Supplementary Materials

The Psychology of Reification: Perceiving Social Reality as Mind-Independent

Study 1

Methods

Participants

Participants were recruited from the Prolific platform, comprising 301 English-speaking U.S. residents (156 men and 145 women). Eligibility criteria included English as a first language, and declaring fluent language, a minimum of 50 prior submissions on the Prolific platform, and an approval rate exceeding 95%. Participants provided informed consent prior to participation.

Participants answered four preliminary clarification questions related to the mind-dependence concept to ensure comprehension. Participants who failed to demonstrate adequate understanding were excluded from the subsequent analysis. Likewise, participants that failed one or more of the attention checks were also excluded. Of the 301 initial participants, 32 failed the attention checks and an additional 53 did not correctly answer all four clarification questions. These participants were excluded from further analysis, leaving a final sample of 216 participants.

Procedure

Participants responded to 18 statements regarding various social facts, such as "Overweight people face discrimination" and "Flowers are used for decoration", presented in a randomized order. Each statement was followed by a series of evaluations: (1) the truthfulness of the statement, (2) confidence in the truthfulness of the assessment, (3) five questions about the perception of the statement's mind-dependence, and (4) naturalness, stability, and inherency to our world.

Upon completion of the statement evaluations, participants completed the 12 Item Social and Economic Conservatism Scale (SECS; [1]) questionnaire, a brief version of the Big Five Inventory consisting of 10 items (BFI; [2]), and demographic questions.

Analysis

Under the hypothesis that mind dependence constitutes a distinct construct from those forming the basis of essentialism, our study hypothesizes that evaluations of social facts concerning their mind dependency, naturalness, stability, and inherency to our world will manifest as two separate constructs: mind-dependence and essentialist traits (i.e., naturalness, stability, and inherency). Mind-dependence is an attribution specific to each individual and statement; thus, to mitigate idiosyncratic influences, we standardized the data relative to both

participants and statements. Subsequently, a factor analysis was conducted to test this hypothesis.

Further, we employed Hierarchical Linear Modeling (HLM) approach to examine intraclass correlations (ICCs), thereby illuminating variances across different statements and between participants. Model 1 was specified as follows:

$$Mind\text{-}Dependence \sim 1 + (1 \mid Person) + (1 \mid Statement) + (1 \mid Question).$$

The interaction between person and statement will reveal the varying patterns of mind-dependence attribution across different individuals and statements. To investigate this effect, we compared Model 1 with Model 2, where Model 2 also incorporates an interaction term:

$$Mind\text{-}Dependence \sim 1 + (1 \mid Person) + (1 \mid Statement) + (1 \mid Person{:}Statement) + (1 \mid Question).$$

We used a base factor to compare the two models.

Additionally, to demonstrate that mind-dependence is not perceived in a binary manner but rather exhibits a gradation, we visualized the distribution of mind-dependency attributions.

The hypotheses and the detailed analysis plan have been pre-registered and are accessible at: https://aspredicted.org/p924-kcxb.pdf.

Results

Factor analysis

The suitability of the data for factor analysis was initially confirmed by Bartlett's test of sphericity, which indicated a significant correlation within the data (χ^2 (28) = 14,145.99, $p < .001$). The exploratory factor analysis, using various methods to determine the number of factors, consistently suggested the extraction of two factors. Specifically, the scree test results from the acceleration factor, the optimal coordinates, parallel analysis, and Kaiser criterion all supported a two-factor model.

Rotated component matrix, using an oblimin rotation, revealed distinct loadings for the two factors, which together explained 56.50% of the total variance in the data (Figure S1). Aligned with our hypothesis, the first factor (mind-dependence), accounting for 37.90% of the variance, strongly loaded on Believe (0.84), Think (0.83), Want (0.82), Desire (0.77), and counter-factuality (0.61). The second factor (essentialism) explaining 18.6% of the variance, included significant loadings on Inherency (0.78), Naturalness (0.71), and Stability (0.60). The inter-factor correlation was (r = -0.10)

Hierarchical Linear Modelling (HLM) Analysis

The intra-class correlation coefficients (ICCs) of model 1 revealed substantial variability attributed to individual differences, with an ICC of 0.316 for Person, indicating that approximately 31.60% of the variance in mind-dependence scores could be attributed to differences between individuals. The variance attributed to statement differences was notably smaller, with an ICC of 0.080, accounting for about 8.00% of the variance. Further examination suggested significant contributions of both random effects to the model. Removing the random effect of Person from the model resulted in a substantial increase in the Akaike Information Criterion (AIC) from 79,403 to 86,706, with a likelihood ratio test (LRT) statistic of 7304.60 (p <.001). Similarly, excluding the random effect of Statement led to an increase in AIC to 81,596 and an LRT statistic of 2195.40 (p < .001), both indicative of the significance of these components in the model.

To assess the significance of the interaction between individual participants and statements, we compared Model 1, which did not include an interaction term, with Model 2, which did. The ICC showed that this interaction (Person:Statement) accounted for 31.70% of the total variance, underscoring its substantive influence alongside the contributions from individual participants (30.10%), and Statements (7.80%). The comparison of Model 1 with Model 2 revealed a significant improvement in model fit with the inclusion of the interaction term χ^2 (1) = 6585.4, p < .001. The AIC and Bayesian Information Criterion (BIC) were notably lower for Model 2; AIC = 72818, BIC = 72865 compared to Model 1 where the AIC = 79401 and the BIC = 79441. Additionally, the bayes factor (BF), computed using the BIC approximation, strongly favored the model including the interaction term, with a BF < .001, against Model 1. This substantial BF indicates a strong preference for the model accounting for interactions between individuals and statements over the model without these interactions.

Figure S1. Factor Analysis of Mind-dependence and Essentialism of Social State-of-Affairs

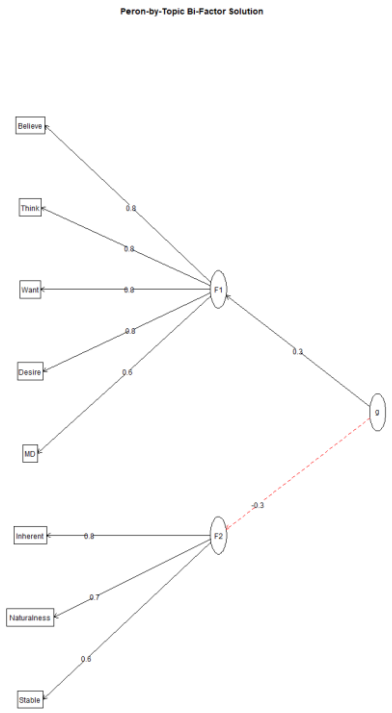

Note. Factor analysis method used oblimin rotation.

Study 2

Methods

Participants

A total of 250 participants were initially invited to complete the mind-dependence questionnaire that included ratings across 10 distinct topics, presented in randomized order. Participants who failed an attention check or did not respond appropriately to clarification questions designed to assess understanding of the key questions, were excluded from further participation. After these exclusions, a total of 182 participants were retained and invited to participate in a second session one week later, during which they completed the same questionnaire. To the second session, 148 participants responded. Following the same exclusion criteria as in the first session, an additional 21 participants were excluded, resulting in a final sample of 127 participants who completed both sessions.

Procedure

The study was designed to evaluate the temporal stability of mind-dependence judgments using a one-week test–retest design. Participants first completed a clarification phase to ensure comprehension of the mind-dependence questions. This phase included a set of conceptual clarification questions identical to those used in the previous study. Participants who failed to correctly answer any of these items were excluded in accordance with the pre-registered criteria.

At Time 1, participants evaluated ten social state-of-affairs selected to represent a diverse sample of social facts. The ten statements were presented in randomized order for each participant. For every statement, participants provided: (a) a truth assessment; (b) a confidence rating; (c) five items measuring perceived mind-dependence (e.g., whether the truth of the statement depends on what people think, believe, desire, or want). During this first session, participants also completed a brief attention-check to ensure engagement.

One week later, all eligible participants were re-contacted to complete a second session (Time 2). In this session, participants were presented with the same ten statements, in newly randomized order, and completed the identical sequence of ratings. No feedback was provided regarding their prior responses, and participants were not informed that the goal of the study was to assess response consistency across time. As in the first session, participants who failed clarification items or the attention check were excluded according to the preregistered rules.

Analysis

All analyses were pre-registered (see, https://aspredicted.org/xv4u9p.pdf). To assess the test-retest reliability of the mind-dependence ratings, we employed two complementary analytic strategies. First, we computed Pearson correlations between Time 0 and Time 1 ratings separately for each of the 10 topics. Second, we estimated generalizability coefficients using hierarchical linear modeling (Shrout & Lane, 2012), enabling us to partition the variance in mind-dependence judgments into person-level, topic-level, and interaction-level components across time. Specifically, we calculated: (a) an overall reliability coefficient reflecting consistency in individual responding across all topics and timepoints; and (b) a coefficient reflecting the reliability of topic-specific differential responding by each participant.

Results

Test-Retest Correlations

Correlations between Time 1 and Time 2 ratings for individual topics ranged from r = 0.475 (for the item: "Females are the primary caregivers for the child") to r = 0.725 (for the item: "When a person dies a funeral is held"). The mean test-retest correlation across all topics was r = 0.640, suggesting a moderate-to-high degree of temporal stability in participants' ratings (see Table 1S for all topic-wise correlations).

Table S1. Test-Retest Correlations by Topic

State of affairs	Test-Retest Correlation (r)
Age	0.639
DIE	0.725
DOG	0.696
FIS	0.611
FM	0.475

FW	0.688
ISR	0.661
OWP	0.647
SKI	0.581
TWI	0.676

Note. Values represent Pearson correlation coefficients between Time 0 and Time 1 ratings for each topic.

Hierarchical Linear Modeling and Generalizability Coefficients

Variance components estimated from the hierarchical linear model indicated substantial between-person variability (*SD* = 1.35), with additional variability attributable to topic-level (*SD* = 0.59) and participant-by-topic interaction effects (*SD* = 0.79). The variance component for topic-by-time interactions was negligible (*SD* = 0.00), suggesting minimal systematic change at the topic level over time (see Table S2 for all results).

Based on these estimates, we computed generalizability coefficients (R_{1R}) at three levels of analysis. The coefficient for participant-level reliability ($R_{1R\text{-}participant} = 0.797$) indicates high stability in individuals' overall mind-dependence judgments across topics and time points. The coefficient for State-of-affairs-level reliability ($R_{1R\text{-}state\text{-}of\text{-}affirs} = 0.890$) reveals a similarly high degree of consistency in how different topics are judged across participants and time, underscoring the reliability of State-of-affairs-level. Finally, the coefficient for participant-by-State-of-affairs interactions ($R_{1R\text{-}interaction} = 0.305$) reflects a more modest level of stability in how individuals differentially judge specific topics over time.

Table S2. Hierarchical Linear Model Estimates for Mind-Dependence Scores

Effect type	**Parameter**	**Estimate**	**SE**	**95% CI**	**t**	***p***
Fixed effect	Intercept	4.26	0.27	[3.74,4.78]	15.98	<.001
						SD
Random effect	Intercept (ID)					1.35
	Intercept (State-of-affairs)					0.59
	Intercept (Time)					0.2
	Intercept (ID: State-of-affairs)					0.79
	Intercept (ID:Time)					0.58
	Intercept (State-of-affairs:Time)					0.00
	Residual					1.01

Note. The fixed effect reflects the average mind-dependence score across all topics and time points. Random effects capture variability across participants, topics, and time, including cross-level interactions. Confidence intervals and p-values for fixed effects were computed using a Wald t-distribution approximation.

Study 3

Methods

Participants

This study involved 150 participants (63 men and 87 women) recruited through the Prolific platform, all of whom were from the United States with English as their primary language. Eligibility criteria required participants to have at least 50 previous submissions on Prolific with an approval rate exceeding 95%.

After providing informed consent, participants completed four initial clarifying questions to confirm their understanding of the study tasks. Those who failed to demonstrate adequate comprehension were excluded from further analysis. Then, participants continued to the main task. In addition, participants who failed one or more attention checks were excluded from the analysis. Of the 150 initial participants, 16 failed the attention checks and an additional 30 did not correctly answer all four clarification questions. These participants were excluded from further analysis, leaving a final sample of 104 participants

Procedure

The primary task involved presenting participants with 20 statements, corresponding to 10 topics, in a randomized order. Each topic (e.g., Dogs) featured a pair of statements: one reflecting a natural fact (e.g., "Dogs can't fly") and another representing a social fact (e.g., "Dogs are pets"; see Table S3). Participants responded to each statement by firstly, stating whether the statement is true, reporting their confidence in their response, and subsequently evaluating mind-dependence and essentialism. Following the main task, participants completed the 12-item Social and Economic Conservatism Scale (SECS; [1]), the short 10-item Big Five Inventory (BFI; [2]), and a demographic questionnaire.

Table S3. Social and Natural state of affairs

N	Topic	Social fact	Natural Fact
1	Overweight	Overweight people face discrimination	Overweight people face health risks
2	Dying	When a person dies a funeral is held	When a person dies their brain stops functioning
3	Females	Females are the primary caregivers for the child	Females give birth
4	Flowers	Flowers are used for decoration	Flowers cannot bloom without water
5	Dogs	Dogs are pets	Dogs can’t fly

6	Identical twins	Identical twins cannot marry each other	Identical twins have the same genes
7	Fish	Fish are eaten at restaurants	Fish can breathe underwater
8	Age of 18	The legal voting age in the US is 18	Puberty occurs before the age of 18
9	Skin Color	People often stereotype others based on their skin color	Skin color varies between people, and is mostly determined by genetics
10	Israel	Israel is a UN member	Israel is located in the middle east

Analysis

We employed Hierarchical Linear Modeling (HLM) approach to examine the differences between the mind dependency attributed to social facts compared to natural facts (Fact-type). The model was set as follows:

Mind-Dependence ~ 1 + Fact-type + (Fact-type| Person) + (Fact-type | Topic) + (Fact-type | Person:Topic) + (Fact-type | Question).

The fixed effect of Fact-type (Social versus Natural) was used to estimate the hypothesis that people perceive social facts as significantly more mind-dependent than natural facts.

Further, to explore the hypothesis that the variance of mind-dependence attribution is greater for social facts than for natural facts, both within topics (i.e., there is a greater variance in mind dependency for social versus natural facts per topic), and within person (i.e., people have greater variance in the mind-dependency they attribute for social facts as opposed to natural facts), we utilized the HLM model and bootstrapping using a parametric method and 200 resamples of the data to estimate the different variances, and to compare them. The analysis plan and all hypotheses have been pre-registered and can be found at: https://aspredicted.org/hrz2-9b84.pdf.

Results

To investigate our first hypothesis that social facts are perceived as more mind dependent then natural facts, we conducted Hierarchical Linear Model (HLM) to examine the impact of Fact-type (i.e., social facts versus natural facts) on Mind-dependency, accounting for and exploring individual variations and topic-specific effects. Model 1 was specified with random intercepts and slopes for Fact-type across participants (Person), topics (Topics), and the

interaction between participants and topics, as well as by the different mind-dependency questions (Question). Model's 1 specification was as follows:

Mind-Dependency ~ 1 + Fact-type + (Fact-type| Person) + (Fact-type | Topic) + (Fact-type | Person:Topic) + (Fact-type | Question).

Variance analysis using Satterthwaite's method revealed significant effects of Fact-type $F(1, 24.685) = 76.845$, $p < .001$. The fixed effects analysis showed that the intercept, representing the baseline response for natural facts, was significant $b = 1.94$, $SE = 0.16$, $t(44.26) = 11.96$, $p < .001$. Additionally and in accordance with our hypothesis, social facts had a significantly higher mind dependency score, $b = 2.58$, $SE = 0.29$, $t(24.69) = 8.77$, $p < .001$ compared to natural facts. In accordance, using model 1 predicted mean mind-dependency score of $M = 1.94$, $p < .001$, 95% CI [1.62,2.25], for natural facts, and $M = 4.52$, $p < .001$, 95% CI [4.09,4.95], for social facts.

Bootstrap Analysis

To test the hypothesis that the variance of mind dependency for social facts exceeds that for natural facts, we conducted a variance comparison of variances extracted from the model above. Subsequently, bootstrap methods were employed (parametric method, using the bootMer function from the lme4 package), involving 200 parametric simulations to estimate the variances. The bootstrap analysis supported our hypothesis, confirming the greater variability of social facts both within topics and within individuals. The mean difference in variance within individuals was $M_{\text{difference}} = 1.37$, $p < .005$, 95% CI [0.858, 1.95], and within topics was $M_{\text{difference}} = 1.23$, $p = .059$, 95% CI [0.12, 2.63]. Distributions of bootstrapped variances are presented in Figure S2.

Figure S2. Distribution of bootstrapped variances and 95% quantile CIs

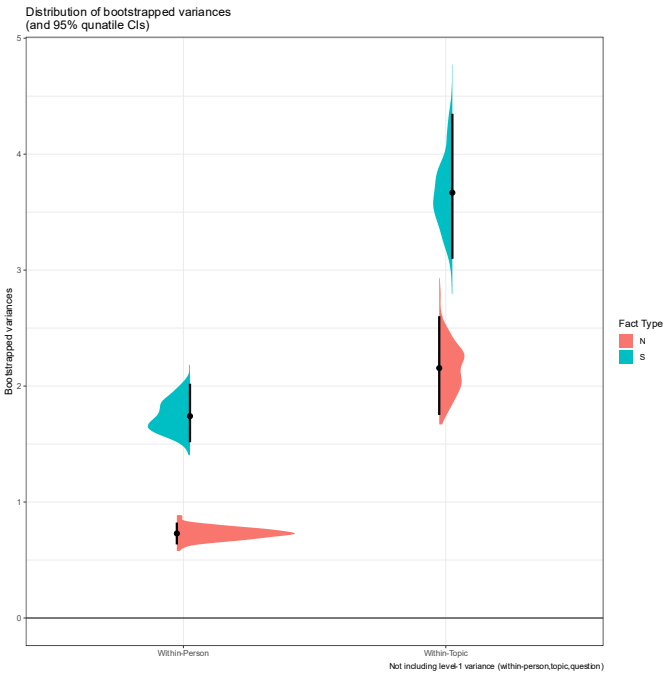

Note. 200 resamples were used.

Study 4

Methods

Participants

A total of 500 English-speaking participants (205 men, 295 women) were recruited through the Prolific platform. Eligibility criteria required at least 50 prior submissions and an approval rate above 95%. Participants who failed one or more attention checks (*n = 70*) or the clarification questions (*n = 69*, 16 of whom failed both) were excluded, leaving a final analytic sample of 377 participants. All participants provided informed consent prior to participation.

Procedure

Each participant was randomly assigned one state of affairs out of twenty, representing ten topical pairs in which each topic contained both a *social* and a *natural* version (e.g., *"Overweight people face discrimination"* vs. *"Overweight people face health risks"*). The complete set of topics is presented in Table S3. Participants were first instructed to explain their assigned state-of-affairs in their own words using a minimum of 50 words. Immediately afterward, they completed the mind-dependence questionnaire for that same state-of-affairs, consisting of five items assessing the extent to which the statement's truth depends on what people think, believe, want, or desire.

Coding of Explanations. Each participant's open-ended explanation was analysed using GPT-4o via the OpenAI Assistants API to extract continuous stance scores corresponding to Dennett's three stances: *Physical*, *Design*, and *Intentional*. Scoring followed the pre-registered rubric (https://aspredicted.org/6jg7xj.pdf). Each stance was rated on a 1–7 scale, independently of the others.

The development of the stances prompt

To generate reliable stance scores for participants' explanations, we developed a structured prompt for GPT-4o through an iterative calibration process. The goal was to ensure that the model could accurately identify the degree to which each explanation reflected Dennett's (1987) three cognitive stances—Physical, Design, and Intentional—when applied to short, open-ended causal accounts of social and natural state-of-affairs.

The development of the prompt proceeded in several stages. We first constructed a validation dataset consisting of 330 think-aloud explanations produced by 10 participants in response to "what are the causes?" prompts for 33 topics (17 social, 16 natural). Using an initial subset of 20 explanations, we iteratively refined the model's

instructions by providing GPT-4o with detailed explanations of each stance, testing preliminary ratings, and clarifying conceptual boundaries when outputs diverged from theoretical definitions. This process yielded an initial set of structured coding instructions that defined each stance, outlined behavioral and linguistic indicators, and specified a 1–7 scoring scale for each stance independently. All model runs used temperature set to 0 to ensure deterministic output.

To evaluate the clarity of the initial instructions, six trained research assistants independently coded the human-produced explanations. Each excerpt was coded by two raters, with each rater evaluating approximately one-third of the dataset. Inter-rater reliability, assessed via two-way random-effects intraclass correlation coefficient $ICC_{(2,1)}$, was moderate to good for the Physical (ICC = 0.70) and Intentional (ICC = 0.60) stances, but low for the Design stance (ICC = 0.20), reflecting conceptual challenge in distinguishing functional reasoning from mechanistic or intentional explanations.

We therefore conducted a calibration phase that included individual debriefing sessions with each rater, a group consensus meeting to discuss ambiguous cases, and the creation of four anchor examples exemplifying consensus ratings for each stance. These refinements substantially improved reliability in a cross-validation round in which excerpts were redistributed and coded anew by different pairs of raters. Reliability increased to $ICC_{(2,1)} = 0.79$ for Physical, 0.68 for Intentional, and 0.45 for Design, indicating successful operationalization of the framework.

The final version of the coding instructions including the four anchor examples was then incorporated into a prompt used with GPT-4o for automated stance coding. Because the Design stance remained the most difficult to identify consistently, we explored several optimization strategies (e.g., varying phrasing, reordering instructions, adjusting structural cues), but none improved performance beyond including the calibrated anchors. API ratings were validated against averaged human ratings using the same $ICC_{(2,1)}$ metric employed for inter-rater reliability, treating API ratings and averaged human ratings as two independent raters. We validated API performance on different random subsets of the data to ensure robustness. This yielded: Physical $ICC_{(2,1)} = 0.82$ (95% CI [0.62, 0.90], $p < .001$), Intentional $ICC_{(2,1)} = 0.74$ (95% CI [0.57, 0.83], $p < .001$), and Design $ICC(2,1) = 0.53$ (95% CI [0.45, 0.60], $p < .001$).

API-human agreement was comparable to or exceeded inter-human reliability across all three stances. For the challenging Design stance, the API achieved ICC =

0.53 compared to human inter-rater ICC = 0.45, demonstrating that automated coding-maintained reliability even for conceptually ambiguous distinctions. Pearson correlations between API and averaged human ratings confirmed these patterns: Physical r = 0.87 ($p < .001$), Intentional r = 0.79 ($p < .001$), Design r = 0.53 ($p < .001$).

Analysis

As pre-registered, the primary analysis examined whether adoption of the intentional stance predicted perceived mind-dependence. To account for variability both between and within states-of-affairs, a hierarchical linear model (HLM) was estimated with the following specification:

Mind-Dependence ~ Intentional-Stance-Between + Intentional-Stance-Within + (Intentional-Stance-Within | State-of-Affairs)

All analyses were conducted in R using *lmeTest* and *parameters* packages. The simple Pearson correlations among stance variables were computed as descriptive statistics only and were not part of the pre-registered hypothesis.

Results

The final dataset comprised 189 social and 188 natural explanations ($N = 377$). mind-dependence ratings were consistently higher for social compared with natural statements, replicating the pattern observed in earlier studies. Number of partisans per state-of-affairs, and means and standard deviations of mind-dependence ratings are summarized in Table S4. Figure 1 displays the distribution of mind-dependence scores across topics, further illustrating that social facts were judged as markedly more mind-dependent than natural ones.

Table S4. Mean (SD) of Mind-Dependence Scores across Topics and State-of-Affairs Type

Topic	*Mean Natural*	*Mean Social*	*SD Natural*	*SD Social*	*N Natural*	*N Social*
Age of 18	2.65	3.15	1.70	1.90	19	16
Dogs	1.28	4.09	0.71	1.94	17	19
Females	1.34	3.56	0.87	1.82	24	19
Fish	1.40	3.36	0.80	1.86	23	18

Table S4. Mean (SD) of Mind-Dependence Scores across Topics and State-of-Affairs Type

Topic	*Mean Natural*	*Mean Social*	*SD Natural*	*SD Social*	*N Natural*	*N Social*
Flowers	1.74	4.12	1.71	1.75	18	18
Identical twins	1.71	3.82	1.16	1.90	16	18
Israel	2.27	4.04	1.69	2.10	18	15
Overweight people	2.38	4.30	1.83	1.50	17	24
Skin color	1.77	3.90	1.37	1.98	19	23
When a person dies	1.74	4.03	1.25	1.96	17	19

Figure 2. Distribution of Mind-Dependence Scores by Topic and Fact Type

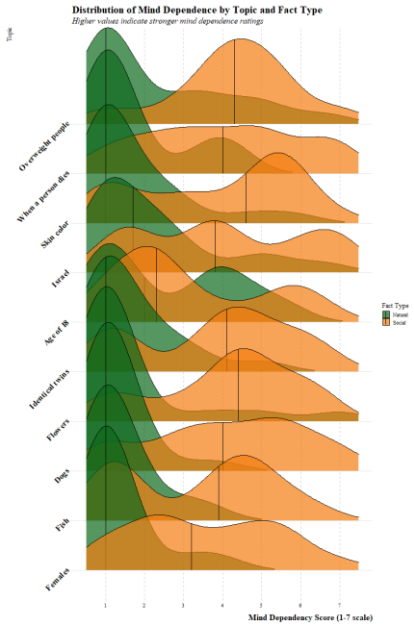

337 *Descriptive Correlations.* For descriptive purposes, we also examined simple
338 Pearson correlations among stance and MD scores (Figure SS).

Figure S3. Correlation matrix of stance scores and Mind-Dependence ratings.

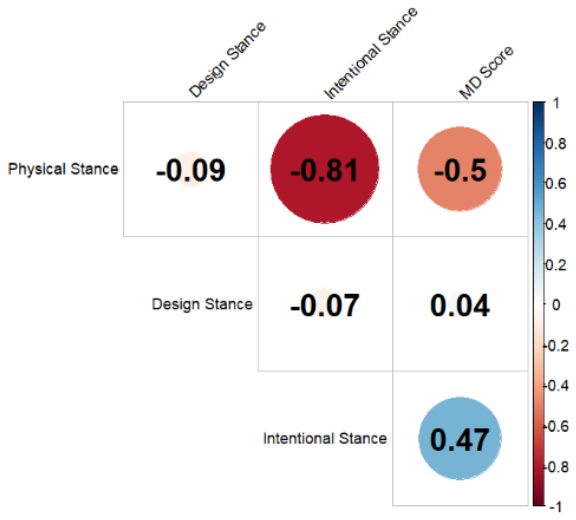

Note. Circle size and color indicate the magnitude and direction of Pearson correlations (blue = positive, red = negative).

339 *Hierarchical Linear Model.* Consistent with the pre-registered hypothesis, the
340 model revealed that adoption of the Intentional Stance positively predicted Mind-
341 Dependence (Table 2). Participants who tended to use more intentional reasoning
342 across topics also tended to judge the corresponding facts as more mind-dependent.

The between–state-of-affairs effect of Intentional Stance was significant (β = 0.51, SE = 0.06, 95% CI [0.39, 0.63], $p < .001$). The within–state-of-affairs effect was smaller but remained significant ($\beta = 0.09$, $SE = 0.05$, 95% CI [0.00, 0.18], $p = .046$).

Table S5. Hierarchical linear model predicting *Mind-Dependence* from *Intentional Stance*

Parameter	**Coefficient**	**SE**	**95% CI**	***t***	***p***
(Intercept)	0.00	0.06	[–0.12, 0.12]	0.04	.968
Intentional Stance (between)	0.51	0.06	[0.39, 0.63]	9.04	< .001
Intentional Stance (within)	0.09	0.05	[0.00, 0.18]	2.05	.046

Study 5

Methods

Participants

The target sample for this study was 1200 participants. A total of 1628 English-speaking adults were recruited through the Prolific platform. Eligibility criteria required participants to have completed at least 50 previous submissions with an approval rate above 95%. Participants who failed one or more attention or comprehension checks were replaced, in accordance with the pre-registered exclusion protocol. In total, 428 participants failed at least one of these checks, leaving a final eligible sample of 1200 participants who successfully completed the study. All participants provided informed consent prior to participation. The study's pre-registration can be accessed at: https://aspredicted.org/h5dm37.pdf.

Procedure

Participants first completed a consent form and were introduced to the general purpose of the study. Each participant was randomly assigned to one of two experimental conditions: an *intentional stance* condition or a d*esign stance* condition. In both conditions, participants were presented with one of 20 distinct states of affairs. These 20 items were selected based on the data collection conducted in *study 6*, from which the states of affairs rated as most mind-independent were chosen for use in the current study. Each state-of-affairs was followed by a short explanatory paragraph, written in either an intentional or design frame in respect to the assigned condition. To ensure engagement, participants were asked to describe in their own words the most compelling argument from the explanation text. Following this, participants responded to three items measuring perceived unthinkability of the state-of-affairs, each rated on a 1–7 Likert scale. Next, participants completed two clarification questions and then completed five items measuring perceived mind-dependence of the state-of-affairs, again rated on 1–7 scales. Finally, participants provided basic demographic information before completing the study.

Results

Main Analysis.

Participants in the intentional stance condition reported higher perceived mind-dependence (M = 3.24, SD = 1.82) than those in the design stance condition (M = 3.01, SD = 1.84). The mixed-effects model confirmed that this difference was statistically significant, B = 0.23, SE = 0.11, 95% CI [0.01, 0.46], $t(18.26) = 2.15$, $p = .045$. This suggests that explanations framed in terms of intentional stance increase the perceived mind-dependence of a state-of-affairs. In contrast, the effect of condition on perceived unthinkability was not statistically significant. Although the design stance condition yielded slightly higher unthinkability scores (M = 4.87, SD = 1.50) compared to the intentional stance condition (M = 4.76, SD = 1.46), the

estimated effect was small and nonsignificant, B = -0.06, SE = 0.10, 95% CI [–0.28, 0.16], t(18.54) = –0.56, p = .583.

Mediation Analysis.

To examine whether mind-dependence mediates the effect of condition (intentional vs. design) on judgments of unthinkability, we conducted a multilevel mediation analysis with facts as a grouping variable. The analysis included 1,200 observations nested within facts. The mediator model regressed within-fact mind dependence attribution on condition. The outcome model regressed unthinkability judgments on condition, the mediator, and the fact-level mean mind dependence, with random slopes for condition and the mediator across facts. The average causal mediation effect (ACME) was small and not significant (ACME = -0.019, 95% CI [-0.047, 0.001], p = .072), indicating that condition did not significantly affect unthinkability through differences in within-fact mind dependence attribution.

Study 6A

Methods

Participants

A total of 1,069 English-speaking participants were recruited through the Prolific platform. All participants were required to have completed at least 50 prior submissions with an approval rate exceeding 95%. 159 participants failed attention checks and were excluded, yielding a final eligible sample of 910 participants. All participants provided informed consent before beginning the study.

Procedure

This study aimed to construct a comprehensive dataset of social state-of-affairs and to map their perceived structure across six conceptual dimensions of social reality. First, a corpus of 300 social facts (e.g., *"People thank others after receiving help"*) was generated. Social facts were designed to correspond systematically to the first 100 categories of the Dewey Decimal Classification System, ensuring broad coverage of domains that characterize the social reality. To fill gaps in topical representation, additional social facts were derived from a taxonomy of university faculties, resulting in a balanced set of 300 social facts.

In the first phase, participants were invited to rate subsets of these facts on truthfulness, confidence in truthfulness, and one of the six social dimensions (described in the measurement section). Each participant evaluated 50 social facts randomly assigned from the full list. Once each social fact reached approximately 30 ratings, facts with less than 90% agreement regarding their truth value were removed, leaving a refined corpus of 233 validated social facts. In the second phase, the remaining participants rated 50 randomly selected social facts (from the 233 list) on one of the six dimensions without providing truthfulness ratings. All procedures were approved by Tel-Aviv University's Institutional Review Board.

Measures

Each social fact was evaluated on one of the following six dimensions, presented with a 7-point scale and examples to ensure consistent interpretation:

Institutional Origin and Support. Assesses dependence on formal institutions for existence or maintenance. (7 = Exists only because of institutional rules or authority; 1 = Exists independently of institutions.) Example (7): "Citizens are issued identification documents by their government." Example (1): "People bring gifts when visiting someone's home."

Hierarchical Function. Assesses whether the fact reflects or reinforces social hierarchies or unequal power. (7 = Strongly reflects hierarchy; 1 = No implication of power or status.) Example (7): "CEOs make much more money than line workers." Example (1): "Coworkers chat with each other during breaks."

Sociopolitical Sensitivity. Captures the degree to which the fact involves politically or emotionally charged topics. (7 = Highly sensitive or controversial; 1 = Broadly accepted and

neutral.) Example (7): “Many countries allow abortion.” Example (1): “People usually wait their turn in shared spaces.”

Embedded Moral Valence. Measures whether the fact involves behavior or conditions treated as morally significant. (7 = Commonly moralized; 1 = Not associated with moral evaluation.) Example (7): “People are often criticized for being unfaithful to their partners.” Example (1): “People typically eat together during major celebrations.”

Scope of Application. Assesses how broadly the fact applies within society. (7 = Applies universally; 1 = Applies narrowly.) Example (7): “All citizens are required to pay taxes.” Example (1): “Surgeons must scrub their hands for a specific duration before an operation.”

Intentionality and Collective Representation. Evaluates whether the fact exists due to collective agreement or emerges unintentionally from individual behavior. (7 = Upheld by explicit collective belief; 1 = Unintended behavioral pattern.) Example (7): “A specific day is a national holiday.” Example (1): “Cars slow down to look at an accident in another lane, causing a traffic jam.”

Results

A mean of 27.44 participant rated the truthfulness of the social facts in the fist phase. Across dimensions, Intentionality and Collective Representation and Scope of Application yielded the highest mean ratings. Sociopolitical Sensitivity received the lowest average rating. The mean number of raters per dimension per social fact ranged from 24 to 37, reflecting balanced sampling across the six dimensions (see Table S6).

Table S6. Descriptives of Dimensions for 233 Facts Rated by 910 Participants

Dimension	N	Mean	SD
Intentionality and Collective Representation	36.47	4.81	0.85
Hierarchical Function	26.16	3.37	1.05
Embedded Moral Valence	26.64	3.53	1.22
Sociopolitical Sensitivity	25.31	2.25	0.89
Scope of Application	25.84	4.90	0.98
Institutional Origin and Support	23.61	3.66	1.71

Note. N is the mean number of raters per social fact.

Interrelations Among Social Dimensions

To explore the relationships among the six dimensions, we computed Pearson correlations based on mean ratings for each of the 233 validated social facts (Figure S4). Institutionalization was strongly correlated with intentionality (r = .66). Hierarchical Function correlated moderately with institutionalization (r = .44); and with moral valence and

sociopolitical sensitivity (both; r = .34). Moral valence was positively related to sociopolitical sensitivity (r = .60).

Figure S4. Correlation matrix of mean dimension scores across 233 social facts.

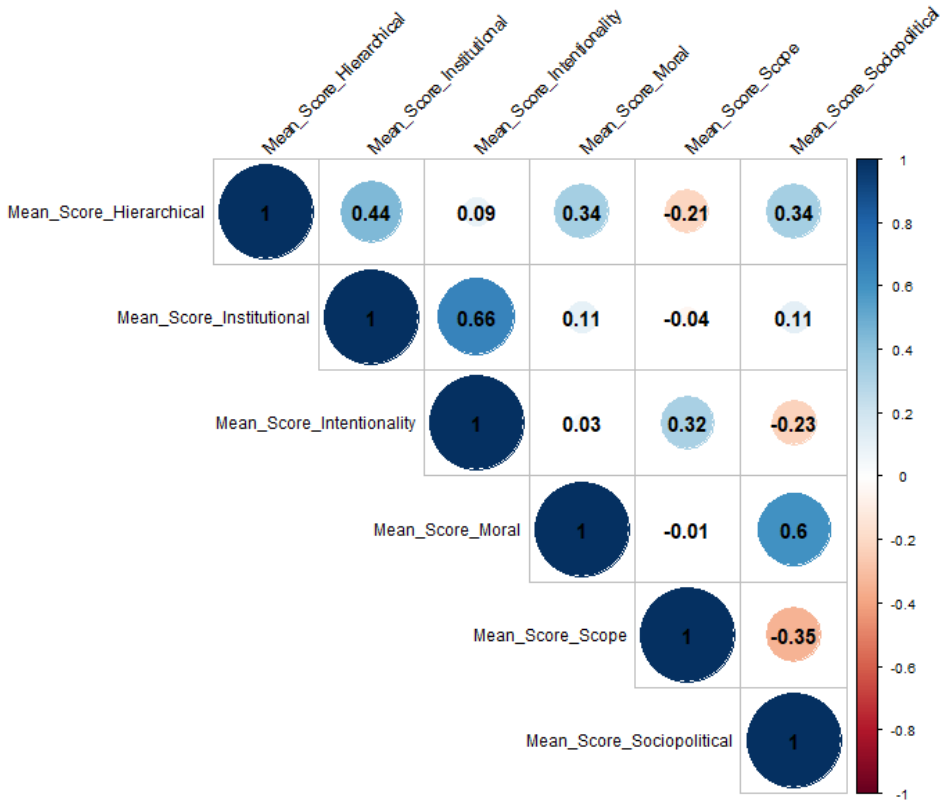

Note. Circle size and color indicate the magnitude and direction of Pearson correlations (blue = positive, red = negative).

Studies 6B and 7

Methods

Participants

The target sample for this study was 1,200 participants. A total of 1,719 English-speaking adults were recruited through the Prolific platform. Eligibility criteria required participants to have completed at least 50 previous submissions with an approval rate above 95%. Participants who failed one or more attention or comprehension checks were replaced, in accordance with the pre-registered exclusion protocol. In total, 519 participants failed at least one of these checks, leaving a final eligible sample of 1,200 participants who successfully completed the study. All participants provided informed consent prior to participation.

Procedure

Participants were first introduced to the study's objectives and completed a brief clarification phase designed to ensure comprehension of the task. This phase included two example scenarios, each followed by clarification questions that verified participants understood the type of evaluations they were about to perform.

In the main assessment, each participant evaluated five randomly selected social state-of-affairs drawn from a corpus of 200 statements, which were themselves selected from the 233 validated social state-of-affairs evaluated in *study 6A*. This sampling ensured broad topical coverage while maintaining the empirical foundation established in the earlier study. For each statement, participants first rated whether they are true and their confidence in that assessment and then answered a series of questions concerning mind-dependence, essentialism, and unthinkability.

After completing the main assessment, participants were randomly assigned to one of several individual-differences modules, each containing a different set of measures. These included Raven's Standard Progressive Matrices (nine-item version), Schwartz Values and System Justification scales combined with the Inherency scale, the Big Five Inventory–2, or a set of items measuring Actively Open-Minded Thinking and Folk Moral Objectivism. Each module also included a brief demographic questionnaire.

All procedures and analytic strategies were pre-registered (see: https://aspredicted.org/9zp4nd.pdf). The primary objective was to evaluate how the six conceptual dimensions of social reality: scope, institutionalization, hierarchical function, sociopolitical sensitivity, moral valence, and intentionality, relate to perceived mind-dependence across a wide range of social state-of-affairs. Multilevel models were used to account for both participant-level and state-of-affairs-level variation. Secondary pre-registered analyses examined whether individual differences influence the tendency to attribute mind-dependence to social facts. All analyses were conducted in R, using the *lmerTest*, and *parameters* packages.

Results

Main analysis. The relation between each of the six social-reality dimensions and perceived mind-dependence was evaluated using a hierarchical linear model (HLM) of the form:

State-of-Affairs Level Mind-Dependence Score ~ Social-Reality-Dimension + (1 +Social-Reality-Dimension | Participant)

Each model estimated the standardized fixed effect of a given dimension on mind-dependence while accounting for random variation across participants. Standardized coefficients for all six models are presented in Table S7. Across dimensions, two of the social-reality characteristics: institutionalization and intentionality, showed significant negative associations with mind-dependence. Specifically, state-of-affairs perceived as more institutionalized were judged as less mind-dependent (β = –0.21, *SE* = 0.07, 95% CI [–0.34, –0.07], $p < .005$), and facts characterized by higher intentionality were also associated with lower mind-dependence ratings (β = –0.17, *SE* = 0.07, 95% CI [–0.32, –0.03], $p < .05$). In contrast, moral valence showed a significant positive relation with mind-dependence (β = 0.26, *SE* = 0.08, 95% CI [0.11, 0.41], $p < .005$), indicating that morally charged social facts were perceived as more dependent on wishes and beliefs. The remaining dimensions, i.e. hierarchical function, sociopolitical sensitivity, and scope, did not significantly predict mind-dependence (all *ps* > .25).

Table S7. Standardized coefficients for the relation between social-reality dimensions and mind-dependence.

Parameter	*β*	SE	CI	Low	High	t	p	sig.
Hierarchical	0.01	0.071	0.95	-0.13	0.15	0.11	0.91	
Institutional	-0.21	0.07	0.95	-0.34	-0.07	-2.97	< 0.005	**
Intentionality	-0.17	0.07	0.95	-0.32	-0.03	-2.37	< 0.05	*
Moral	0.26	0.08	0.95	0.11	0.41	3.38	< 0.005	**
Sociopolitical	0.09	0.08	0.95	-0.07	0.24	1.17	0.25	
Scope	-0.08	0.07	0.95	-0.22	0.06	-1.14	0.26	

Individual differences analysis. The relation between each measurement of individual differences and perceived mind-dependence was evaluated using a hierarchical linear model (HLM) of the form:

Participant Level Mind-Dependence Score ~ Measurement-Score + (1 + Measurement-Score-Centered || Stata-of-Affairs)

Across the big five personality domains, conscientiousness showed a significant negative association with perceived mind-dependence (B = –0.44, 95% CI [–0.66, –0.22], p < .001), indicating that more conscientious individuals tended to attribute less mind-dependence to social state-of-affairs. Conversely, negative emotionality was positively related to mind-dependence (B = 0.23, 95% CI [0.05, 0.41], p = .015), suggesting that individuals higher in negative emotionality were more likely to view social facts as mind-dependent. No significant effects were observed for extraversion, agreeableness, or open-mindedness (all ps > .08).

Facet-level analyses revealed a similar pattern; within conscientiousness, all three facets were negatively associated with perceived mind-dependence: higher organization predicted lower mind-dependence (B = −0.34, 95% CI [−0.52, −0.15], p < .001), as did productiveness (b = −0.32, 95% CI [−0.51, −0.13], p = .001) and responsibility (B = −0.37, 95% CI [−0.58, −0.16], p = .001). For negative emotionality, all three facets were positively related to mind-dependence; anxiety (B = 0.19, 95% CI [0.02, 0.36], p = .025), depression (B = 0.17, 95% CI [0.01, 0.33], p = .042), and emotional volatility (B = 0.20, 95% CI [0.03, 0.38], p = .020). Within agreeableness, respectfulness exhibited a reliable negative association (B = −0.41, 95% CI [−0.70, −0.12], p = .005), whereas compassion and trust were not significant (both p > .30). Extraversion facets, sociability, assertiveness, and energy level, did not predict mind-dependence (all p > .17). Finally, within openness to experience, intellectual curiosity showed a small, non-significant positive trend (B = 0.21, p = .083), while aesthetic sensitivity and creative imagination were not associated with mind-dependence (both p > .36).

Among the cognitive and attitudinal measures, actively open-minded thinking (AOT) scores were unrelated to mind-dependence across all subscales (all ps > .25). In contrast, measures of folk moral objectivism (FMO) revealed reliable positive associations; both relativism (B = 0.23, 95% CI [0.06, 0.40], p = .007) and no-truth orientation (B = 0.25, 95% CI [0.09, 0.41], p = .002) predicted higher mind-dependence. Examining values, self-direction from the schwartz value survey showed a small but significant positive association (B = 0.12, 95% CI [0.00, 0.23], p = .048). Other value dimensions, including power, achievement, hedonism, stimulation, universalism, benevolence, tradition, conformity and security, were not significantly related. Inherency did not predict mind-dependence (B = −0.02, 95% CI [−0.15, 0.11], p = .732), indicating that a general tendency to view social arrangements as inherent features of the world was unrelated to whether those arrangements were seen as mind-dependent. Likewise, system-justifying attitudes showed no reliable association (B = −0.03,

95% CI [−0.12, 0.06], p = .501), suggesting that motivation to defend existing social systems does not extend to construing them as more or less dependent on collective mental states. Likewise, cognitive ability, as measured by performance on the nine-item version of raven's standard progressive matrices, was not significantly associated with perceived mind-dependence. The fixed-effects model indicated a small, nonsignificant coefficient for raven's score (B = 0.006, 95% CI [–0.01, 0.02], p = .389), suggesting that general cognitive ability did not meaningfully predict how participants perceive the extent to which social facts are mind-dependent.

Finally, demographic variables showed small associations with perceived mind-dependence. Age was a negative predictor (B = –0.02, 95% CI [–0.02, –0.01], p < .001), indicating that older participants tended to attribute less mind-dependence to social state-of-affairs. In contrast, education showed a positive association (B = 0.10, 95% CI [0.04, 0.17], p = .002), suggesting that more highly educated individuals were more inclined to view social state-of-affairs as mind-dependent. Income was unrelated to mind-dependence (B = –0.02, 95% CI [–0.07, 0.03], p = .45).

Additional analysis. Relation between essentialism and mind dependence within and between topics and participants in the large sample on 1200 participants and 200 state of affairs was assessed and presented in table S8.

Table S8. Results of multilevel relation between essentialism and mind-dependence.

Parameter	β	SE	95% CI	t	df	p
Fixed Effects						
(Intercept)	5.37	0.33	(4.72, 6.02)	16.20	368.81	< .001
Eessentialising Between Participants	-0.26	0.04	(-0.34, -0.18)	-6.73	944.10	< .001
Eessentialising Between State-of-affairs Mean	-0.23	0.06	(-0.36, -0.11)	-3.78	203.48	< .001
Eessentialising Within Participants and State-of-affairs	-0.17	0.02	(-0.21, -0.13)	-8.25	203.64	< .001

Study 8

Methods

Participants

A total of 163 English-speaking participants (67 men, 96 women) were recruited via the Prolific platform. Eligibility criteria included English as a first language, a minimum of 50 prior submissions, and an approval rate above 95%. Ten participants who failed one or more attention checks and an additional thirteen participants who failed the clarification questions were excluded, resulting in a final analytic sample of 140 participants. All participants provided informed consent before beginning the study.

Procedure

Participants evaluated six statements concerning the value of Bitcoin, Dollar bills, Shares of stock, Works of art, Diamonds, and Gold bullion. For each statement, participants first indicated whether they believed the statement "has value" to be true and rated their confidence in this judgment. Next, participants were asked to rank the six statements according to their perceived mind-dependence of value, using the five core items of the mind-dependence scale (e.g., "Rank the statements from the one whose value is most dependent on what people believe to the one whose value is least dependent on what people believe"). Finally, participants rated the same six statements, presented in random order, on how likely each was to increase, decrease, or preserve its value in the future.

Results

Table S9. Means and standard deviations of MD Scores per topic

Topic	Mean Score	SD
Bitcoin	3.83	1.33
Shares of Stock	3.59	0.96
Dollar Bills	3.54	1.35
Works of Art	3.53	1.43
Diamonds	3.44	1.14
Gold Bullion	3.08	1.22

Figure S5. Violin plot of Mind-dependence scores of the fact an asset holds value

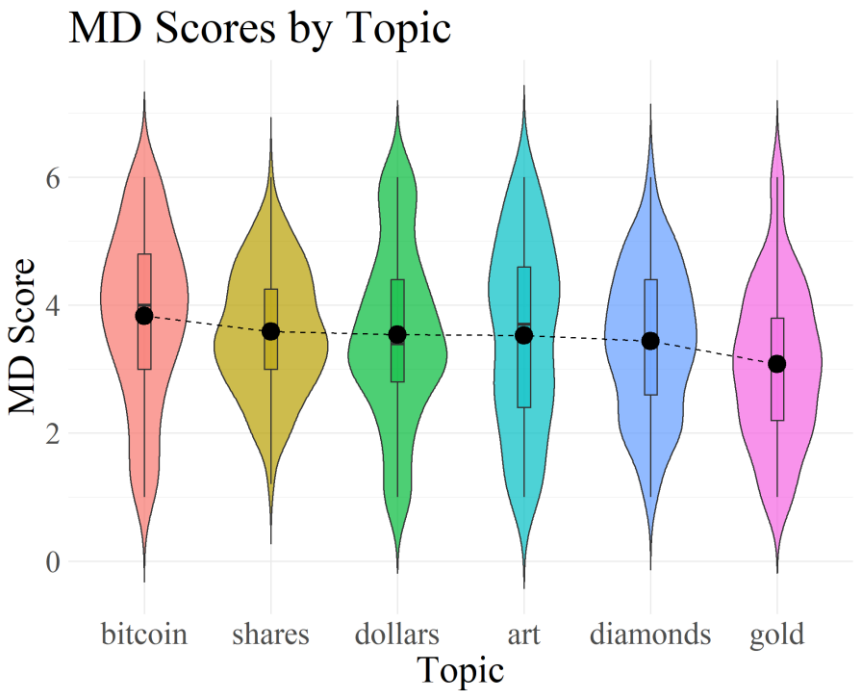

Note. MD Score = Mind dependence score. Point represent the mean of each asset.

To examine topic-level differences, a linear model was estimated with Topic as a fixed factor (Bitcoin as the reference). Results revealed significant differences in perceived mind-dependence across topics, $F(5, 834) = 25.4$, $p < .001$. Results are presented in table S10.

Table S10. linear regression of MD predicted by Topic

Parameter	**Coefficient**	**SE**	**95% CI**	**t(834)**	***p***
Intercept (Bitcoin)	3.83	0.11	[3.63, 4.04]	36.36	< .001
Shares of Stock	−0.25	0.15	[−0.54, 0.05]	−1.65	.100
Dollar Bills	−0.29	0.15	[−0.59, 0.00]	−1.97	.049
Works of Art	−0.31	0.15	[−0.60, −0.01]	−2.05	.041
Diamonds	−0.40	0.15	[−0.69, −0.10]	−2.65	.008
Gold Bullion	−0.76	0.15	[−1.05, −0.46]	−5.07	< .001

Subsequently, we examined whether perceived mind-dependence predicted expectations about future changes in value. Three separate linear models regressed Increase, Decrease, and Preservation ratings on mind-dependence scores (see table 3, and figures 2-4).

Table S11. Linear models Predicting Increase, Decrease, and Preservation from MD

N	Outcome	Predictor	B	SE	95% CI	t(838)	p
1	Increase	mind-dependence	0.12	0.05	[0.02, 0.21]	2.48	.013

	2	Decrease	mind-dependence	0.35	0.05	[0.26, 0.44]	7.72	< .001
	3	Preservation	mind-dependence	−0.19	0.05	[−0.28, −0.10]	−4.04	< .001

Study 9

Methods

Participants

The target sample for this study was 620 participants. A total of 685 English-speaking adults were recruited through the Prolific platform. Eligibility criteria required participants to have completed at least 50 previous submissions with an approval rate above 95%. Participants who failed one or more attention or comprehension checks were replaced, in accordance with the pre-registered exclusion protocol. In total, 65 participants failed at least one of these checks, leaving a final eligible sample of 620 participants who successfully completed the study. All participants provided informed consent prior to participation.

Procedure

Participants read a vignette describing a gentrification scenario in which a real-estate developer, Dave Green, purchases a property that ultimately leads to the displacement of a long-term tenant, Sam Coleman. To ensure comprehension of the narrative, participants completed two memory items prior to the experimental manipulation. Each item displayed a photograph and asked participants to identify whether the person depicted was the developer or the tenant.

The experimental manipulation followed the memory check. All participants were presented with two arguments concerning whether gentrification should be understood as mind-dependent or mind-independent. Participants assigned to the mind-dependent (MD) condition viewed one strong and one weak argument supporting the mind-dependent interpretation. Participants assigned to the mind-independent (MID) condition viewed one strong and one weak MID argument. While reading both arguments, participants were asked to label each sentence as a good or bad argument. They then indicated which of the two arguments they found more persuasive and provided an open-ended explanation of at least 20 words describing the reasoning that influenced their choice.

Participants subsequently completed the key dependent measures. Responsibility attribution was assessed using a 0–100 slider asking how responsible Dave Green was for the residents losing their homes. A second slider assessed the extent to which participants liked the developer, ranging from 0 (“not at all”) to 100 (“very much”). Political orientation was measured on a 1–7 scale and was included for exploratory analyses. All procedures and analytic strategies were pre-registered and can be accessed via the following link: https://aspredicted.org/8r97tn.pdf. Primary analysis consisted of a one-tailed independent-samples t-test predicting higher responsibility ratings in the MD condition than in the MID condition. A secondary preregistered t-test examined whether liking of the developer would be lower in the MD condition.

Results

Main Analysis. The primary hypothesis concerned whether participants in the MD condition would attribute greater responsibility to the developer than those in the MID condition. Consistent with this prediction, responsibility ratings were higher in the MD condition ($M = 79.27$) than in the MID condition ($M = 75.49$). A one-tailed t-test confirmed that this difference was statistically significant, $t(603.37) = 2.00$, $p = .023$, 95% CI [0.67, ∞]. The effect size indicated a small difference in the predicted direction, Cohen's d = −0.16, 95% CI [−0.32, −0.00].

Secondary analysis. The preregistered secondary hypothesis predicted that liking of the developer would be lower in the MD condition. This prediction was also supported. Participants in the MD condition reported significantly lower liking ($M = 22.39$) than those in the MID condition ($M = 27.89$). A one-tailed t-test indicated that this difference was significant, $t(603.53) = −2.81$, $p = .003$, 95% CI [−∞, −2.27]. The effect size was small, Cohen's d = 0.23, 95% CI [0.07, 0.38], suggesting that presenting the scenario as mind-dependent not only increased perceived responsibility but also reduced affective evaluations of the developer.

Mediation Analysis. To examine whether responsibility attribution mediated the effect of condition (MID → MD) on liking of the developer, we estimated a mediation model using structural equation modelling with standardized coefficients. Confidence intervals were derived from the bootstrap distribution (1,000 resamples). Results indicated that condition significantly predicted liking of the developer, such that moving from the MID condition to the MD condition was associated with lower liking ($\beta = −.08$, SE = .04, 95% CI [−.16, −.01], $p = .028$,). Condition also significantly predicted responsibility attribution, with participants in the MD condition attributing greater responsibility to the developer ($\beta = .08$, SE = .04, 95% CI [.00, .16], $p = .040$). In turn, higher responsibility attribution was strongly associated with lower liking of the developer ($\beta = −.35$, SE = .05, 95% CI [−.44, −.27], $p < .001$). The indirect effect of condition on liking via responsibility attribution was statistically significant ($\beta = −.03$, SE = .01, $p = .049$, 95% CI [−.06, .00]), indicating potential partial mediation. The total effect of condition on liking was also significant ($\beta = −.11$, SE = .04, 95% CI [−.19, −.03], $p < .005$). Together, these results suggest that framing the scenario as more mind-dependent reduces liking of the developer in part by increasing perceived responsibility, while a substantial portion of the effect remains direct.

References

1. Everett, J. A. C. The 12 Item Social and Economic Conservatism Scale (SECS). *PLoS One* **8**, e82131 (2013).

2. DeYoung, C. G., Quilty, L. C. & Peterson, J. B. Between facets and domains: 10 aspects of the Big Five. *J. Pers. Soc. Psychol.* **93**, 880–896 (2007).