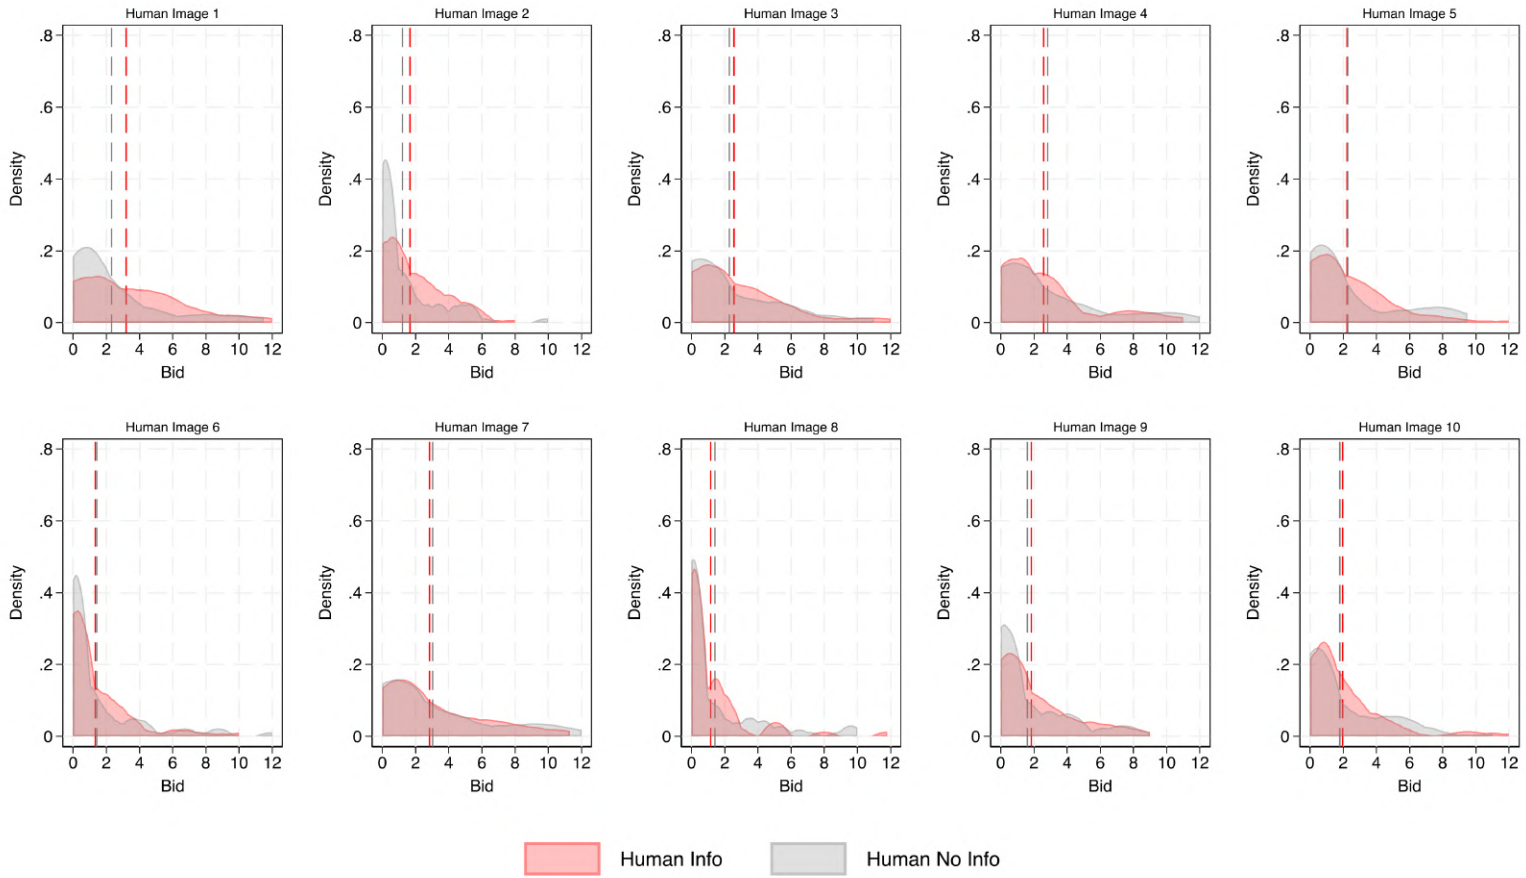
Notes: The table shows the mean bid under each treatment condition for each artwork, the difference in the means, the p-value from a two-tailed t-test, the Kolmogorov-Smirnov statistic and p-value.

Figure E.5: Distribution of Bids by Treatment: AI Images



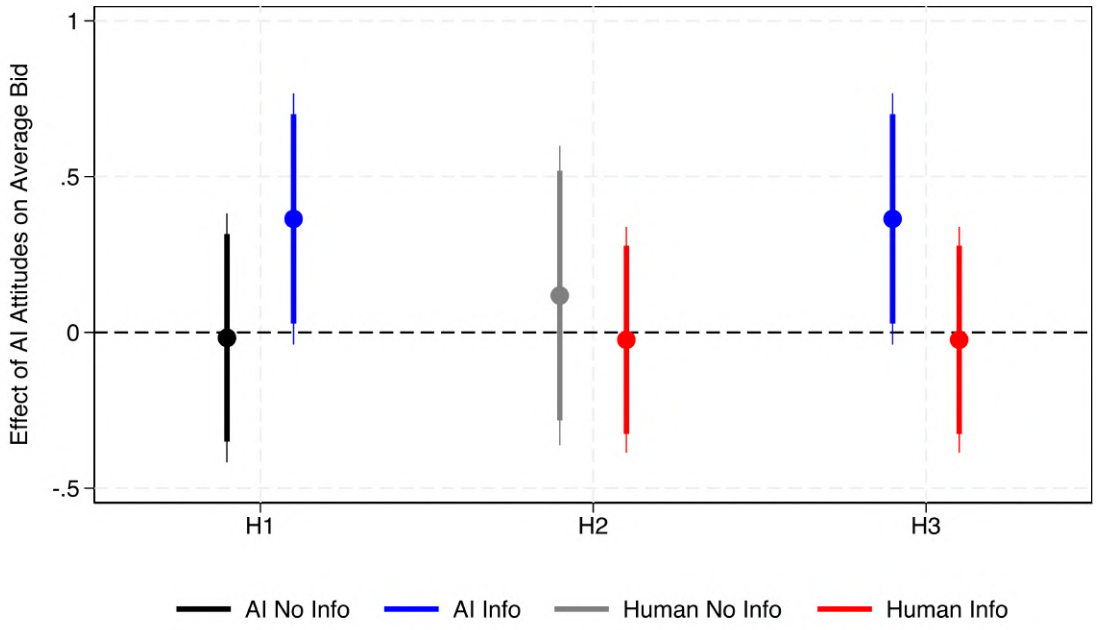
Notes: The figure shows kernel density estimates of the distribution of bids (£) per artwork. Dashed vertical lines show mean bids for each treatment condition.

Figure E.6: Distribution of Bids by Treatment: Human Images



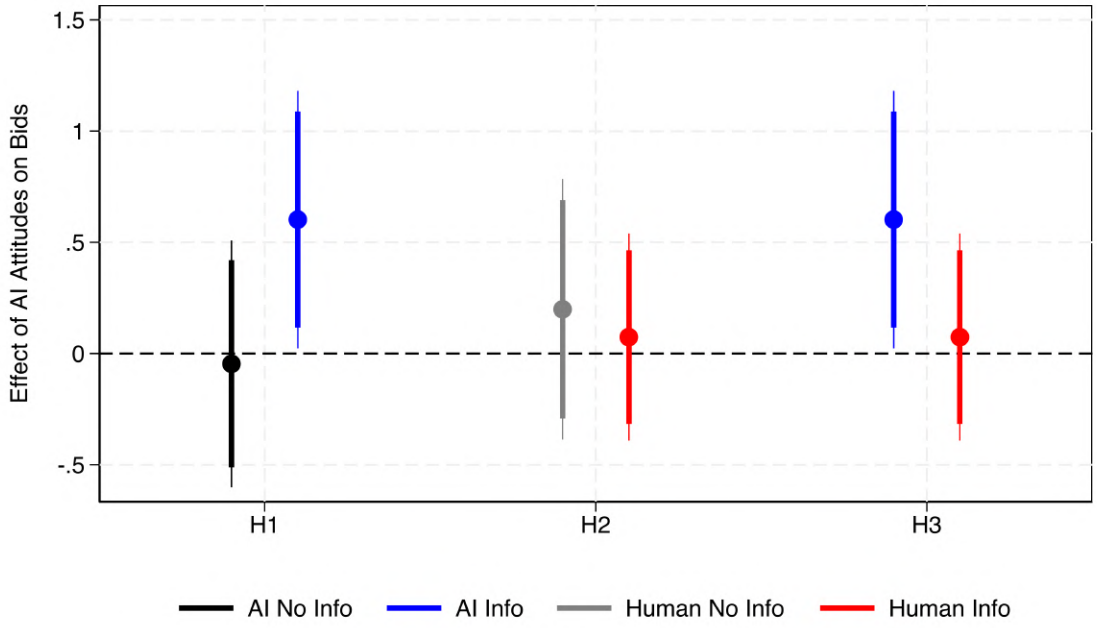
Notes: The figure shows kernel density estimates of the distribution of bids (£) per artwork. Dashed vertical lines show mean bids for each treatment condition.

Figure E.7: Effect of AI Attitudes on Bids by Treatment: OLS estimates



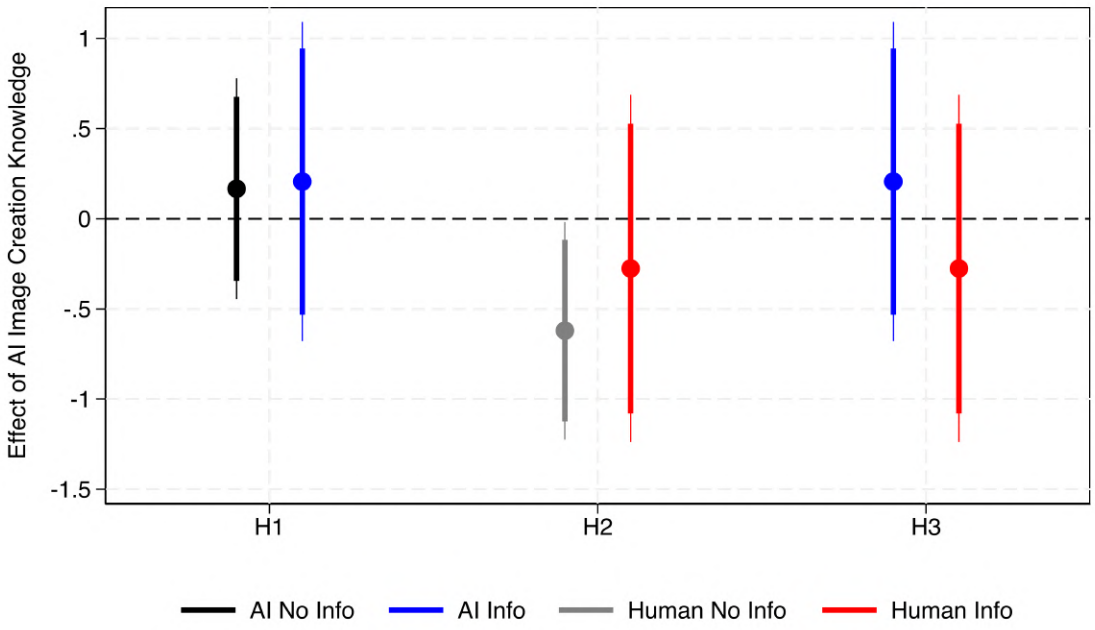
Notes: The dependent variable is the average bid amount (£) per individual. Points show the marginal effects of the AI attitudes index in each treatment condition, estimated from four separate OLS regressions. All specifications include individual-level controls (age 18-21, male, UK national, above average income, studying economics or finance, AI development awareness, AI use frequency, and art knowledge). Thick (thin) lines denote 90% (95%) confidence intervals.

Figure E.8: Effect of AI Attitudes on Bids by Treatment: Bid-level regressions



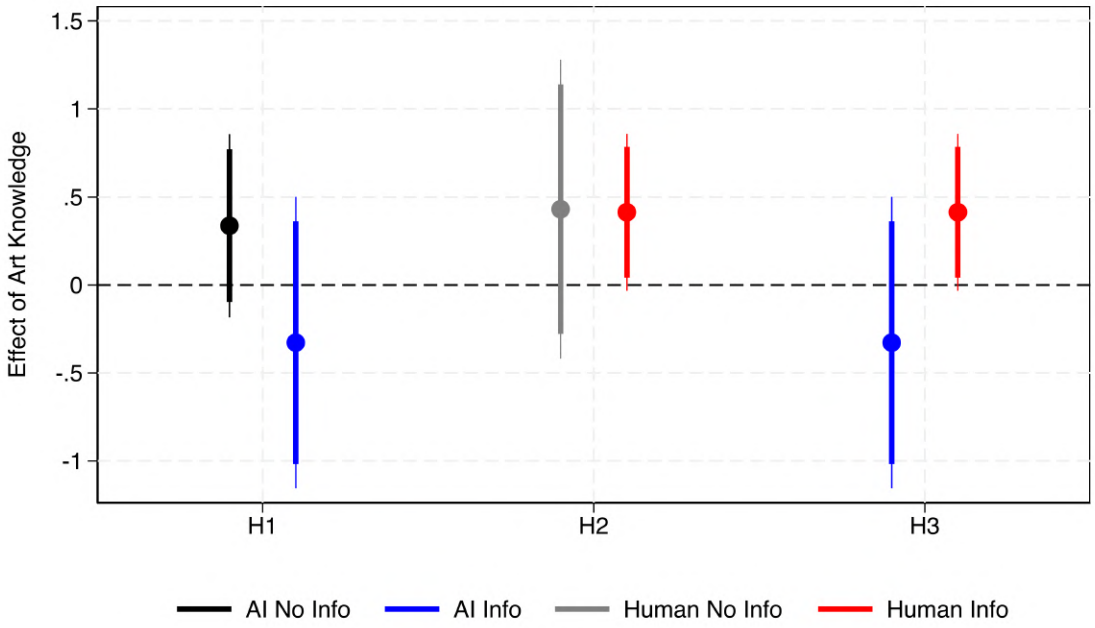
Notes: The dependent variable is the bid amount (£). Points show the marginal effects of the AI attitudes index in each treatment condition, estimated from three separate random-effect Tobit regressions. All specifications include individual-level controls (age 18-21, male, UK national, above average income, studying economics or finance, AI development awareness, AI use frequency, and art knowledge). Thick (thin) lines denote 90% (95%) confidence intervals.

Figure E.9: Effect of AI Image Creation Knowledge on Bids by Treatment



Notes: The dependent variable is the average bid amount (£) per individual. Points show the marginal effects of how aware an individual is that AI can generate images in each treatment condition, estimated from four separate Tobit regressions. All specifications include individual-level controls (age 18-21, male, UK national, above average income, studying economics or finance, AI development awareness, AI use frequency, and art knowledge). Thick (thin) lines denote 90% (95%) confidence intervals.

Figure E.10: Effect of Art Knowledge on Bids by Treatment



Notes: The dependent variable is the average bid amount (£) per individual. Points show the marginal effects of art knowledge in each treatment condition, estimated from four separate Tobit regressions. All specifications include individual-level controls (age 18-21, male, UK national, above average income, studying economics or finance, AI development awareness and AI use frequency). Thick (thin) lines denote 90% (95%) confidence intervals.

F Descriptives and additional data analyses: Online experiments

This section presents descriptive statistics and supplementary analyses that complement the findings from our online experiments. We begin by providing summary statistics for each of the two experiments distinguished by item type, artworks and haiku mugs, and balance tests across experimental conditions, followed by a comparison of the composition of our achieved sample against the recruitment quota targets. We then report statistical tests of our hypotheses at the participant level for the full sample, re-estimate our bid-level regressions using alternative estimators, and repeat the participant-level hypothesis tests for the subsample of 18 to 34 year olds. We then test whether bidding behaviour changes as participants gain experience across the ten auction rounds. We next examine the distributions of participants' AI attitudes, both for individual components and the composite index. Finally, we examine correlations between the hypothetical willingness-to-pay measures elicited in pre-testing and actual bids in the experiment.

Table F.1: Summary Statistics: Online Experiments

	Count	Mean	Std. Dev.	Min.	25th pctl.	50th pctl.	75th pctl.	Max.
Panel A: Artworks								
Average bid	525	5.48	3.42	0	2.70	5.69	8.24	12.00
18-24 years old	525	0.13	–	0	–	–	–	1
25-34 years old	525	0.18	–	0	–	–	–	1
35-44 years old	525	0.19	–	0	–	–	–	1
45-54 years old	525	0.16	–	0	–	–	–	1
55+ years old	525	0.33	–	0	–	–	–	1
Female	525	0.49	–	0	–	–	–	1
UK national	524	0.94	–	0	–	–	–	1
Income above average	525	0.27	–	0	–	–	–	1
Student or degree holder	524	0.56	–	0	–	–	–	1
Awareness of AI development	525	3.35	1.02	1.00	3.00	3.00	4.00	5.00
Frequency of AI/LLM use	525	3.49	1.78	1.00	1.00	4.00	5.00	6.00
Awareness of AI art generation	525	3.31	1.22	1.00	2.00	3.00	4.00	5.00
Art/poetry knowledge	525	2.36	1.13	1.00	1.00	2.00	3.00	5.00
AI Attitudes	525	5.73	2.68	1.00	3.50	6.00	8.00	10.00
AI will improve life	525	5.61	2.82	1.00	3.00	6.00	8.00	10.00
AI will improve work	525	5.41	3.07	1.00	2.00	6.00	8.00	10.00
Will use AI in future	525	6.35	2.99	1.00	4.00	7.00	9.00	10.00
AI positive for humanity	525	5.56	2.75	1.00	3.00	6.00	8.00	10.00
Panel B: Haiku mugs								
Average bid	541	5.29	3.51	0	2.17	5.18	8.16	12.00
18-24 years old	541	0.09	–	0	–	–	–	1
25-34 years old	541	0.18	–	0	–	–	–	1
35-44 years old	541	0.15	–	0	–	–	–	1
45-54 years old	541	0.18	–	0	–	–	–	1
55+ years old	541	0.39	–	0	–	–	–	1
Female	541	0.51	–	0	–	–	–	1
UK national	540	0.94	–	0	–	–	–	1
Income above average	541	0.23	–	0	–	–	–	1
Student or degree holder	539	0.55	–	0	–	–	–	1
Awareness of AI development	541	3.29	0.96	1.00	3.00	3.00	4.00	5.00
Frequency of AI/LLM use	541	3.38	1.74	1.00	1.00	4.00	5.00	6.00
Art/poetry knowledge	541	2.35	1.01	1.00	2.00	2.00	3.00	5.00
AI Attitudes	541	5.61	2.60	1.00	3.75	5.50	7.75	10.00
AI will improve life	541	5.47	2.71	1.00	3.00	5.00	8.00	10.00
AI will improve work	541	5.22	2.99	1.00	2.00	5.00	8.00	10.00
Will use AI in future	541	6.26	2.88	1.00	5.00	7.00	9.00	10.00
AI positive for humanity	541	5.49	2.71	1.00	4.00	5.00	8.00	10.00

Notes: The table reports summary statistics for the samples of the online experiments. Panel A covers participants who bid on artworks and Panel B those who bid on haiku mugs. Average bid is the participant's mean bid (£) across the ten auction rounds. Standard deviations and percentiles are omitted for binary variables. AI Attitudes is the mean of the four component items listed beneath it, each measured on a 1-10 scale. Awareness of AI art generation was elicited in the artwork experiment only.

Table F.2: Covariate Balance by Treatment: Online Experiments

	H4: Artworks AI Info vs Human Info			H5: Haiku mugs AI Info vs Human Info		
	β	S.E.	p-value	β	S.E.	p-value
25-34 years old	0.101	0.081	0.216	-0.099	0.087	0.253
35-44 years old	0.114	0.080	0.155	-0.058	0.091	0.520
45-54 years old	0.023	0.084	0.783	0.076	0.089	0.394
55+ years old	0.111	0.079	0.163	-0.038	0.083	0.645
Female	0.041	0.045	0.354	-0.023	0.044	0.600
UK national	0.005	0.091	0.953	-0.023	0.088	0.793
Income above average	-0.085	0.055	0.122	0.098	0.056	0.082
Student or degree holder	0.088	0.051	0.084	-0.064	0.049	0.197
Awareness of AI development	-0.014	0.026	0.602	-0.050	0.027	0.066
Frequency of AI/LLM use	0.008	0.015	0.615	-0.006	0.015	0.699
Art/poetry knowledge	0.026	0.025	0.286	0.067	0.025	0.007
Constant	0.325	0.150	0.031	0.597	0.146	0.000
Observations	523			538		
F-test p-value	0.470			0.082		

Notes: The dependent variable is a dummy for treatment status (*AI Info*). The table shows estimates from a linear probability model, standard errors and p-values. Age enters as band dummies with 18-24 year olds as the omitted category. The F-test is a test of the joint significance of all controls.

Table F.3: Sample representativeness: Online Experiments (Combined)

	Our sample (%)	Target (%)	Diff. (pp)
Age			
18-24	11.0	10.6	0.4
25-34	18.5	16.8	1.6
35-44	17.3	16.5	0.7
45-54	16.9	16.2	0.7
55+	36.4	39.9	-3.5
Region			
North East, England	4.4	4.1	0.3
North West, England	11.3	11.0	0.2
Yorkshire and the Humber, England	8.2	8.2	-0.0
East Midlands, England	7.5	7.3	0.2
West Midlands, England	9.2	8.8	0.4
East of England	8.3	9.4	-1.1
London	13.9	13.0	0.9
South East, England	13.3	13.8	-0.4
South West, England	8.3	8.7	-0.4
Wales	5.2	4.7	0.5
Scotland	8.2	8.3	-0.1
Northern Ireland	2.4	2.8	-0.4
Sex			
Female	50.3	51.6	-1.3
Male	49.7	48.4	1.3

Notes: The table compares the composition of the analytic sample for the online experiments with the recruitment quota targets. Diff. is the difference between our sample share and the target share, in percentage points.

Table F.4: Hypothesis Tests: Online Experiments

	AI Info mean	Human Info mean	Diff. in means	T-test p-value	K.S. statistic	K.S. p-value
H4: Artworks, AI Info vs Human Info	5.29	5.68	-0.39	0.191	0.093	0.207
H5: Haiku mugs, AI Info vs Human Info	4.94	5.64	-0.70	0.020	0.136	0.013

Notes: The table shows the mean bid (£) under each treatment condition, the difference in the means (*AI Info* minus *Human Info*), the p-value from a two-tailed t-test, the Kolmogorov-Smirnov statistic and p-value. Bids are averaged across the ten auction rounds at the participant level.

Table F.5: Alternative Estimators: OLS Regressions, Online Experiments

	Bid														
	Artworks					Haiku mugs					Pooled				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Panel A: Random effects															
AI Info	-0.391	-0.696**	-0.721**	-0.720**	-0.781***	-0.703**	-0.718**	-0.713**	-0.715**	-0.880***	-0.391	-0.660**	-0.703**	-0.702**	-0.729**
	(0.299)	(0.299)	(0.299)	(0.299)	(0.296)	(0.301)	(0.301)	(0.301)	(0.301)	(0.299)	(0.302)	(0.302)	(0.302)	(0.302)	(0.298)
AI Info × Mug											-0.311	-0.069	-0.015	-0.016	-0.111
											(0.424)	(0.424)	(0.424)	(0.424)	(0.420)
Panel B: Random effects, standard errors clustered on the bidder															
AI Info	-0.391	-0.696**	-0.721**	-0.720**	-0.781***	-0.703**	-0.718**	-0.713**	-0.715**	-0.880***	-0.391	-0.660**	-0.703**	-0.702**	-0.729**
	(0.298)	(0.299)	(0.299)	(0.299)	(0.294)	(0.301)	(0.301)	(0.301)	(0.301)	(0.298)	(0.298)	(0.298)	(0.298)	(0.299)	(0.294)
AI Info × Mug											-0.311	-0.069	-0.015	-0.016	-0.111
											(0.423)	(0.424)	(0.424)	(0.424)	(0.419)
Aesthetic rating (pre-test)		✓		✓	✓		✓		✓	✓		✓		✓	✓
Hypothetical WTP (pre-test)			✓	✓	✓			✓	✓	✓			✓	✓	✓
Individual controls					✓					✓					✓
Observations	5,250	5,250	5,250	5,250	5,230	5,410	5,410	5,410	5,410	5,380	10,660	10,660	10,660	10,660	10,610
Respondents	525	525	525	525	523	541	541	541	541	538	1,066	1,066	1,066	1,066	1,061

Notes: The dependent variable is the bid amount (£). The table re-estimates the specifications of Table 2 by OLS. Panel A reports random-effects estimates; Panel B reports the same estimates with standard errors clustered on the bidder. Only the treatment terms are shown; the rows beneath record which pre-test item scores and individual-level controls each column includes, and the samples and control sets are as in Table 2. The pooled columns include a mug dummy and its interaction with *AI Info*. Standard errors are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table F.6: Hypothesis Tests: online experiments, 18-34 year olds

	AI Info mean	Human Info mean	Diff. in means	T-test p-value	K.S. statistic	K.S. p-value
H4: Artworks, AI Info vs Human Info	6.23	6.38	-0.16	0.769	0.100	0.809
H5: Haiku mugs, AI Info vs Human Info	5.79	5.81	-0.03	0.963	0.113	0.727

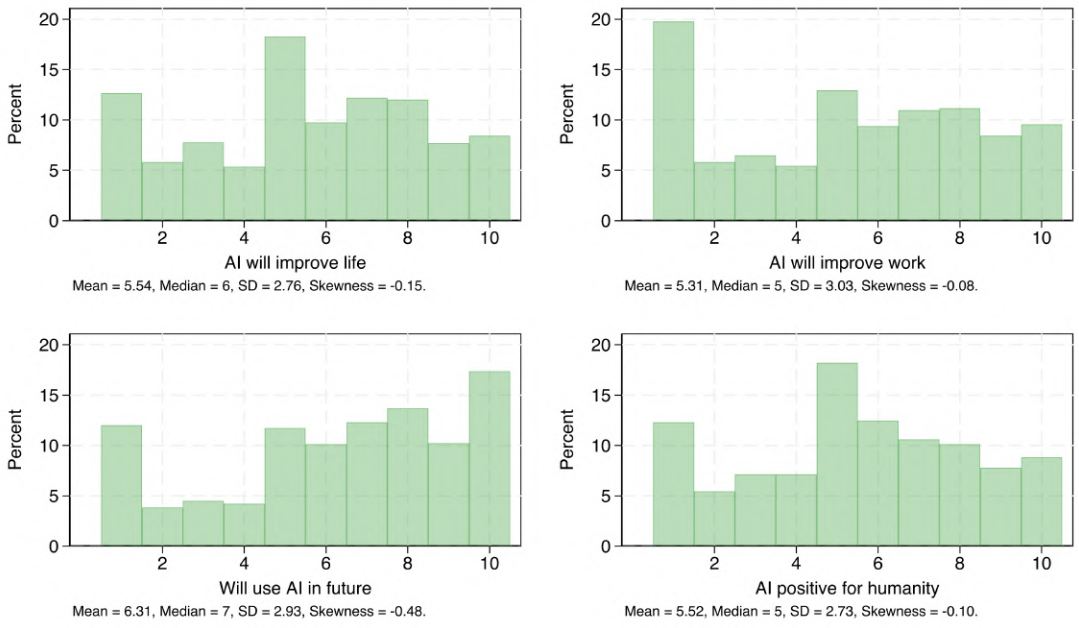
Notes: The table repeats the tests in Table F.4 for the subsample of participants aged 18 to 34. The table shows the mean bid (£) under each treatment condition, the difference in the means (*AI Info* minus *Human Info*), the p-value from a two-tailed t-test, the Kolmogorov-Smirnov statistic and p-value. Bids are averaged across the ten auction rounds at the participant level.

Table F.7: Learning Across the Ten Rounds: Online Experiments

	Bid	
	Artworks	Haiku mugs
	H4: Human Info vs AI Info	H5: Human Info vs AI Info
	(1)	(2)
Auction round	-0.064***	-0.016
	(0.017)	(0.014)
AI Info × Round	0.004	-0.011
	(0.024)	(0.019)
Subject fixed effects	✓	✓
Aesthetic rating (pre-test)	✓	✓
Observations	5,250	5,410
Subjects	525	541

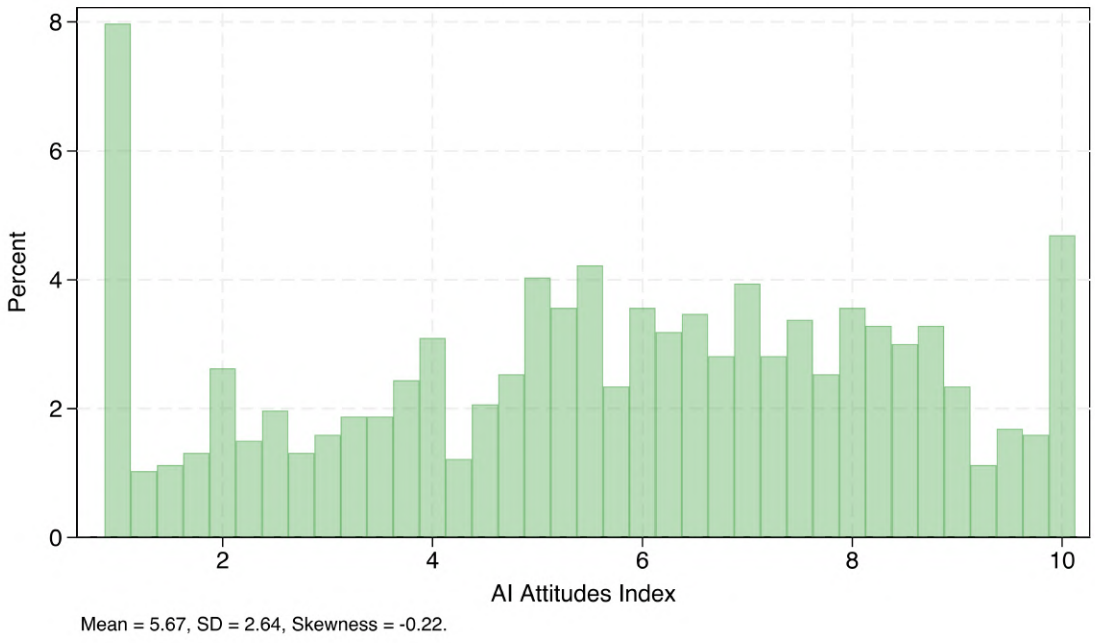
Notes: The dependent variable is the bid amount (£). The table shows estimates from linear regressions with subject fixed effects, estimated separately for each hypothesis. Column (1) uses the artwork experiment and column (2) the haiku mug experiment. All specifications control for the item's aesthetic rating, elicited in the Prolific pre-testing phase. Standard errors, clustered at the subject level, are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure F.1: The Distribution of the AI Attitudes Index Components: Online Experiments



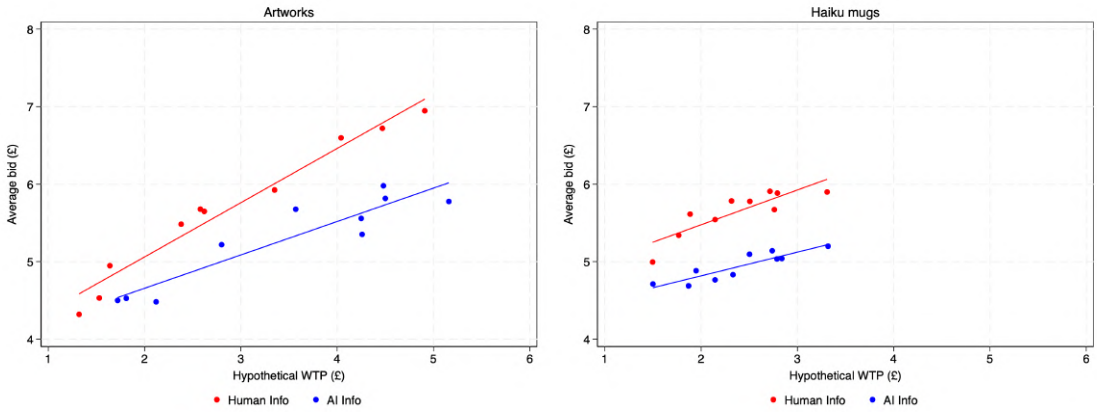
Notes: The figure shows the distribution of the four components of the AI Attitudes index. All responses range from 1 to 10, where higher values indicate a more positive view, or frequent future use, of AI.

Figure F.2: The Distribution of the AI Attitudes Index: Online Experiments



Notes: The figure shows the AI Attitudes index, which is the average of four components.

Figure F.3: Relationship Between Hypothetical WTP and Auction Bids by Treatment: Online Experiments



Notes: Each point shows the mean bid (£) and hypothetical WTP for one of the 10 artworks or 10 haiku mugs per treatment condition. The hypothetical WTP measure is the average willingness to pay (£) for each item elicited in the Prolific pre-testing phase. Lines show linear fits. The left panel compares bids for artworks when the source is AI vs human (H4) and the right panel compares bids for haiku mugs when the source is AI vs human (H5).

G References

Celebi, Can, Christine Exley, Soren Harrs, Hannu Kivimaki, Marta Serra-Garcia and Jeffrey Yusof. 2026. "Mission possible: The collection of high-quality data." *CESifo Working Paper No. 12535* .

Grassini, Simone. 2023. "Development and validation of the AI attitude scale (AIAS-4): A brief measure of general attitude toward artificial intelligence." *Frontiers in Psychology* 14.