eLife

Reviewed Preprint
v1 • March 4, 2025
Not revised

Neuroscience

The resource elasticity of control

Levi Solomyak, Aviv Emanuel, Eran Eldar

Department of Cognitive Sciences, Hebrew University of Jerusalem, Jerusalem, Israel • Department of Psychology Hebrew University of Jerusalem, Jerusalem, Israel • Department of Psychology, McGill University, Montreal, Canada

eLife Assessment

This study makes the **important** claims that people track, specifically, the elasticity of control (rather than the more general parameter of controllability) and that control elasticity is specifically impaired in certain types of psychopathology. These claims will have implications for the fields of computational psychiatry and computational cognitive neuroscience. However the evidence for the claim that people infer control elasticity is **incomplete**, given that it is not clear that the task allows the elasticity construct to be distinguished from more general learning processes, the chosen models aren't well justified, and it is unclear that the findings generalize to tasks that aren't biased to find overestimates of elasticity. Moreover, the claim about psychopathology relies on an invalid interpretation of CCA; a more straightforward analysis of the correlation between the model parameters and the psychopathology measures would provide stronger evidence.

https://doi.org/10.7554/eLife.105507.1.sa4

Abstract

The ability to determine how much the environment can be controlled through our actions has long been viewed as fundamental to adaptive behavior. While traditional accounts treat controllability as a fixed property of the environment, we argue that real-world controllability often depends on the effort, time and money we are able and willing to invest. In such cases, controllability can be said to be *elastic* to invested resources. Here we propose that inferring this elasticity is essential for efficient resource allocation, and thus, elasticity misestimations result in maladaptive behavior. To test these hypotheses, we developed a novel treasure hunt game where participants encountered environments with varying degrees of controllability and elasticity. Across two pre-registered studies (N=514), we first demonstrate that people infer elasticity and adapt their resource allocation accordingly. We then present a computational model that explains how people make this inference, and identify individual elasticity biases that lead to suboptimal resource allocation. Finally, we show that overestimation of elasticity is associated with elevated psychopathology involving an impaired sense of control. These findings establish the elasticity of control as a distinct cognitive construct guiding adaptive behavior, and a computational marker for control-related maladaptive behavior.

Introduction

Taking action requires us to invest resources, such as time, money, and effort, and is thus only sensible if our actions have the potential to change the environment to our advantage. Accurate estimation of this potential – that is, of the environment's controllability[1-4] – enables us to make informed decisions as to whether to engage in action[5-7]. Misestimations of controllability thus result in maladaptive behavior, and contribute to mental health disorders such as depression, anxiety, and obsessive-compulsive disorder[8-12].

The degree of control we possess over our environment, however, may itself depend on the resources we are willing and able to invest. For example, the control a biker has over their commute time depends on the effort they are willing and able to invest in pedaling. Likewise, the control a diner in a restaurant has over their meal may depend on how much money they have to spend. In such situations, controllability is not fixed but rather *elastic* to invested resources (i.e., in the same sense that supply and demand may be elastic to changing prices[13]). Importantly, high controllability need not imply high or low elasticity. Controllability can be high, yet inelastic to invested resources – for example, when choosing among equally priced bus routes to work (**Figure 1**).

Knowing the elasticity of control is necessary for deciding whether and how to act. Failing to achieve one's goal, for example, should only lead one to invest more resources in an elastic environment. Critically, in many real-life domains (e.g., studying for an exam), we must use our experiences to infer the degree to which we can improve our outcomes by investing more resources[16-20]. Failure to do so is bound to result in maladaptive behavior, and may thus contribute to mental health disorders. Depressed patients, for instance, are less willing to expend their resources to obtain control[21,26-27], which could result from an underestimation of elasticity. Conversely, the repetitive actions that characterize obsessive compulsive disorder patients[24] could indicate that such patients overestimate elasticity, believing that the additional efforts they invest will yield progressively better outcomes. Ultimately, not having a good grasp of how one's additional efforts influence outcomes is bound to undermine one's sense of agency[14-15].

Thus, here we ask whether and how people infer the elasticity of control over their environment. We hypothesize that individual differences in elasticity inference lead to a mismanagement of one's resources when attempting to control the environment, and are thus associated with psychopathologies involving a distorted sense of control[21-23].

Experimental paradigms to date have tested elasticity and controllability as coupled, such that controllability was either low or elastic[16-20]. Elasticity, however, must be dissociated from controllability to accurately diagnose mismanagement of resources. A given individual, for instance, may tend to overinvest resources because they overestimate controllability – for example, exercising due to a misguided belief that this can make one taller. Alternatively, they may do so because they overestimate the elasticity of control – for example, a chess expert practicing unnecessarily hard to win against a novice. In this latter example, the amount of effort the expert chooses to invest would not change their level of control over the match's outcome. Diagnosing the former source of mismanagement of resources requires a manipulation of controllability, whereas diagnosing the latter requires manipulating the elasticity of control.

Therefore, here we dissociate elastic and inelastic controllability using a novel treasure-hunt game where participants encountered three distinct environments: The first offered *high elastic controllability*, where a high level of control can only be attained by investing extra resources. The second offered *high inelastic controllability*, where control is attainable with any level of resource investment. And the third offered *low controllability*, where control was not attainable despite any

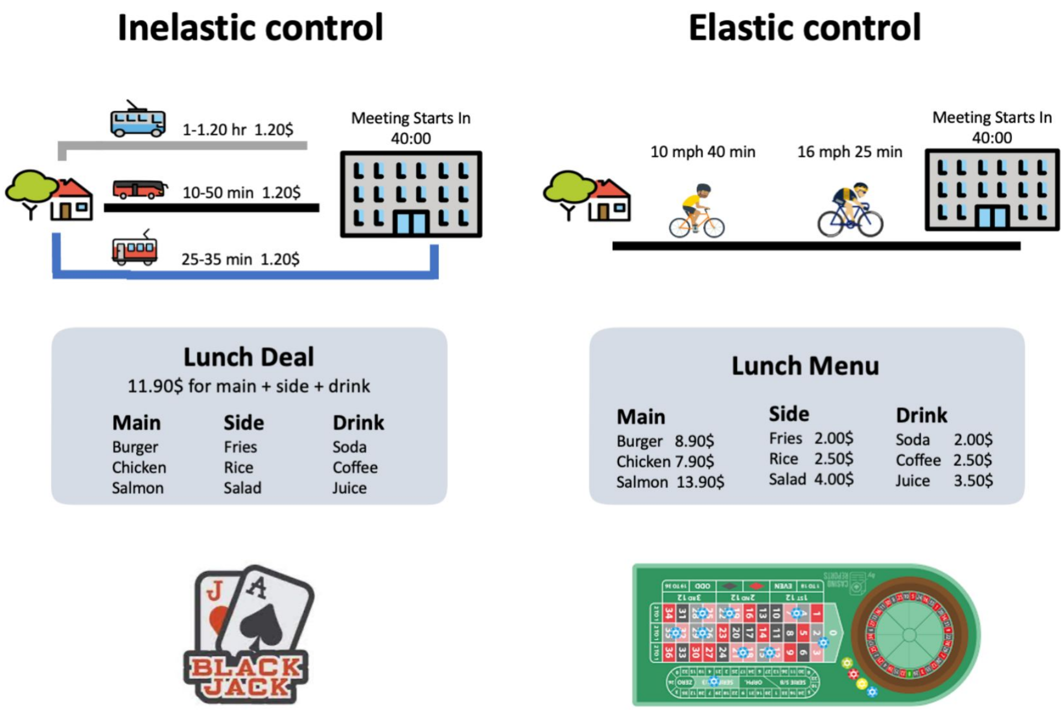

Figure 1.

Examples of inelastic vs elastic control:

Top – Choosing one of three equally priced public transport routes provides inelastic control over commute time, whereas the control of biker over commute time is elastic to the effort they invest in pedaling. **Middle –** A diner has either inelastic or elastic control over their lunch depending on whether the restaurant offers a fixed-price lunch deal, or a standard menu where dishes vary in price. **Bottom –** the probability of winning in blackjack is inelastic to money, as it only depends on whether one hits or stands, whereas when playing roulette, one can increase the probability of winning by investing more coins to cover more possible outcomes. Importantly, most real-world scenarios lie on a spectrum between these illustrative extremes, such that controllability is partly elastic and partly inelastic.

level of resource investment. In each environment, participants were allowed to invest various amounts of resources to pursue reward or to altogether forgo pursuing reward. Examining how participants adapted their resource investment to observed outcomes enabled us to determine whether and how participants inferred each environment's controllability and its elasticity.

Across two pre-registered studies (N = 514), we establish that people do indeed infer the elasticity of control. We present a computational model that explains how people make this inference, and identify individual biases in its implementation. We thus show that elasticity overestimation leads to an overinvestment of resources and is associated with elevated psychopathology involving an impaired sense of control. These findings establish the elasticity of control as a distinct inference guiding adaptive behavior, and a computational marker of control-related psychopathology.

Results

We had online participants (first study = 264, replication study = 250, ages =18-45, Mean 33 ± .6) play a novel treasure-hunt game set across multiple planets. On each trip (trial), participants attempted to travel from an initial location ('desert' or 'fountain') to a treasure (worth 150 coins), located at either the house or the mountain (**Figure 2A**). Participants could exercise control by boarding the train to reach the house, or the plane to reach the mountain (right-side **Figure 2B**). The present planet's controllability determined the probability that participants would succeed in boarding the vehicle they selected. In high-controllability planets, participants consistently boarded their preferred vehicle, enabling them to reliably reach their destination. Conversely, in low-controllability planets, participants frequently failed to board and were thus forced to walk to the nearest location, where the treasure was located in only 20% of trials.

To investigate elasticity learning, we allowed participants to invest different amounts of resources in attempting to board their preferred vehicle. Specifically, participants could purchase one (40 coins), two (60 coins), or three tickets (80 coins) or otherwise walk for free to the nearest location. We defined inelastic controllability as the probability that even one ticket would lead to successfully boarding the vehicle, and elastic controllability as the degree to which two extra tickets would increase that probability. Participants were informed that a single ticket allowed them to board the vehicle only if it stopped for them at the station, but that they may be able to increase their chances of boarding by purchasing extra tickets. Each additional ticket awarded participants an attempt to jump onto the moving vehicle after it had already left the main platform (**Figure 2C**). At the end of each trip, participants were shown where they arrived at, enabling them to infer whether they successfully boarded their vehicle.

To dissociate elasticity from inelastic controllability, each participant visited three randomly sampled planets. In the first planet, controllability was high and inelastic, meaning that no matter how many tickets you purchased, you had a good chance of making your ride (sampled from green area in **Figure 2D**). In the second planet, controllability was high and elastic, meaning that you needed to purchase two extra tickets to have a good chance of making your ride (**Figure 2D**, blue). And in the third planet, controllability was low, meaning that no matter how many tickets you purchased, you did not have a good chance of making your ride (**Figure 2D**, red).

Importantly, the order of these planets was randomized, and participants did not know in advance which planet they were visiting. Thus, they could only infer a planet's controllability and elasticity from the outcomes of their actions. Moreover, to increase sensitivity to individual biases, we had all participants visit a fourth planet where purchasing any number of tickets (0, 1, 2, 3) was equally optimal (**Figure 2D**, black circle). We also homogenized participants' initial learning experiences by giving them three free tickets for the first five visits to every planet. As we shall see below, this latter measure also helped dissociate between different models of how participants solved the task.

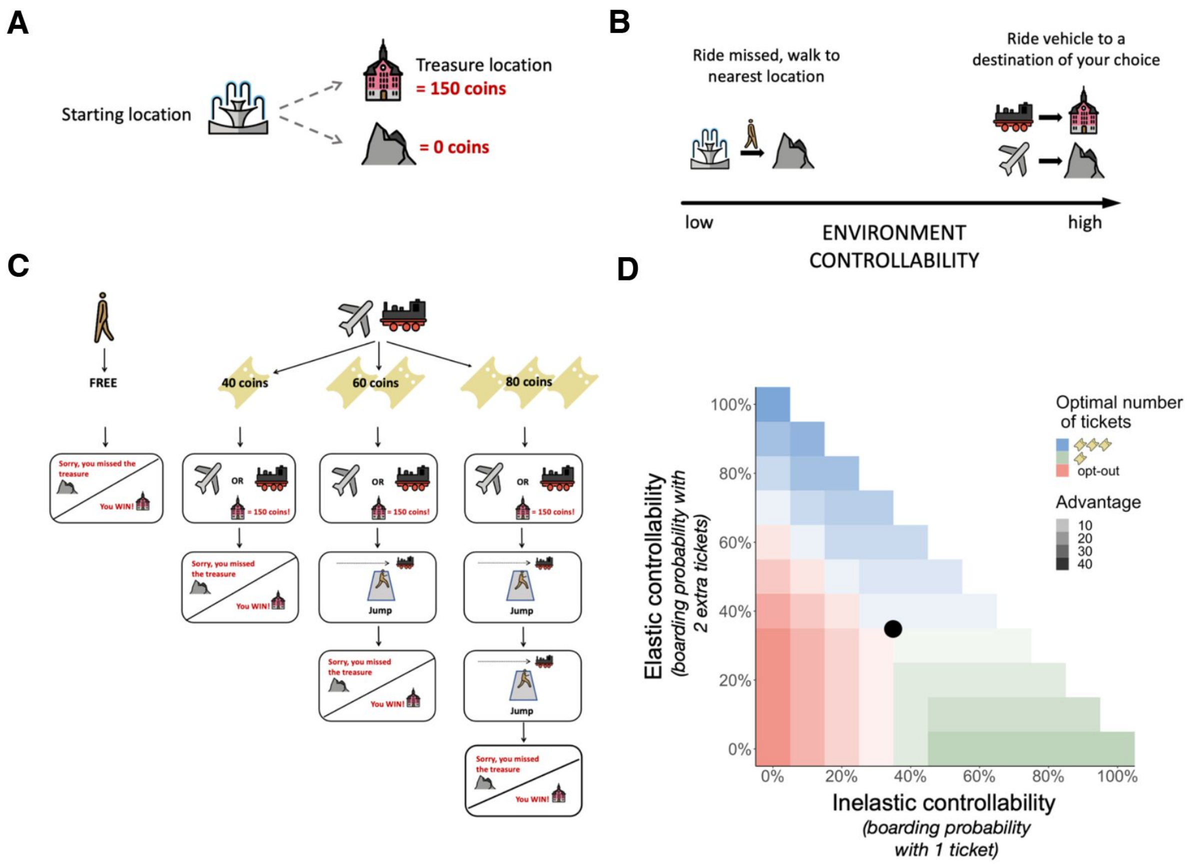

Figure 2.

Experimental design:

(A) **Goal**. On each trip to a planet, participants' goal was to reach a treasure located either at the house or the mountain (B) **Transition rules**. Participants could exercise control by boarding either the plane or the train to a destination of their choice, whereas missing the ride sent the participant walking to the nearest location (which happened to be the treasure location in only 20% of trials). (C) **Trip structure**. At the beginning of each trip, participants selected whether to purchase 1, 2, or 3 tickets to attempt to board their vehicle of choice, or walk for free to the nearest location. If the participant purchased at least one ticket, they were allowed to choose between the plane and the train. Then, for each additional ticket purchased, the participant was given an opportunity to increase their chances of boarding the vehicle. Specifically, the chosen vehicle appeared moving from left to right across the screen, and the participant attempted to board it by pressing the spacebar when it reached the center of the screen. At the end of the trip, participants were shown where they arrived at, allowing them to infer whether they successfully boarded the vehicle. (D) **The space of possible planets**. Planets varied in inelastic controllability – the probability that even one ticket would lead to successfully boarding the vehicle – and in elastic controllability – the degree to which two extra tickets would improve the probability of successfully boarding the vehicle (one extra ticket provided half the benefit). The color corresponds to the optimal number of tickets (0, 1 or 3) in each planet. The darker the color the higher the advantage in expected value gained by purchasing the optimal number of tickets relative to the second-best option. Each participant made 30 consecutive trips in one planet from the green area, one planet from the blue area, one planet from the red area, and one planet with identical characteristic across all participants (black circle).

People adapt to the elasticity of control

To determine whether participants inferred elasticity, we first examined how they adapted their resource investment across planets. We found that participants were more likely to purchase tickets ('opt in') in more controllable planets, whether controllability was elastic (Initial and Replication studies, p<.001; mixed logistic regression) or inelastic (Initial and Replication studies, p <.001; **Figure 3A&C**). Conversely, participants purchased more extra tickets only in planets with high elastic controllability (Initial p =.01, Replication p=.004; mixed probit regression; **Figure 3B&C**). Although substantial individual differences were observed in this regard (**Figure 3D**), these results show that most participants successfully adapted to the controllability and elasticity of the environments they visited.

People infer the elasticity of control

That participants adapted to elasticity does not necessarily mean they inferred elasticity. Successful adaptation may also be achieved merely by treating each possible level of resource investment as a separate course of action. One could thus learn the degree of control afforded by each course of action (i.e., purchasing 1, 2, or 3 tickets), and then choose the one that affords the highest expected value considering its cost. This simple strategy could achieve optimal performance given unlimited experience. However, it would not be as efficient as a strategy that employs elasticity inference, because it does not make use of the dependencies that exist between different levels of resource investment in the degrees of control they afford. Thus, upon failing to board a vehicle despite purchasing 3 tickets, the simpler strategy would need to purchase 1 or 2 tickets to know these options would fail as well, whereas an agent equipped with the concept of elasticity would conclude that controllability is low, and it is thus best to opt out. Conversely, upon succeeding to board with 3 tickets, the simpler strategy would tend to favor purchasing 3 tickets, over 1 or 2 tickets, even if the extra tickets confer no benefit, whereas an agent that infers elasticity would not favor 3 over 1 or 2 tickets because it received no evidence that controllability is elastic (i.e., that the extra tickets help).

To determine which of these models best describes the way participants solved the task, we formalized the simpler strategy as the 'controllability model', and an elasticity-inferring strategy as the 'elastic controllability model' (pre-registration: *https://aspredicted.org/CHW_12H*). Both models use beta distributions, characterized by parameters a and b, to represent the probability of boarding when purchasing different numbers of tickets, but only the elastic controllability model learns latent elasticity estimates. Specifically, in the 'controllability model', three beta distributions represent the expected level of control from purchasing 1, 2, or 3 tickets. Thus, the parameters of these distributions (a_n, b_n, $n \in \{1, 2, 3\}$) accumulate the number of times each number of tickets led (a_n) or did not lead (b_n) to successful boarding. Conversely, in the 'elastic controllability model', the beta distributions represent a belief about the maximum achievable level of control (a_control,b_control) coupled with two elasticity estimates, specify the degree to which successful boarding requires purchasing at least one ($a_{\mathrm{elastic}\geq1}$, $b_{\mathrm{elastic}\geq1}$) or specifically two ($a_\mathrm{elastic2}$, b_elastic2) extra ticket. **Figure 4** shows schematically how the two models differ in updating their estimates in response to different types of outcomes. Both models similarly use these estimates to compute an expected value for purchasing 0, 1, 2, or 3 tickets, and accordingly choose between these options (see Methods, pg. 20-23, for model equations).

Examining the predictions of both models, each simulated using the parameter setting that best fit participants' choices, showed that the elastic-controllability model was better than the controllability model at opting-out in low controllability planets (elastic-controllability model % opt in = 49% ± .01, controllability model % opt in = 63% ± .01; **Figure 5C**, top panel). Most importantly, this higher level of performance closely matched participants' behavior (% opt in = 47%, ±.02). Furthermore, the elastic controllability model was also better at not purchasing extra

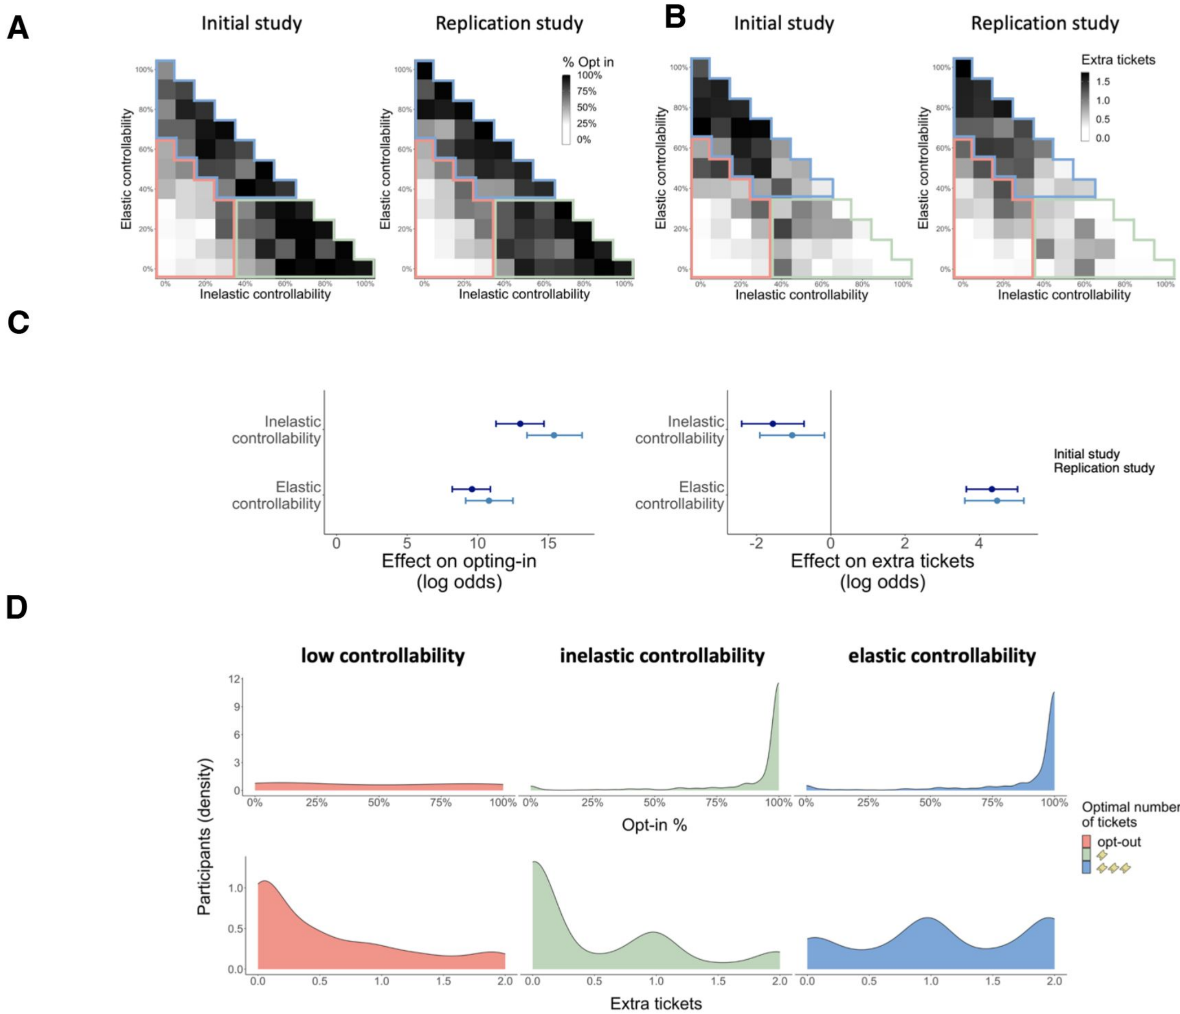

Figure 3.

Participants adapted to the elasticity of control:

Results from initial (N = 264) and replication (N = 250) studies. (A) **Opt-in percentage across all planets**. Participants opted-in more frequently on controllable planets (outlined in blue and green) than on uncontrollable planets (outlined in red) (B) **Extra tickets purchased across all planets**. Participants purchased more extra tickets in planets with high elastic controllability (outlined in blue). An average of 11 to 12 participants visited each planet in each of the studies. (C) **Effect of elastic and inelastic controllability on opting in and the purchasing of extra-tickets**. Bars show estimated fixed effects and 95% CIs from mixed logistic (opt-in) and probit (extra tickets) regressions. (D) **Individual differences**. Distribution of participants by opt-in rate (top) and average number of extra tickets purchased (bottom) across both studies, shown separately for planets with low controllability (red), high inelastic controllability (green), and high elastic controllability (blue).

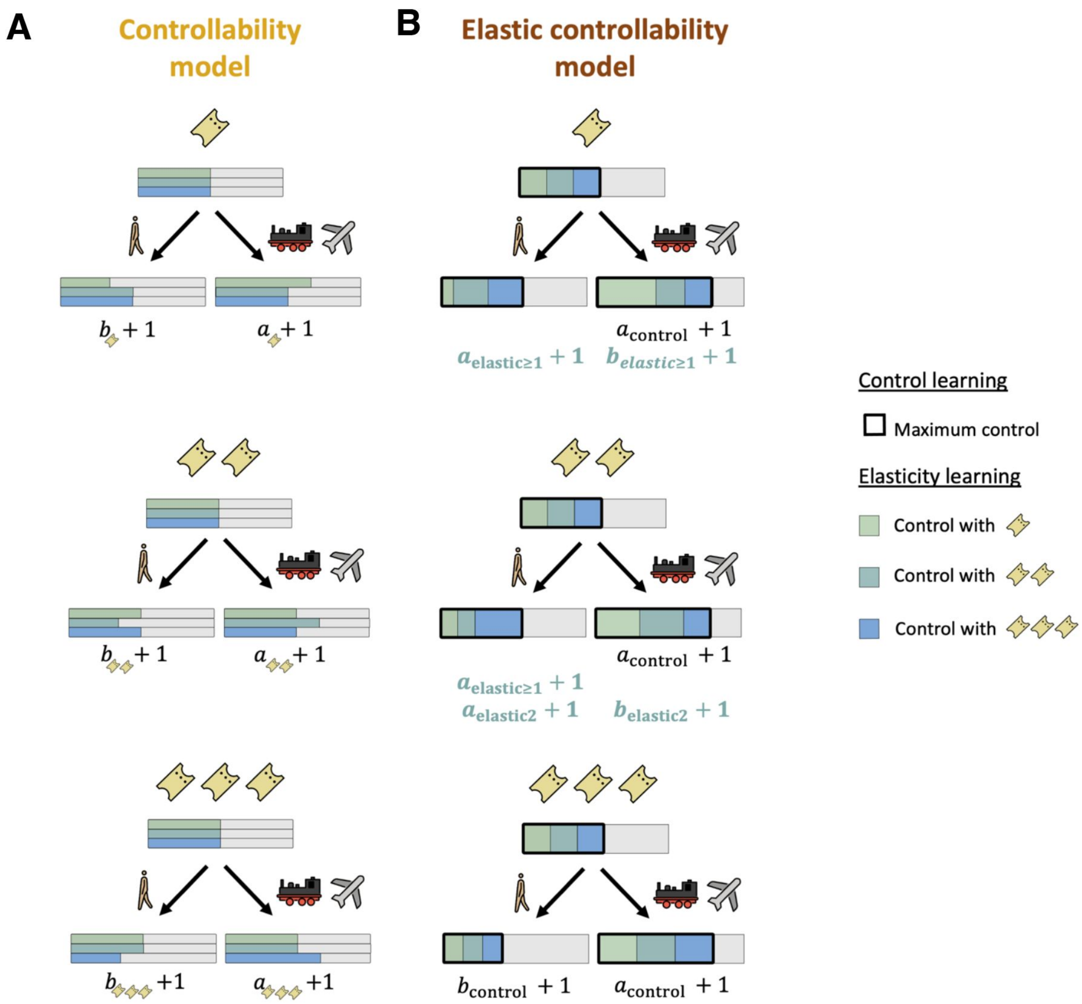

Figure 4.

Computational models:

Illustration of the learning rules of the *controllability* and *elastic controllability* models. The width of the colored region represents the estimated control, shown as a percentage of absolute control (gray area) along with the update rules based on each outcome. (A) **Controllability model:** This model treats the purchase of 1, 2, and 3 tickets as distinct actions. It accumulates evidence for the effectiveness (a) or lack of effectiveness (b) of each action based solely on whether or not the action led to successful boarding, illustrated by changes in the width of the shaded region corresponding to each ticket amount. (B) **Elastic Controllability model:** Represents beliefs about maximum control (black outline) and the degree to which at least one or specifically two extra tickets are necessary to obtain control. These are combined in calculating the expected control with 1, 2, and 3 tickets, demonstrated by the bars. The arrangement of the bars illustrates that, only in the elastic controllability model, controllability offered by extra tickets is added up on top of what is offered by a single ticket. The model updates its beliefs as follows: **Top –** Successfully boarding with one ticket provides evidence of high maximum controllability (a_{control} + 1; expanded shaded region) and reduces the perceived need to purchase extra tickets ($b_{\text{elastic}\geq 1}$ + 1), increasing expectations of purchasing 1, relative to 2 or 3, tickets (light green expanded). A failure to board does not change estimated maximum controllability, but rather suggests that 1 ticket might not suffice to obtain control ($a_{\text{elastic}\geq 1}$ + 1; light green diminished). **Middle –** Successfully boarding with 2 tickets provides evidence of high maximum controllability (a_{control} + 1; expanded shaded region), and reduces the perceived need to purchase two extra tickets (b_{elastic2} + 1), increasing expectations of purchasing 1 or 2, relative to 3, tickets (light & dark green expanded). Here too, a failure to board does not change estimated maximum controllability, but rather suggests that 2 tickets might not be sufficient (a_{elastic2} + 1; light & dark green diminished). **Bottom** – Successfully boarding with 3 tickets provides evidence of high maximum controllability (a_{control} + 1; expanded shaded region), but offers no insight about whether fewer tickets would suffice to obtain it. Failure to board with 3 tickets provides evidence of low controllability (b_{control} + 1; diminished shaded region).

tickets in low-controllability (elastic controllability extra tickets = .49±.01 controllability extra tickets = .65±.01) and inelastic-controllability (elastic controllability extra tickets = .58±.01 controllability extra tickets = .66±.01) planets, and this too was more closely aligned with participant behavior (low-controllability extra tickets = 51±.02, inelastic-controllability extra tickets = 53±.02; **Figure 5C**, bottom panel).

Beyond these general aspects of task behavior, having participants start on every planet with free 3-ticket trips also enabled us to examine their response to succeeding or failing with 3 tickets, prior to any other experiences on the planet (**Figure 5D**). As noted above, failure with 3 tickets makes the controllability model favor 1 and 2 tickets over 3 tickets, whereas success with 3 tickets makes it favor 3 tickets over 1 and 2 tickets. Critically, neither the participants nor the simulation of the elastic controllability model showed these effects. Upon failing with 3 tickets, both participants and the elastic model primarily favored to opt out. Conversely, if they succeeded with 3 tickets, they did not become less likely to purchase 1 or 2 tickets.

In fact, failure with 3 tickets even made participants favor 3, over 1 and 2, tickets. Presumably, failure resulted in increased uncertainty about whether it is at all possible to control one's destination, and thus, participants who nevertheless opted in invested as much resources as they could to resolve this uncertainty. This directly contradicts the controllability model, but also goes beyond the behavior exhibited by the elastic controllability model. To capture this aspect of participants' behavior, we modified the elastic controllability model (elastic controllability*) such that it would increase its elasticity estimates when controllability is reduced. Model comparison consistently supported this modification (Initial study: log Bayes Factor = 143 Replication study: log Bayes factor = 155). More importantly, both the original and modified elastic controllability models explained participants' behavior better than the controllability model (Initial: log Bayes Factor ≥ 270, Replication: log Bayes Factor ≥ 148; **Figure 5E**). These results establish that people infer the elasticity of control, and thereby adapt their resource allocation to the environment's controllability and its elasticity.

Individual elasticity biases explain misallocation of resources

To determine whether the elastic controllability* model could identify individual biases in elasticity inference, we fitted the model to participants' choices and thus derived each participant's beliefs about controllability ($\gamma_{\text{controllability}}$) and elasticity ($\gamma_{\text{elasticity}}$) prior to visiting a given planet. $\gamma_{\text{controllability}}$ is used to initialize the model's belief about the maximum available control (a_{control}, b_{control}), with higher values reflecting stronger prior beliefs that planets are controllable. $\gamma_{\text{elasticity}}$ is used to initialize the model's elasticity estimates (a_{elastic}, b_{elastic}), with higher values reflecting stronger prior beliefs that control is elastic (see Methods, pg. 22). The stronger the belief, the greater its influence on choices to opt in ($\gamma_{\text{controllability}}$) and purchase extra tickets ($\gamma_{\text{elasticity}}$) regardless of observed outcomes. In addition to these two model parameters, we examined four additional parameters that directly influence opting-in and the purchase of extra tickets. These include an inverse temperature parameter (β), which specifies the stochasticity with which expected values affect resource investment choices, and baseline-preference parameters for purchasing 1, 2, or 3 tickets (α1, α2, α3). Across both samples, shared variance among parameters never exceeded 14% (**Supplementary Figure S1**), and was particularly low for the controllability and elasticity priors (Initial study: 5%, Replication study: 3%). Thus, a tendency to assume the environment is controllable did not imply a tendency to assume that control is elastic, establishing elasticity and controllability inferences as distinct facets of individual differences.

To examine the variance in resource investment explained by each of the parameters, we regressed participants' opt-in rates and extra ticket purchases on their fitted model parameters (**Figure 6A**). We found that a higher controllability bias (higher $\gamma_{\text{controllability}}$) was associated with more opting in (Initial study p<.001, Replication study p<.02; logistic regression) and more

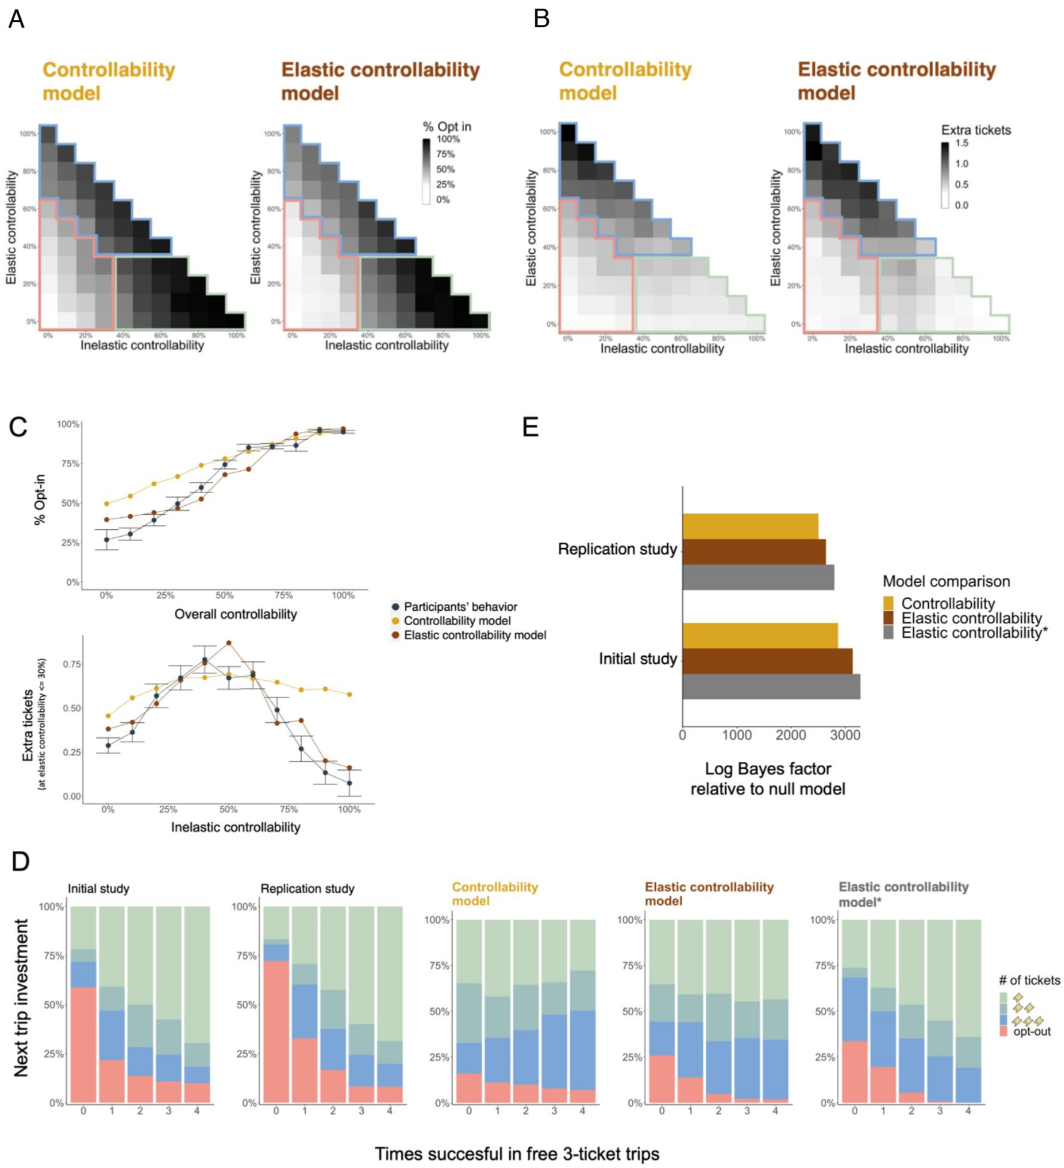

Figure 5.

Model simulations:

(A) **Opt-in**. Proportion of opt-in by the controllability and elastic controllability models across all planets. Each region is outlined by the color corresponding to the optimal ticket strategy (Blue: 3 tickets, Green: 1 ticket, Red: opt-out). (B) **Extra ticket purchases**. Extra tickets purchased by the controllability and elastic controllability models across all planets. Here too, each region is outlined by the color corresponding to the optimal ticket strategy. (C) **Comparison of opt-in rates and extra ticket purchases made by models and participants. Top -** Proportion of opting-in as a function of overall controllability. **Bottom -** The number of extra tickets purchased for each level of inelastic controllability when elasticity did not warrant purchasing extra tickets (i.e. below 30%). (D) **Resource investment following free 3-ticket trips**. The distribution of resource investment choices made immediately following the free 3-ticket trips, as a function of the outcomes of those trips. Results are shown for the two studies (Initial, Replication), the controllability model, and the original and modified variant of the elastic controllability models. (E) **Model comparison**. Consistency of each model with participant behavior is shown as log Bayes factor compared to a null model wherein participants have distinct preferences among different number of tickets to purchase but no latent controllability estimates.

extra tickets purchased (Initial $p<.001$, Replication $p<.001$; probit regression) on low-controllability, as compared to high-controllability, planets. Thus, a prior assumption that control is likely available was reflected in a futile investment of resources in uncontrollable environments.

In contrast, a higher elasticity bias was associated with the purchasing of more extra tickets on planets with inelastic, as compared to elastic, controllability (Initial study: $\beta = .18$, 95% CI = [0.03, 0.34], p = .02; Replication study: $\beta = .15$, 95%CI = [0.02 0.28], p = .02; linear regression). Thus, a prior assumption that control is likely elastic was primarily reflected in a needless investment of extra resources in environments where control could be obtained more cheaply.

Other parameters were also associated with individual differences that are consistent with the parameters' roles in the model. Specifically, the inverse temperature parameter (β) was associated with a more optimal allocation of resources, all three baseline preference parameters (α1, α2, α3) were associated with opting in, and α2 and α3, in particular, were also associated with the purchasing of more extra tickets (**Figure 6A**).

In sum, participants differed in how they allocated their resources across environments, with prior beliefs about controllability explaining misallocation of resources in uncontrollable environments, and prior beliefs about elasticity explaining misallocation of resources in environments where control was inelastic.

Individual elasticity biases are associated with psychopathology

To examine whether the individual biases in controllability and elasticity inference have psychopathological ramifications, we assayed participants on a range of self-report measures of psychopathologies previously linked to a distorted sense of control (see Methods, pg. 24). We then ran a canonical correlation analysis (CCA) to examine the relationship between the self-report measures and the model parameters underlying individual biases in task behavior. CCA is a useful method of characterizing relationships between two sets of variables, by assigning a canonical loading to each variable that maximizes the correlation between the linear combinations of each set of variables. The canonical loading of each measure thus indicates the contribution of that variable to the overall correlation. We assessed statistical significance against a null distribution obtained by permuting the model parameter relative to the psychopathology scores. This showed that the canonical correlation was significant in both the Initial (p=.01) and the Replication (p=.02) studies.

To obtain more robust estimates of the CCA loadings, we next ran a canonical correlation analysis on both datasets combined (**Figure 6C**, ρ=.33, p=.004). On the side of model parameters, the elasticity prior parameter ($\gamma_{\text{elasticity}}$) received the highest loading. By comparison, the controllability prior parameter received a lower loading with the opposite sign. A further permutation test showed that the contribution of the elasticity prior parameter to the parameter-psychopathology correlation was itself significant (p=.04; **Figure 6D**).

Loadings on the side of psychopathology were dominated by an impaired sense of agency (SOA), obsessive compulsive symptoms (OCD), and social anxiety (LSAS) – all symptoms that have been linked to an impaired sense of control[22-25]. Conversely, depression scores (SDS) received a negative loading, consistent with evidence of a lower willingness to expend one's resources in depression[21,26-27]. This psychopathological profile also manifested in raw task measures, since, like the elasticity prior parameter, it too was associated with the purchasing of more tickets on planets where controllability was less elastic (Initial: $\beta = .19$, 95% CI = [. 06, .32], p=.004; Replication: $\beta = .13$, 95% CI = [. 01, .25], p = .03; linear regression; **Figure 6A**). Thus, we found a distinct psychopathological profile involving a distorted sense of agency associated with a bias in elasticity inference.

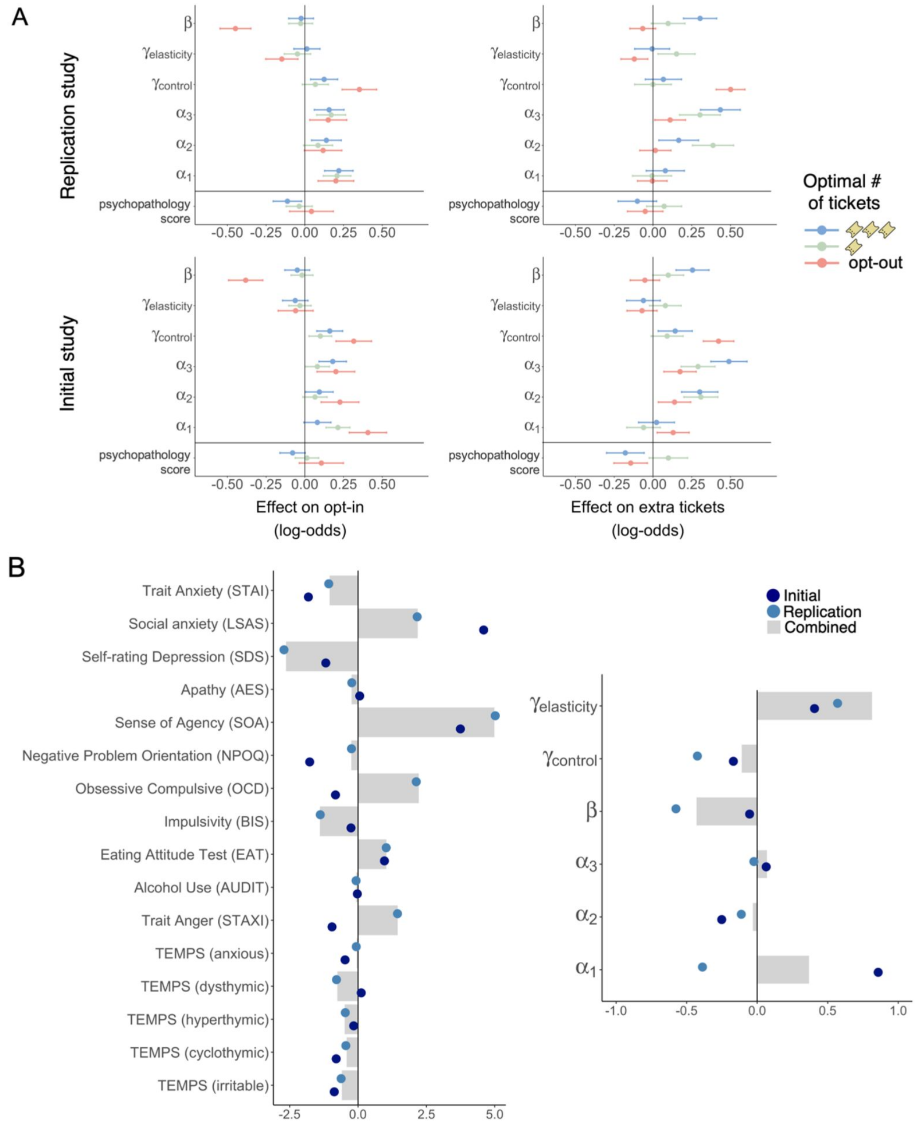

Figure 6.

Individual biases in controllability and elasticity inference:

(A) **Model parameters and task behavior**. Model parameters fitted to each participant's task behavior were used as predictors of opt-in (left; multiple logistic regressions) and extra ticket purchases (right; multiple probit regressions) separately for low controllability (red), high elastic controllability (blue) and high inelastic controllability (green) planets. In a separate analysis, participants' composite psychopathology scores (see panel C) were similarly used to predict opt-in (left; logistic regression) and extra ticket purchases (right; probit regression). Bars represent regression coefficients and 95% CIs. (B) **Model parameters and psychopathology**. Loadings from a Canonical Correlation Analysis (CCA) between model parameters and self-report psychopathology measures. Self-report measures were scored such that higher values reflect higher levels of psychopathology. Loadings are shown for the Initial (dark blue), Replication (light blue), and combined (gray bars) datasets. The magnitude of the loading reflects the degree to which each measure contributes to the correlation between parameters and self-reports. (C) **Model parameters and psychopathology composite scores**. For each participant, a composite psychopathology score was calculated by multiplying their self-report scores with their CCA-derived loadings, and a composite model parameter score was calculated by multiplying their parameter fits and corresponding loadings. Each dot represents an individual participant. Black line: group-level posterior slope. Gray shading: standard error. (D) **Significance testing**. p-values for the observed CCA correlation (circle) were computed against empirical null distributions (shaded area and bar showing 95% high-density interval). The top null distribution was obtained by shuffling the model parameter fits (top) relative to the psychopathology scores (top), and the bottom null distribution was obtained by shuffling only the elasticity prior parameter ($\gamma_{elasticity}$) while keeping all other parameters and self-report scores matched (bottom).

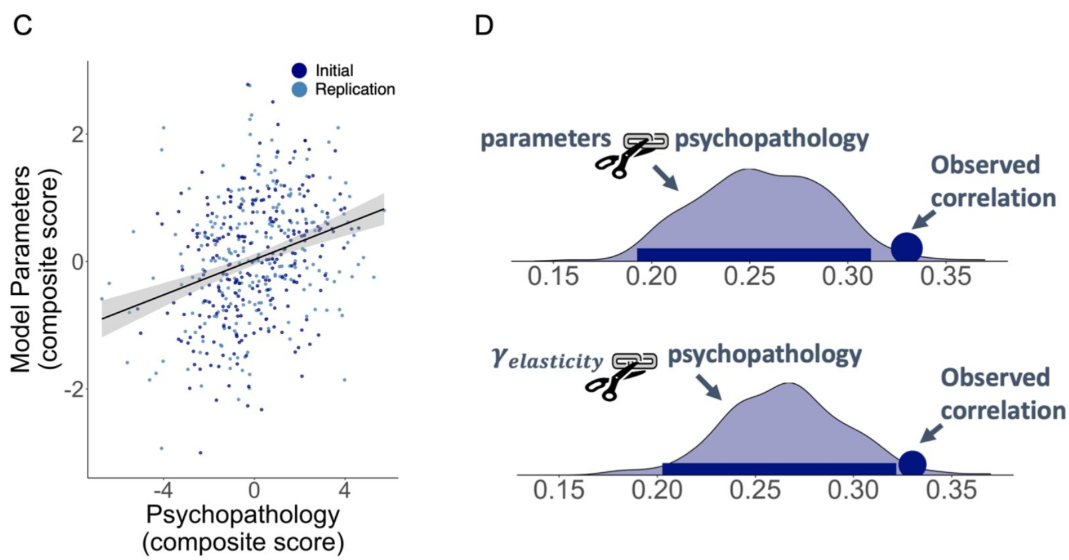

Figure 6. (continued)

Discussion

Across two pre-registered studies, we demonstrate that humans infer the elasticity of control over their present environment, and this informs their decisions of whether and how to act. Accordingly, we found that a bias to overestimate elasticity leads to a mismanagement of resources, and is associated with elevated psychopathology involving an impaired sense of agency.

Unlike prior work that treated controllability and elasticity as coupled, our studies dissociated elastic and inelastic controllability, revealing that individuals who typically assume that taking action is worthwhile (overall controllability) do not necessarily assume that investing more resources in taking action is worthwhile (the elasticity of control). In fact, elasticity and controllability biases were oppositely related to psychopathology, revealing a psychopathological profile that combines underestimation of control with overestimation of the degree to which investing more resources is necessary to obtain control.

The observed elasticity biases across individuals are particularly striking given the simplicity of the experimental task in comparison to real-life scenarios. Unlike the task, where successful boarding directly determined a trip's outcome, in many real situations many actions may contribute to a single outcome. In this regard, the present work introduces a new kind of credit-assignment problem[28], which involves identifying not only which actions influenced the outcome but also which of these actions was made more effective by investing more resources in its execution. Moreover, whereas in the task elasticity increased linearly with the number of tickets, resource investment in real life often exhibits non-linear effects. Specifically, further increases in resource investment often result in diminishing returns, and beyond a certain point, even worse outcomes (e.g. hiring too many employees, making the work less efficient). Additionally, real life typically doesn't offer the streamlined recurrence of homogenized experiences that makes learning easier in experimental tasks. These complexities introduce substantial additional uncertainty into inferences of elasticity in naturalistic settings, thus allowing more room for prior biases to exert their influences. The elasticity biases observed in the present studies are therefore likely to be amplified in real-life behavior.

Our finding that elasticity biases are associated with agency-related psychopathologies raises interesting questions for future research on the causal nature of this association. One possibility is that biased elasticity inferences distort one's sense of agency by leading to maladaptive resource allocation that results in consistently suboptimal outcomes. Another possibility is that people who already have a reduced sense of agency, attempt to compensate for this feeling by investing more resources despite minimal benefit, thus demonstrating an elasticity bias. Importantly, these possibilities are not mutually exclusive. That is, elasticity biases and a distorted sense of agency could reinforce each other in a cyclical manner. Longitudinal studies tracking or manipulating elasticity beliefs while monitoring changes in real-life resource investment and subjective sense of agency, could resolve how elasticity biases and psychopathology influence each other.

Our operationalization of controllability aligns with previous work that defined controllability as the degree to which actions influence the probability of obtaining a reward[3-4,6]. Other notable studies defined controllability as the degree to which actions predict subsequent states, independent of reward[1a]. Here we adhered to the former definition due to our present focus on resource elasticity, since individuals are unlikely to invest resources to gain control if control does not yield more reward. That said, we followed the latter work in ensuring that controllability is not confounded with predictability[1b], in the sense that in planets with high controllability, outcomes were not necessarily more predictable than in planets with low controllability (See Methods). This enabled us to specifically assess individual inferences concerning controllability and its resource elasticity.

In interpreting the present findings, it needs to be noted that we designed our task to be especially sensitive to overestimation of elasticity. We did so by giving participants free 3 tickets at their initial visits to each planet, which meant that upon success with 3 tickets, people who overestimate elasticity were more likely to continue purchasing extra tickets unnecessarily. Following the same logic, had we first had participants experience 1 ticket trips, this could have increased the sensitivity of our task to underestimation of elasticity in elastic environments. Such underestimation could potentially relate to a distinct psychopathological profile that more heavily loads on depressive symptoms. Thus, by altering the initial exposure, future studies could disambiguate the dissociable contributions of overestimating versus underestimating elasticity to different forms of psychopathology.

Another interesting possibility is that individual elasticity biases vary across different resource types. For instance, a given individual may assume that controllability tends to be highly elastic to money but inelastic to effort. Future studies could explore this possibility by developing tasks that compare individuals' elasticity beliefs across different resource types (e.g., money, time, effort). This would clarify whether elasticity biases are domain-specific or domain-general, and thus elucidate their impact on everyday decision-making.

Methods

Pre-registration

Both the Initial and Replication studies were pre-registered prior to data collection. The pre-registrations included hypotheses, sample sizes, exclusion criteria, task design, analysis plans (regressions, CCA), and computational models. Detailed pre-registration documents are available at *https://aspredicted.org/Q7M_BHR* and *https://aspredicted.org/CHW_12H*.

Participants

We used the Prolific online platform (*prolific.com*) to recruit participants (Initial study N = 264, Replication study N = 250). Participants were selected based on their demonstrated high engagement levels in prior studies, fluency in English, and absence of psychiatric diagnoses. Due to cultural differences in the perception of controllability[29] participants were recruited from English speaking western countries. The average age of participants across both datasets was 33 years old (SE =.6), and 38% were female.

The sample size was determined using a power analysis as described below. Exclusion criteria were applied to ensure data quality and participant engagement. Specifically, participants were excluded if they demonstrated inattentiveness (defined as more than 8 mistakes on comprehension quizzes during the instruction phase), displayed less than 90% accuracy in selecting the correct vehicle corresponding to its pre-learned destination (indicating a lack of task comprehension), or chose not to purchase any tickets in 90% or more of trials across all experimental conditions (suggesting a lack of engagement with the task). These criteria led to the exclusion of 22 participants (8.3% exclusion rate) from the initial study and 23 participants (8.9% exclusion rate) from the replication study. Participants gave written informed consent before taking part of the study, which was approved by the university's ethics review board (approval number: 2022-06213). Participants were paid at a rate of 9 euros per hour.

Experimental Task

Participants were informed about the location of a treasure (150 coins, worth 15¢ of bonus monetary compensation) and could exercise control by boarding the train to reach the house, or the plane to reach the mountain. (see main text and **Figure 2**). Forgoing or failing to board resulted in walking to the nearest location which happened to be where the treasure was located 20% of the time.

The first ticket (cost = 40 coins) allowed participants to select their desired vehicle, but it only stopped for them at the platform some percentage of the time, corresponding to the inelastic controllability of the planet. Participants could purchase up to 2 additional tickets (cost = 20 coins each) to attempt boarding the moving vehicle after it had left the platform. This improved the probability of boarding in correspondence with the elastic controllability of the planet.

When participants purchased a second or third ticket, the chosen vehicle appeared moving from left to right across the screen, and participants attempted to board it by pressing the spacebar when it reached the center of the screen. The precision of each boarding attempt was quantified as the absolute distance of the vehicle from the screen's center at the time of the press. Unless elastic control was below 15%, boarding attempts with precision in the top 15 percentile relative to the participant's prior performance automatically succeeded, and unless elastic control was above 85%, those in the bottom 15 percentile automatically failed. For the remaining attempts, the outcome was probabilistically determined such that the overall probability of successfully boarding matched the elasticity of control on that planet. To verify that participants' pressing ability did not bias their preference for purchasing extra tickets, we ran a mixed probit regression of the number of extra tickets purchased on participants' average press precision. We found no significant effect ($p = .45$).

At the end of each trip, participants were shown where they arrived, allowing them to infer whether they successfully boarded their vehicle. Since this feedback is only interpretable if participants knew each vehicle's destination, we only included participants in subsequent analyses if they chose the vehicle corresponding to the destination of the treasure with at least 90% accuracy (Mean 97% ±2%, see first exclusion criteria).

To allow sufficient learning of each planet's elasticity and controllability, participants took 30 consecutive trips on each planet. This trip count was determined based on pilot data (N=19) analyzing when participants' strategies stabilized. We examined the standard deviation in ticket choices over the final 5 trips and found stabilization after around 15 trips (STD = 0.51). Based on this, we provided an initial 15 trips for learning, then analyzed ticket choices starting from trip 16 onward per planet. This approach followed our pre-registered plan.

To homogenize learning experience across participants, we gave participants three free tickets for the first five visits to every planet. However, only four of those trips were informative, since on one of every five trips (randomly placed), the treasure was located in the nearby location, making boarding inconsequential for reaching it.

Planet types

To study elasticity learning across all possible degrees of controllability and elasticity, 66 unique planet conditions were generated by systematically varying inelastic controllability (the probability of the chosen vehicle stopping at the platform if purchasing a single ticket) and elasticity (the added probability of boarding from purchasing 2 extra tickets) from 0 to 100% in 10% increments. The treasure value (150 coins) and ticket costs (40 coins for the first ticket, 20 coins for extras) were chosen so that the optimal strategy, as determined by expected value calculations (**equation 11**), was to purchase 0 tickets on 22 planets (low controllability condition), 1 ticket on a different set of 22 planets (high inelastic controllability condition), and 3 tickets on the remaining 22 planets (high elastic controllability condition). Participants were randomly assigned one planet from each condition, while ensuring they could earn roughly the same total number of coins if they made optimal choices (Initial study= 45±.3, Replication study 46 ± .3). To address potential confounds between controllability and predictability[1b], we designed the experiment such that high-controllability planets (where purchasing 1 or 3 tickets was optimal) were not inherently more predictable than low-controllability planets. We quantified predictability as the entropy of the probability distribution determining whether the participant will or will not arrive at the rewarded location. For 60% of participants, purchasing tickets on high-controllability planets decreased entropy (i.e., increased predictability) compared to low-controllability planets, whereas for the other 40% it increased entropy (i.e., decreased predictability). To increase sensitivity to individual differences, all participants first encountered a planet where purchasing any number of tickets (0, 1, 2, or 3) yielded the same expected value (EV = 30 coins).

Initial training

Prior to engaging in the task, participants completed 90 practice trials. These practice trials were divided equally across planets with high inelastic controllability, high elastic controllability, and low controllability. To reinforce participants' understanding of how elasticity and controllability were manifested in each planet, they were informed of the planet type they had visited after every 15 trips.

Power analysis

The sample sizes for both the initial and the replication studies were pre-registered before data collection commenced. For the initial study, we derived the sample size of 264 participants using a power analysis aiming to estimate with high accuracy (credible interval width < 0.1) the likelihood that participants will opt-in for any combination of elastic and inelastic control (66 total combinations). For this purpose, we used Kruschke's (2014) procedure for Bayesian Power Analysis[30]. Based on pilot data (N=19), we fitted a Bayesian Beta Binomial model to the proportion of opt-ins in the planet with no optimal strategy (and therefore the highest variability in behavior would be expected), and found that 12 participants were required for the credible interval of the binomial proportion parameter to be lower than 0.1. Since each participant completes 3 combinations of elastic and inelastic control, we multiplied the required sample size by the total number of planets in each condition (22) which results in an estimate of n=264.

To examine whether this sample size is sufficient to detect significant effects of elastic and inelastic controllability on opting in (analysis 1), and the purchasing of extra tickets (analysis 2; see pg. 4-5), we used a bootstrap procedure whereby we sampled 264 participants from the pilot data and ran the regression analyses on them. We repeated this procedure for 50 iterations, each time recording $\beta_{elastic}$, $\beta_{inelastic}$, and their p-values. For analysis 1, we found that both $\beta_{elastic}$ *and* $\beta_{inelastic}$ were significantly above 0 in all iterations, and for analysis 2, $\beta_{elastic}$ was significantly above 0 in all iterations, thereby giving us >98% power to detect all of our effects of interest. Confidence intervals were derived from a binomial test in which a positive significant coefficient counted as a success.

For the replication study, we derived the sample size of 250 participants using a power analysis aiming for 95% power to detect a significant ($p <.05$) canonical correlation between the model-derived parameters and questionnaire scores (see analysis on pg. 12). For this purpose, we sampled with replacement from the initial dataset (n=264) which had showed a significant canonical correlation (p=.01; see results pg. 13), and calculated the p-value of a CCA in the subsample. We repeated this procedure for 100 iterations, each time recording the calculated p-value. To ensure our sample size was sufficient to detect significant effects in the regression models predicting opt-in and extra ticket purchases from the composite psychopathology score with over 95% power (described on pg. 13) we conducted another bootstrap analysis. We sampled 250 participants with replacement each time calculating the composite score and ran the regression analyses on that score, repeating this procedure 100 times. We found that $\beta_{composite}$ was significant ($p<.05$) in 95% of iterations (95% power).

Regression analyses

Mixed model regression analyses were performed using the "ordinal" package which fits cumulative mixed models via Laplace approximations[31]. The model predicts the likelihood of purchasing tickets (logistic model) and how many extra tickets are purchased (probit model) on the last 15 trips on each planet. Predictors included fixed and random effects for elastic and inelastic controllability and fixed and random intercepts per participant and planet number, with the latter included to account for potential learning or fatigue effects. Statistical significance was assessed using the Wald Z Test[32] with normality assessed using the Shapiro-Wilk Test[33].

Computational Modeling

To assess whether participants inferred elasticity or merely learned the control afforded by each course of action separately, we fit two competing reinforcement-learning models ('*controllability*' and '*elastic controllability*'). Both models use beta distributions defined by parameters (*a*, *b*), to form controllability beliefs for each ticket quantity n (n $\in$ {1, 2, 3}), but only the elastic controllability model leverages the dependencies that exist between different levels of resource investment. Below, we expand on the computations underlying both models.

Controllability model

The controllability model accumulates the number of times that each ticket quantity *n* led to successful or unsuccessful boarding:

$$\begin{aligned} a_n &\leftarrow a_n + O \\ b_n &\leftarrow b_n + (1 - O) \end{aligned} \qquad \textit{(equation 1)}$$

Where O=1 if boarding was successful and O=0 if not. The probability of boarding with 1,2 or 3 tickets on planet s is thus simply the expected value of that particular beta distribution:

$$p(\text{boarding}|n, s) = \frac{a_n}{a_n + b_n}. \qquad \textit{(equation 2)}$$

Here and in all subsequent models, learning only takes place when the treasure is not located on the nearby planet, which makes it possible to identify boarding success.

Elastic controllability model

The *elastic controllability model* explains ticket choices based on latent beliefs about the controllability and elasticity of each planet (see **Figure 4**). These beliefs are represented by three beta distributions, each defined by two parameters (*a*, *b*). The expected value of one beta distribution (defined by $a_{\text{control}}, b_{\text{control}}$) represents the belief that boarding is possible (controllability) whereas the expected value of the other two beta distributions represent the belief that successful boarding is added by purchasing at least one ($a_{\text{elastic}\geq 1}$, b_{elastic2}) or specifically two (a_{elastic2}, b_{elastic2}), extra tickets (elasticity).

Thus, the parameter a_{control} accumulates the number of successful boarding attempts with any number of tickets, whereas b_{control} accumulates the number of failed boarding attempts despite purchasing the maximum number of tickets (3 tickets):

$$\begin{aligned} a_{\text{control}} &\leftarrow a_{\text{control}} + 1 \text{ if win with 1,2, or 3 tickets} \\ b_{\text{control}} &\leftarrow b_{\text{control}} + 1 \text{ if lose with 3 tickets} \end{aligned} \qquad \textit{(equation 3)}$$

A failure to board with fewer tickets doesn't reduce the maximum controllability estimate but instead suggests that more tickets or the maximum number of tickets may be needed to obtain control:

$$\begin{aligned} a_{\text{elastic}\geq 1} &\leftarrow a_{\text{elastic}\geq 1} + 1 \text{ if lose with 1 or 2 tickets} \\ a_{\text{elastic2}} &\leftarrow a_{\text{elastic2}} + 1 \text{ if lose with 2 tickets} \end{aligned} \quad \textit{(equation 4)}$$

Conversely, successfully boarding with fewer tickets is evidence that obtaining controllability does not require maximal or any extra tickets:

$$\begin{aligned} b_{\text{elastic}\geq 1} &\leftarrow b_{\text{elastic}\geq 1} + \text{if win with 1 ticket} \\ b_{\text{elastic2}} &\leftarrow b_{\text{elastic2}} + 1 \text{ if win with 2 tickets} \end{aligned} \quad \textit{(equation 5)}$$

To account for our finding that failure with 3 tickets made participants favor 3, over 1 and 2, tickets, we introduced a modified elastic controllability* model, wherein purchasing extra tickets is also favored upon receiving evidence of low controllability (loss with 3 tickets). This effect was modulated by a free parameter *k*:

$$\begin{aligned} a_{\text{more}} &\leftarrow a_{\text{elastic1}} + \kappa \text{ if lose with 3 tickets} \\ a_{\text{elastic2}} &\leftarrow a_{\text{elastic2}} + \kappa \text{ if lose with 3 tickets} \end{aligned} \quad \textit{(equation 6)}$$

The estimated controllability and elasticity, as characterized by the expected value of the three beta distributions ($C_{\text{total control}}, C_{\text{more}}, C_{\text{full}}$), is used to derive the probability of boarding the selected vehicle with 1, 2, or 3 tickets. Thus, the estimated probability of boarding with 1 ticket is an estimate of inelastic control, that is, the percentage of the total control $C_{\text{total control}}$ that does not require investing more resources:

$$P(\text{boarding}|n = 1) = C_{\text{total control}} \times (1 - C_{\text{more}}) \quad \textit{(equation 7)}$$

For 2 tickets, the added probability of boarding is an estimate of what percentage of total control is elastic (C_{more}) but does not require investing maximal resources (1 ™ C_{full}):

$$P(\text{boarding}|n = 2) = P_s(\text{boarding}|n = 1) + C_{\text{total control}} \times C_{\text{more}} \times (1 - C_{\text{full}}) \quad \textit{(equation 8)}$$

Finally, since 3 tickets is the maximum resource investment, the estimated probability of boarding with 3 tickets is equivalent to an estimate of total attainable control.

$$P(\text{boarding}|n = 3, s) = C_{\text{total control}} \quad \textit{(equation 9)}$$

Models' resource investment

Because participants win 20% of the time even when they fail to board their selected ride, we calculate the probability of arriving at the treasure with *n* tickets as:

$$P(\text{win}|n) = P(\text{boarding}|n) + .2 \times (1 - P(\text{boarding}|n)) \quad \textit{(equation 10)}$$

Then, taking the possible reward (R = 150) and ticket costs (C = {40,20,20}) into account, the expected value for purchasing n tickets (*n* = {1,2,3}) tickets is:

$$V(n) = (R - \textstyle\sum_{i=1}^{n} C_i) \times P(\text{win}|n) - \textstyle\sum_{i=1}^{n} C_i \times (1 - P(\text{win}|n)) \quad \textit{(equation 11)}$$

One exception to this rule is the expected value of opting out, V(0), which is similar across all planets, and equals the potential reward (150) times the probability that walking will lead to the reward (0.2).

The probability that a participant will purchase 0, 1, 2, or 3 tickets is calculated using the SoftMax function such that:

$$p(\mathrm{n}) \propto \beta \times \mathrm{V}(\mathrm{n}) + \alpha_{\mathrm{n}} + \rho \times \mathrm{I}(\mathrm{n}_{\mathrm{t}-1} == \mathrm{n}) \qquad \textit{(equation 12)}$$

where β is an inverse temperature parameter, $\alpha_{1-;3}$ are parameters that represent participants' fixed tendencies to purchase 1, 2, or 3 tickets respectively, and ρ is a perseveration parameter that represents participant tendency to purchase the same number of tickets as in the last trial (n_{t-1}). A_0 is set to 0 to maintain parameter identifiability.

We compared the aforementioned models to a null model wherein participants have distinct preferences among different number of tickets (α_{1-3}) and perseveration (ρ) but no latent controllability estimates (β = 0 in **equation 12**). We also evaluated variants of both models that replaced the SoftMax-based exploration (governed by inverse temperature β) with a decaying epsilon-greedy approach. This epsilon-greedy variant, which begins with a fixed exploration rate that decreases with the number of trials[57], performed significantly worse at explaining participants' choices compared to the SoftMax-based models.

Prior controllability and elasticity beliefs

To capture prior biases that planets are controllable and elastic, we introduced parameters $\gamma_{\text{controllability}}$ and $\gamma_{\text{elasticity}}$, each computed by multiplying the direction (λ – 0.5) and strength (ϵ) of individuals' prior belief. $\Lambda_{\text{controllability}}$ and $\lambda_{\text{elasticity}}$ range between 0 and 1, where values above 0.5 indicate a bias towards high controllability and elasticity, and values below 0.5 indicate a bias towards low controllability and elasticity, whereas $\epsilon_{\text{controllability}}$, $\epsilon_{\text{elasticity}}$ (ϵ > 0) capture the confidence in the bias. These parameters were thus used to initialize the beta distributions representing participants' beliefs:

$$\begin{aligned}
&\mathrm{a}_{\text{controllability}} = 2 \times \epsilon_{\text{controllability}} \times \lambda_{\text{controllability}} \\
&\mathrm{b}_{\text{controllability}} = \epsilon_{\text{controllability}} \times (1 - \lambda_{\text{controllability}}) \\
&\mathrm{a}_{\text{more,full}} = \epsilon_{\text{elasticity}} \times \lambda_{\text{elasticity}} \\
&\mathrm{b}_{\text{more,full,}} = \epsilon_{\text{elasticity}} \times (1 - \lambda_{\text{elasticity}}) \qquad \textit{(equation 13)}
\end{aligned}$$

$a_{\text{controllability}}$ is multiplied by 2 so that in the absence of a bias, it would have equal preference for purchasing 0 to 3 tickets (this requires that $\mathrm{C}_{\text{controllability}} = \frac{2}{3}$, $\mathrm{C}_{\text{more,full}} = \frac{1}{2}$.

For consistency, in the ‘controllability model', the controllability estimates for each ticket quantity are initialized for each ticket quantity n, where n={1,2,3}, as:

$$\begin{aligned}
&\mathrm{a}_{\mathrm{n}} = \epsilon_{\mathrm{n}} \times \lambda_{\mathrm{n}} \\
&\mathrm{b}_{\mathrm{n}} = \epsilon_{\mathrm{n}} \times (1 - \lambda_{\mathrm{n}}) \qquad \textit{(equation 14)}
\end{aligned}$$

Here, to ensure no preference between different ticket quantities in the absence of a bias, a_1 is multiplied by 2 and a_2-is multiplied by $\frac{3}{2}$.

Model fitting

Reinforcement-learning models were fitted to participants' choices via an expectation maximization approach used in previous work[34-36]. We generated random parameter settings drawn from predefined group-level distributions and calculated the likelihood of observing the participants' choices given each parameter setting. We then approximated the posterior estimates of the group-level prior distributions for each parameter by resampling the parameter values, weighted by their respective likelihoods, and refitted the data based on the updated priors. This iterative process of resampling and refitting continued until no further improvement to model fit.

To determine the best-fitting parameters for each individual participant, we computed a weighted average of the final batch of parameter settings, where each setting was weighted by the likelihood it assigned to the individual participant's choices.

Model parameter β was initialized by sampling from a Gamma distribution ($k = 1$, $\theta = 1$), λ was initialized by sampling from a Beta distribution ($\alpha = 1$ *and* $\beta = 1$), ϵ was initialized by sampling from a Lognormal distribution ($\mu = 0$, $\sigma = 1$)., and all other parameters (α_{1-3}, ρ, k) were initialized by sampling from a Normal distribution ($\mu = 0$, $\sigma = 1$).

Validation of model fitting procedure

To validate the model fitting procedure, we simulated data using each model and using the model comparison procedure to recover the correct model (S1 table). Furthermore, we validated our parameter fits by simulating data using the best fitting parameters for each parameter and then recovering those parameters. The correlations for the parameters of interest were all equal or greater than $r = .74$ and at least .63 for all other parameters (see Table S2).

Model comparison

To compare between models, in terms of how well each model accounted for participants' choices, we employed the integrated Bayesian Information Criterion (iBIC)[37-38]. First, we estimated the evidence for each model (L) by calculating the mean likelihood of the model given 200,000 random parameter settings sampled from the fitted group-level prior distributions. Subsequently, we computed the iBIC by penalizing the model evidence to account for model complexity, following this formula: iBIC = 2 ln L + k ln n, where k represents the number of fitted parameters, and n is the number of subject choices used to compute the likelihood. Lower iBIC values indicate a more parsimonious and better-fitting model.

Model simulations

To examine the predictions of the 'controllability' and 'elastic controllability' models, we simulated choices using the best-fitting parameter settings obtained from participants' actual choices. For generality, planets were randomized in the simulation as they were in the actual experiment, such that a simulated participant did not visit precisely the same planets as the actual participant. Furthermore, to avoid discrepancies between models, each planet was repeated the same number of times for both models. 2000 simulations were performed with each model.

Individual differences

We examined individual biases in elasticity inference using participants' best-fitting parameters for beliefs about controllability ($\gamma_{controllability}$) and elasticity ($\gamma_{elasticity}$). In addition to these focal parameters, we also examined four other parameters that directly influence opt-in decisions and extra ticket purchases: an inverse temperature/noise parameter (β) and three baseline preference parameters (α1, α2, α3).

We then regressed participants' actual opt-in (logistic regression) and extra ticket purchases (probit regression) in each of the 3 controllability conditions on participants' parameters. Regressions were performed in the lme4 package in R using Rstudio[39]. Statistical significance was assessed using the Wald Z Test[32] with normality assessed using the Shapiro-Wilk Test[33].

Self-report measures

To assess trait-level agency beliefs and related psychopathologies, we administered the Sense of Agency Scale (SOA)[40]. Additionally, we measured dysfunctional attitudes toward problem-solving via the Negative Problem Orientation Questionnaire (NPOQ) [41], and temperamental differences (five domains: hyperthymic, dysthymic, cyclothymic, irritable, and anxious) via the Temperament Evaluation Questionnaire (TEMPS-A)[42-43].

Additionally, we incorporated a selection of questionnaires used by Gillan et al. (2016)[44], which assess three factors—Anxious-depression, Compulsive behavior and intrusive thought, and Social withdrawal—linked to goal-directed control deficits. These include the Apathy Evaluation Scale[45] (AES), Eating Attitudes Test[46] (EAT), Alcohol Use Disorders Identification Test[47] (AUDIT), Barratt Impulsiveness Scale[48] (BIS), Liebowitz Social Anxiety Scale[49] (LSAS), Self-Rating Depression Scale[50] (SDS), Trait Subscale of the State-Trait Anxiety Inventory for Adults[51] (STAI; Spielberger, 1983), and the revised Obsessive-Compulsive Inventory[52] (OCI-R). We utilized abbreviated versions of these questionnaires following the methodology of Wise and Dolan (2020)[53], who refined the questionnaires using a classifier to retain the items most contributory to variance in the three factors without significant information loss.

As preregistered, we aggregated the scores of the selected items from each questionnaire independently for inclusion in our CCA. Questionnaires were scored following the corresponding scoring guide, but in such a way that larger scores always indicate greater psychopathology. Based on best practice recommendations[54-55], participants were screened for inattentive responding with infrequency items ('I competed in the 1917 Summer Olympic Games') which resulted in the removal of 11 participants (Initial: 6, Replication: 5).

Directly using the scores from questionnaires of varying lengths in the CCA can give undue influence to short questionnaires, whose scores are likely more noisy. To mitigate this problem, we adjusted the standard deviation of each questionnaire's score proportionally to the square root of its number of items before conducting the CCA (in accordance with the formula for standard error). After performing the CCA on these weighted scores and normalized model parameters, the resulting canonical loadings were transformed back to their original standardized scale, ensuring that the reported loadings accurately reflected the original, unweighted scale of the questionnaire scores.

Canonical Correlation Analysis (CCA)

Canonical correlation analysis was implemented using the robust CCA method based on projection pursuit ('ccAPP' package;)[56]. We searched for the single canonical dimension that maximizes the Spearman correlation between the model parameters and psychopathology scores. Significance testing was conducted using a two-tailed permutation test with 1,000 iterations. In each iteration, the model parameters were shuffled relative to the psychopathology scores, and the canonical correlation was recalculated. The observed canonical correlation was then compared to this null distribution to obtain a p-value. To evaluate the contribution of the elasticity prior parameter ($\gamma_{\text{elasticity}}$) to the parameter-psychopathology correlation, we performed an additional permutation test. In this test, only the elasticity bias estimates ($\gamma_{\text{elasticity}}$) were shuffled, while all other model parameters and self-report scores remained constant.

To assess how the identified psychopathological profile manifests in raw task measures, we calculated a composite psychopathology score for each participant. This score was derived by multiplying individual participant scores by CCA-derived loadings, representing their position on the psychopathology spectrum. We then used this composite score as a predictor ($\beta_{\text{composite}}$) in a linear regression model to forecast the difference in average opt-in and extra ticket behavior between conditions where controllability was high and inelastic versus high and elastic.

Supplementary Information

S1 Table.

Model validation.

10 full experimental datasets were simulated using each model. Rows indicate the model used to simulate data and columns indicate the model recovered from the data using the model comparison procedure.

	Controllability	**Elastic controllability**	**Null**
Controllability	10	0	0
Elastic controllability	0	10	0
Null	0	0	10

S2 Table.

Parameter Recovery.

We validated our model fitting procedure by simulating data using the best fitting parameters for each subject and then recovering those parameters. Our correlation between simulated and recovered parameters was at least .74 for all parameters of interest that capture the effects of the experimental conditions, and at least .63 for all other parameters.

Parameter	**Correlation between Simulated and Recovered Parameters**
$\boldsymbol{\beta}$	**.91**
$\boldsymbol{\gamma}_{\mathbf{controllability}}$	**.77**
$\boldsymbol{\gamma}_{\mathbf{elasticity}}$	**.74**
$\boldsymbol{\alpha}_{\mathbf{1}}$	**.82**
$\boldsymbol{\alpha}_{\mathbf{2}}$	**.81**
$\boldsymbol{\alpha}_{\mathbf{3}}$	**.81**
$\boldsymbol{\rho}$	**.89**
$\boldsymbol{\kappa}$	**.63**

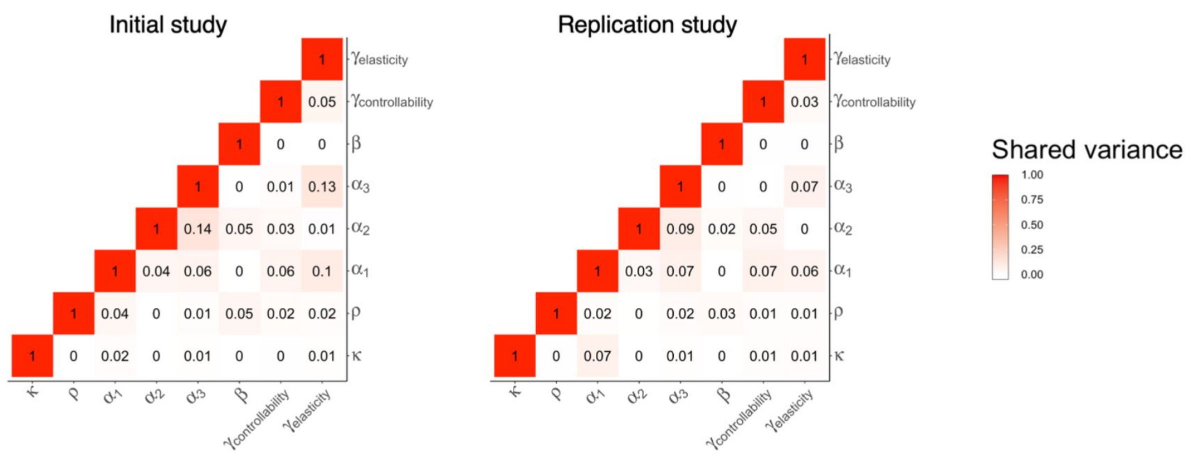

S1 Figure.

Shared variance between elastic controllability* model parameters.

Heatmaps display squared Pearson correlations between parameter pairs for initial (left) and replication (right) studies. Color intensity indicates shared variance magnitude (0-1). Parameters: $\gamma_{\text{controllability}}$, $\gamma_{\text{elasticity}}$, β, α_{1-3}, ρ, κ. Diagonal elements=1; off-diagonal elements reveal parameter interdependencies.

References

1. Ligneul R, Mainen ZF, Ly V, Cools R. 2022) **Stress-sensitive inference of task controllability** *Nat Hum Behav* **6**:812–822 https://doi.org/10.1038/s41562-022-01306-w

2. Raab HA, Foord C, Ligneul R, Hartley CA 2022) **Developmental shifts in computations used to detect environmental controllability** *PLoS Comput Biol* **18**:e1010120

3. Liljeholm M. 2021) **Agency and goal-directed choice** *Current Opinion in Behavioral Sciences* **41**:78–84

4. Moscarello J. M., Hartley C. A. 2017) **Agency and the Calibration of Motivated Behavior** *Trends in Cognitive Sciences* **21**:725–735

5. Dorfman H. M., Gershman S. J. 2019) **Controllability governs the balance between Pavlovian and instrumental action selection** *Nature Communications* **10**:1–8 https://doi.org/10.1038/s41467-019-13737-7

6. Wang K. S., Yang Y. Y., Delgado M. R. 2021) **How perception of control shapes decision making** *Current Opinion in Behavioral Sciences* **41**:85–91

7. Chambon V, The'ro H, Vidal M, Vandendriessche H, Haggard P, Palminteri S 2020) **Information about action outcomes differentially affects learning from self-determined versus imposed choices** *Nat Hum Behav* **4**:1067–1079

8. Diener C., Kuehner C., Brusniak W., Struve M., Flor H. 2009) **Effects of stressor controllability on psychophysiological, cognitive and behavioural responses in patients with major depression and dysthymia** *Psychol. Med* **39**:77–86

9. Moulding R., Kyrios M. 2006) **Anxiety disorders and control related beliefs: the exemplar of Obsessive– Compulsive Disorder (OCD)** *Clin. Psychol. Rev* **26**:573–583

10. Gallagher MW, Naragon-Gainey K, Brown TA 2014) **Perceived Control is a Transdiagnostic Predictor of Cognitive-Behavior Therapy Outcome for Anxiety Disorders** *Cognit Ther Res* **38**:10–22 https://doi.org/10.1007/s10608-013-9587-3

11. Salkovskis P. M., Rapee R.M. 1996) **Cognitive-behavioral approaches to the understanding of obsessional problems** *Current Controversies in the Anxiety Disorders* New York: Guilford Press :103–133

12. McLaren S., Crowe S. F. 2003) **The contribution of perceived control of stressful life events and thought suppression to the symptoms of obsessive–compulsive disorder in both non-clinical and clinical samples** *J. Anxiety Disord* **17**:389–403 https://doi.org/10.1016/S0887-6185(02)00224-4

13. Pindyck R. S. 2018) **Microeconomics**

14. Patall E. A., Cooper H., Robinson J. 2016) **Control and Effort Costs Influence the Motivational Consequences of Choice** *Frontiers in Psychology* **8**:1–12 https://doi.org/10.3389/fpsyg.2016.00675

15. Schunk D. H. 1982) **Effects of effort attributional feedback on children's perceived self-efficacy and achievement** *Journal of Educational Psychology* **74**:548–556

16. Frömer R., Lin H., Dean Wolf C.K., Inzlicht M., Shenhav A. 2021) **Expectations of reward and efficacy guide cognitive control allocation** *Nature communications* **12**:1030

17. Cao S., Tang C., Wu H., Liu X. 2022) **Value analysis determines when and how to strive** *Advances in Psychological Science* **30**:877–887

18. Grahek I., Frömer R., Prater Fahey M., Shenhav A. 2023) **Learning when effort matters: Neural dynamics underlying updating and adaptation to changes in performance efficacy** *Cerebral Cortex* **33**:2395–2411

19. Shenhav A., Prater Fahey M., Grahek I. 2021) **Decomposing the motivation to exert mental effort** *Current Directions in Psychological Science* **30**:307–314

20. Inzlicht M, Shenhav A, Olivola CY 2018) **The Effort Paradox: Effort Is Both Costly and Valued** *Trends Cogn Sci* **22**:337–349 https://doi.org/10.1016/j.tics.2018.01.007

21. Grahek I, Shenhav A, Musslick S, Krebs RM, Koster EHW 2019) **Motivation and cognitive control in depression** *Neurosci Biobehav Rev* **102**:371–381 https://doi.org/10.1016/j.neubiorev.2019.04.011

22. Zorowitz S., Momennejad I., Daw N. D. 2020) **Anxiety, Avoidance, and Sequential Evaluation** *Computational Psychiatry* **4**:1

23. Gillan C. M., et al. 2014) **Obsessive-compulsive disorder patients have a reduced sense of control on the illusion of control task** *Front. Psychol* **5**:204

24. Stein DJ, Costa DLC, Lochner C, Miguel EC, Reddy YCJ, Shavitt RG, van den Heuvel OA, Simpson HB 2019) **Obsessive-compulsive disorder** *Nat Rev Dis Primers* **5**:52Erratum in: Nat Rev Dis Primers. 2024 Oct 16;10(1):79. doi: 10.1038/s41572-024-00569-z. PMID: 31371720; PMCID: PMC7370844 https://doi.org/10.1038/s41572-019-0102-3

25. Beyer F, Sidarus N, Fleming S, Haggard P 2018) **Losing control in social situations: how the presence of others affects neural processes related to sense of agency** *eNeuro* **5**

26. Cohen RM, Weingartner H, Smallberg SA, Pickar D, Murphy DL 1982) **Effort and cognition in depression** *Arch Gen Psychiatry* **39**:593–7 https://doi.org/10.1001/archpsyc.1982.04290050061012

27. Bi R, Dong W, Zheng Z, Li S, Zhang D. 2022) **Altered motivation of effortful decision-making for self and others in subthreshold depression** *Depress Anxiety* **39**:633–645 https://doi.org/10.1002/da.23267

28. Sutton R. S. 1984) **Temporal Credit Assignment in Reinforcement Learning** Amherst: University of Massachusetts

29. Msetfi R. M., Kornbrot D. E., Matute H., Murphy R. A. 2015) **The relationship between mood state and perceived control in contingency learning: effects of individualist and collectivist values** *Frontiers in Psychology* **6**:1–18

30. Kruschke J. K. 2014) **Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan (2nd ed** Academic Press

31. Christensen R. H. B. 2019) **ordinal - Regression Models for Ordinal Data** *CRAN*

32. Agresti A. 2002) **Categorical Data Analysis (2nd ed** John Wiley & Sons

33. Shapiro S. S., Wilk M. B. 1965) **An Analysis of Variance Test for Normality (Complete Samples)** *Biometrika* **52**:591–611

34. Eldar E., Roth C., Dayan P., Dolan R. J. 2018) **Decodability of Reward Learning Signals Predicts Mood Fluctuations** *Curr. Biol* **28**:1433–1439

35. Eldar E., Lièvre G., Dayan P., Dolan R. J. 2020) **The roles of online and offline replay in planning** *eLife* **9**:e56911 https://doi.org/10.7554/eLife.56911

36. Solomyak L., Sharp P. B., Eldar E. 2022) **Training diversity promotes absolute-valueguided choice** *PLoS Comput. Biol* **18**:e1010664 https://doi.org/10.1371/journal.pcbi.1010664

37. Eldar E., Hauser T. U., Dayan P., Dolan R. J. 2016) **Striatal structure and function predict individual biases in learning to avoid pain** *Proc. Natl. Acad. Sci. U. S. A* https://doi.org/10.1073/pnas.1519829113

38. Huys Q. J. M., et al. 2012) **Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees** *PLoS Comput. Biol* **8**

39. R Core Team 2022) **R: A language and environment for statistical computing** Vienna, Austria: R Foundation for Statistical Computing https://www.R-project.org/

40. Tapai A, Oren E, Dar R, Eitam B. 2017) **The Sense of Agency Scale: A Measure of Consciously Perceived Control over One's Mind, Body, and the Immediate Environment** *Front Psychol* **8**:1552 https://doi.org/10.3389/fpsyg.2017.01552

41. Melisa Robichaud, Dugas Michel 2005) **Negative problem orientation (Part I): Psychometric properties of a new measure** *Behaviour research and therapy* **43**:391–401 https://doi.org/10.1016/j.brat.2004.02.007

42. Akiskal K. K., Akiskal H. S. 2005) **The theoretical underpinnings of affective temperaments: Implications for evolutionary foundations of bipolarity** *Journal of Affective Disorders* **85**:231–239

43. Akiskal HS, Mendlowicz MV, Jean-Louis G, Rapaport MH, Kelsoe JR, Gillin JC, Smith TL 2005) **TEMPS-A: validation of a short version of a self-rated instrument designed to measure variations** *J Affect Disord.* **85**:45–52

44. Gillan C. M., Kosinski M., Whelan R., Phelps E. A., Daw N. D. 2016) **Characterizing a psychiatric symptom dimension related to deficits in goaldirected control** *eLife* **5**:e11305 https://doi.org/10.7554/eLife.11305

45. Marin RS, Biedrzycki RC, Firinciogullari S. 1991) **Reliability and validity of the Apathy Evaluation Scale** *Psychiatry Res* **38**:143–62 https://doi.org/10.1016/0165-1781(91)90040-v

46. Garner DM, Olmsted MP, Bohr Y, Garfinkel PE 1982) **The eating attitudes test: psychometric features and clinical correlates** *Psychol Med* **12**:871–8 https://doi.org/10.1017/s0033291700049163

47. Saunders JB, Aasland OG, Babor TF, de la Fuente JR, Grant M. 1993) **Development of the Alcohol Use Disorders Identification Test (AUDIT): WHO Collaborative Project on Early Detection of Persons with Harmful Alcohol Consumption--II** *Addiction* **88**:791–804

48. Patton JH, Stanford MS, Barratt ES 1995) **Factor structure of the Barratt impulsiveness scale** *J ClinPsychol* **51**:768–74 https://doi.org/10.1002/1097-4679(199511)51:6<768::aidjclp2270510607>3.0.co;2-1

49. Heimberg RG, Horner KJ, Juster HR, Safren SA, Brown EJ, Schneier FR, Liebowitz MR 1999) **Psychometric properties of the Liebowitz Social Anxiety Scale** *Psychol Med* **29**:199–212 https://doi.org/10.1017/s0033291798007879

50. Zung WW 1965) **A self-rating depression scale** *Arch Gen Psychiatry* **12**:63–70 https://doi.org/10.1001/archpsyc.1965.01720310065008

51. Spielberger C. D. 1983) **State-Trait Anxiety Inventory for Adults (STAI-AD) [Database record]** *APA PsycTests*

52. Foa EB, Huppert JD, Leiberg S, Langner R, Kichic R, Hajcak G, Salkovskis PM 2002) **The Obsessive Compulsive Inventory: development and validation of a short version** *Psychol Assess* **14**:485–96

53. Wise T., Dolan R.J. 2020) **Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample** *Nat Commun* **11**:4179

54. Thomas K. A., Clifford S. 2017) **Validity and Mechanical Turk: an assessment of exclusion methods and interactive experiments** *Comput. Hum. Behav* **77**:184–197

55. Barends A. J., de Vries R. E. 2019) **Noncompliant responding: comparing exclusion criteria in MTurk personality research to improve data quality** *Pers. Individ. Differ* **143**:84–89

56. Alfons C. Croux, Filzmoser P. 2016) **Robust maximum association between data sets: The RPackage ccaPP** *Austrian Journal of Statistics* **45**:71–79

57. Kalidindi K, Bowman H. 2007) **Using epsilon-greedy reinforcement learning methods to further understand ventromedial prefrontal patients' deficits on the Iowa Gambling Task** *Neural Netw* **20**:676–89 https://doi.org/10.1016/j.neunet.2007.04.026

Author information

Levi Solomyak*
Department of Cognitive Sciences, Hebrew University of Jerusalem, Jerusalem, Israel
ORCID iD: 0000-0003-2933-607X

For correspondence: levi.solomyak@mail.huji.ac.il
*These authors contributed equally to this work

Aviv Emanuel*
Department of Cognitive Sciences, Hebrew University of Jerusalem, Jerusalem, Israel, Department of Psychology Hebrew University of Jerusalem, Jerusalem, Israel, Department of Psychology, McGill University, Montreal, Canada

*These authors contributed equally to this work

Eran Eldar
Department of Cognitive Sciences, Hebrew University of Jerusalem, Jerusalem, Israel, Department of Psychology Hebrew University of Jerusalem, Jerusalem, Israel

Editors

Reviewing Editor
Roshan Cools
Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, Netherlands

Senior Editor
Michael Frank
Brown University, Providence, United States of America

Reviewer #1 (Public review):

Summary:

The authors investigated the elasticity of controllability by developing a task that manipulates the probability of achieving a goal with a baseline investment (which they refer to as inelastic controllability) and the probability that additional investment would increase the probability of achieving a goal (which they refer to as elastic controllability). They found that a computational model representing the controllability and elasticity of the environment accounted better for the data than a model representing only the controllability. They also found that prior biases about the controllability and elasticity of the environment were associated with a composite psychopathology score. The authors conclude that elasticity inference and bias guide resource allocation.

Strengths:

This research takes a novel theoretical and methodological approach to understanding how people estimate the level of control they have over their environment, and how they adjust their actions accordingly. The task is innovative and both it and the findings are well-described (with excellent visuals). They also offer thorough validation for the particular model they develop. The research has the potential to theoretically inform the understanding of control across domains, which is a topic of great importance.

Weaknesses:

An overarching concern is that this paper is framed as addressing resource investments across domains that include time, money, and effort, and the introductory examples focus heavily on effort-based resources (e.g., exercising, studying, practicing). The experiments, though, focus entirely on the equivalent of monetary resources - participants make discrete actions based on the number of points they want to use on a given turn. While the same ideas might generalize to decisions about other kinds of resources (e.g., if participants were having to invest the effort to reach a goal), this seems like the kind of speculation that would be better reserved for the Discussion section rather than using effort investment as a means of introducing a new concept (elasticity of control) that the paper will go on to test.

Setting aside the framing of the core concepts, my understanding of the task is that it effectively captures people's estimates of the likelihood of achieving their goal (Pr(success)) conditional on a given investment of resources. The ground truth across the different environments varies such that this function is sometimes flat (low controllability), sometimes increases linearly (elastic controllability), and sometimes increases as a step function (inelastic controllability). If this is accurate, then it raises two questions.

First, on the modeling front, I wonder if a suitable alternative to the current model would be to assume that the participants are simply considering different continuous functions like these and, within a Bayesian framework, evaluating the probabilistic evidence for each function based on each trial's outcome. This would give participants an estimate of the marginal increase in Pr(success) for each ticket, and they could then weigh the expected value of that ticket choice (Pr(success)*150 points) against the marginal increase in point cost for each ticket. This should yield similar predictions for optimal performance (e.g., opt-out for lower controllability environments, i.e., flatter functions), and the continuous nature of this form of function approximation also has the benefit of enabling tests of generalization to predict changes in behavior if there was, for instance, changes in available tickets for purchase (e.g., up to 4 or 5) or changes in ticket prices. Such a model would of course also maintain a critical role for priors based on one's experience within the task as well as over longer timescales, and could be meaningfully interpreted as such (e.g., priors related to the likelihood of success/failure and whether one's actions influence these). It could also potentially reduce the complexity of the model by replacing controllability-specific parameters with multiple candidate functions (presumably learned through past experience, and/or tuned by experience in this task environment), each of which is being updated simultaneously.

Second, if the reframing above is apt (regardless of the best model for implementing it), it seems like the taxonomy being offered by the authors risks a form of "jangle fallacy," in particular by positing distinct constructs (controllability and elasticity) for processes that ultimately comprise aspects of the same process (estimation of the relationship between investment and outcome likelihood). Which of these two frames is used doesn't bear on the rigor of the approach or the strength of the findings, but it does bear on how readers will digest and draw inferences from this work. It is ultimately up to the authors which of these they choose to favor, but I think the paper would benefit from some discussion of a common-process alternative, at least to prevent too strong of inferences about separate processes/modes that may not exist. I personally think the approach and findings in this paper would also be easier to digest under a common-construct approach rather than forcing new terminology but, again, I defer to the authors on this.

https://doi.org/10.7554/eLife.105507.1.sa3

Reviewer #2 (Public review):

Summary:

In this paper, the authors test whether controllability beliefs and associated actions/resource allocation are modulated by things like time, effort, and monetary costs (what they call "elastic" as opposed to "inelastic" controllability). Using a novel behavioral task and computational modeling, they find that participants do indeed modulate their resources depending on whether they are in an "elastic," "inelastic," or "low controllability" environment. The authors also find evidence that psychopathology is related to specific biases in controllability.

Strengths:

This research investigates how people might value different factors that contribute to controllability in a creative and thorough way. The authors use computational modeling to try to dissociate "elasticity" from "overall controllability," and find some differential associations with psychopathology. This was a convincing justification for using modeling above and beyond behavioral output and yielded interesting results. Interestingly, the authors conclude that these findings suggest that biased elasticity could distort agency beliefs via maladaptive resource allocation. Overall, this paper reveals some important findings about how people consider components of controllability.

Weaknesses:

The primary weakness of this research is that it is not entirely clear what is meant by "elastic" and "inelastic" and how these constructs differ from existing considerations of various factors/calculations that contribute to perceptions of and decisions about controllability. I think this weakness is primarily an issue of framing, where it's not clear whether elasticity is, in fact, theoretically dissociable from controllability. Instead, it seems that the elements that make up "elasticity" are simply some of the many calculations that contribute to controllability. In other words, an "elastic" environment is inherently more controllable than an "inelastic" one, since both environments might have the same level of predictability, but in an "elastic" environment, one can also partake in additional actions to have additional control over achieving the goal (i.e., expend effort, money, time).

https://doi.org/10.7554/eLife.105507.1.sa2

Reviewer #3 (Public review):

A bias in how people infer the amount of control they have over their environment is widely believed to be a key component of several mental illnesses including depression, anxiety, and addiction. Accordingly, this bias has been a major focus in computational models of those disorders. However, all of these models treat control as a unidimensional property, roughly, how strongly outcomes depend on action. This paper proposes---correctly, I think---that the intuitive notion of "control" captures multiple dimensions in the relationship between action and outcome is multi-dimensional. In particular, the authors propose that the degree to which outcome depends on how much *effort* we exert, calling this dimension the "elasticity of control". They additionally propose that this dimension (rather than the more holistic notion of controllability) may be specifically impaired in certain types of psychopathology. This idea thus has the potential to change how we think about mental disorders in a substantial way, and could even help us better understand how healthy people navigate challenging decision-making problems.

Unfortunately, my view is that neither the theoretical nor empirical aspects of the paper really deliver on that promise. In particular, most (perhaps all) of the interesting claims in the paper have weak empirical support.

Starting with theory, the elasticity idea does not truly "extend" the standard control model in the way the authors suggest. The reason is that effort is simply one dimension of action. Thus, the proposed model ultimately grounds out in how strongly our outcomes depend on our actions (as in the standard model). Contrary to the authors' claims, the elasticity of control is still a fixed property of the environment. Consistent with this, the computational model proposed here is a learning model of this fixed environmental property. The idea is still valuable, however, because it identifies a key dimension of action (namely, effort) that is particularly relevant to the notion of perceived control. Expressing the elasticity idea in this way might support a more general theoretical formulation of the idea that could be applied in other contexts. See Huys & Dayan (2009), Zorowitz, Momennejad, & Daw (2018), and Gagne & Dayan (2022) for examples of generalizable formulations of perceived control.

Turning to experiment, the authors make two key claims: (1) people infer the elasticity of control, and (2) individual differences in how people make this inference are importantly related to psychopathology.

Starting with claim 1, there are three sub-claims here; implicitly, the authors make all three. (1A) People's behavior is sensitive to differences in elasticity, (1B) people actually represent/track something like elasticity, and (1C) people do so naturally as they go about their daily lives. The results clearly support 1A. However, 1B and 1C are not supported.

Starting with 1B, the experiment cannot support the claim that people represent or track elasticity because the effort is the only dimension over which participants can engage in any meaningful decision-making (the other dimension, selecting which destination to visit, simply amounts to selecting the location where you were just told the treasure lies). Thus, any adaptive behavior will necessarily come out in a sensitivity to how outcomes depend on effort. More concretely, any model that captures the fact that you are more likely to succeed in two attempts than one will produce the observed behavior. The null models do not make this basic assumption and thus do not provide a useful comparison.

For 1C, the claim that people infer elasticity outside of the experimental task cannot be supported because the authors explicitly tell people about the two notions of control as part of the training phase: "To reinforce participants' understanding of how elasticity and controllability were manifested in each planet, [participants] were informed of the planet type they had visited after every 15 trips." (line 384).

Finally, I turn to claim 2, that individual differences in how people infer elasticity are importantly related to psychopathology. There is much to say about the decision to treat psychopathology as a unidimensional construct. However, I will keep it concrete and simply note that CCA (by design) obscures the relationship between any two variables. Thus, as suggestive as Figure 6B is, we cannot conclude that there is a strong relationship between Sense of Agency and the elasticity bias---this result is consistent with any possible relationship (even a negative one). The fact that the direct relationship between these two variables is not shown or reported leads me to infer that they do not have a significant or strong relationship in the data.

There is also a feature of the task that limits our ability to draw strong conclusions about individual differences in elasticity inference. As the authors clearly acknowledge, the task was designed "to be especially sensitive to overestimation of elasticity" (line 287). A straightforward consequence of this is that the resulting *empirical* estimate of estimation bias (i.e., the gamma_elasticity parameter) is itself biased. This immediately undermines any claim that references the directionality of the elasticity bias (e.g. in the abstract). Concretely, an undirected deficit such as slower learning of elasticity would appear as a directed overestimation bias.

When we further consider that elasticity inference is the only meaningful learning/decision-making problem in the task (argued above), the situation becomes much worse. Many general deficits in learning or decision-making would be captured by the elasticity bias parameter. Thus, a conservative interpretation of the results is simply that psychopathology is associated with impaired learning and decision-making.

Minor comments:

Showing that a model parameter correlates with the data it was fit to does not provide any new information, and cannot support claims like "a prior assumption that control is likely available was reflected in a futile investment of resources in uncontrollable environments." To make that claim, one must collect independent measures of the assumption and the investment.

Did participants always make two attempts when purchasing tickets? This seems to violate the intuitive model, in which you would sometimes succeed on the first jump. If so, why was this choice made? Relatedly, it is not clear to me after a close reading how the outcome of each trial was actually determined.

It should be noted that the model is heuristically defined and does not reflect Bayesian updating. In particular, it overestimates control by not using losses with less than 3 tickets (intuitively, the inference here depends on your beliefs about elasticity). I wonder if the forced three-ticket trials in the task might be historically related to this modeling choice.

https://doi.org/10.7554/eLife.105507.1.sa1

Author response:

We thank the reviewers for their thorough reading and thoughtful feedback. Below, we provisionally address each of the concerns raised in the public reviews, and outline our planned revision that aims to further clarify and strengthen the manuscript.

In our response, we clarify our conceptualization of elasticity as a dimension of controllability, formalizing it within an information-theoretic framework, and demonstrating that controllability and its elasticity are partially dissociable. Furthermore, we provide clarifications and additional modeling results showing that our experimental design and modeling approach are well-suited to dissociating elasticity inference from more general learning processes, and are not inherently biased to find overestimates of elasticity. Finally, we clarify the advantages and disadvantages of our canonical correlation analysis (CCA) approach for identifying latent relationships between multidimensional data sets, and provide additional analyses that strengthen the link between elasticity estimation biases and a specific psychopathology profile.

> ***Reviewer 1:***
>
> *This research takes a novel theoretical and methodological approach to understanding how people estimate the level of control they have over their environment, and how they adjust their actions accordingly. The task is innovative and both it and the findings are well-described (with excellent visuals). They also offer thorough validation for the particular model they develop. The research has the potential to theoretically inform the understanding of control across domains, which is a topic of great importance.*

We thank the reviewer for their favorable appraisal and valuable suggestions, which have helped clarify and strengthen the study’s conclusion.

> *An overarching concern is that this paper is framed as addressing resource investments across domains that include time, money, and effort, and the introductory examples focus heavily on effort-based resources (e.g., exercising, studying, practicing). The experiments, though, focus entirely on the equivalent of monetary resources - participants make discrete actions based on the number of points they want to use on a given turn. While the same ideas might generalize to decisions about other kinds of resources (e.g., if participants were having to invest the effort to reach a goal), this seems*

> *like the kind of speculation that would be better reserved for the Discussion section rather than using effort investment as a means of introducing a new concept (elasticity of control) that the paper will go on to test.*

We thank the reviewer for pointing out a lack of clarity regarding the kinds of resources tested in the present experiment. Investing additional resources in the form of extra tickets did not only require participants to pay more money. It also required them to invest additional time – since each additional ticket meant making another attempt to board the vehicle, extending the duration of the trial, and attentional effort – since every attempt required precisely timing a spacebar press as the vehicle crossed the screen. Given this involvement of money, time, and effort resources, we believe it would be imprecise to present the study as concerning monetary resources in particular. That said, we agree with the Reviewer that results might differ depending on the resource type that the experiment or the participant considers most. Thus, in our revision of the manuscript, we will make sure to clarify the kinds of resources the experiment involved, and highlight the open question of whether inferences concerning the elasticity of control generalize across different resource domains.

> *Setting aside the framing of the core concepts, my understanding of the task is that it effectively captures people's estimates of the likelihood of achieving their goal (Pr(success)) conditional on a given investment of resources. The ground truth across the different environments varies such that this function is sometimes flat (low controllability), sometimes increases linearly (elastic controllability), and sometimes increases as a step function (inelastic controllability). If this is accurate, then it raises two questions.*
>
> *First, on the modeling front, I wonder if a suitable alternative to the current model would be to assume that the participants are simply considering different continuous functions like these and, within a Bayesian framework, evaluating the probabilistic evidence for each function based on each trial's outcome. This would give participants an estimate of the marginal increase in Pr(success) for each ticket, and they could then weigh the expected value of that ticket choice (Pr(success)*150 points) against the marginal increase in point cost for each ticket. This should yield similar predictions for optimal performance (e.g., opt-out for lower controllability environments, i.e., flatter functions), and the continuous nature of this form of function approximation also has the benefit of enabling tests of generalization to predict changes in behavior if there was, for instance, changes in available tickets for purchase (e.g., up to 4 or 5) or changes in ticket prices. Such a model would of course also maintain a critical role for priors based on one's experience within the task as well as over longer timescales, and could be meaningfully interpreted as such (e.g., priors related to the likelihood of success/failure and whether one's actions influence these). It could also potentially reduce the complexity of the model by replacing controllability-specific parameters with multiple candidate functions (presumably learned through past experience, and/or tuned by experience in this task environment), each of which is being updated simultaneously.*
>
> *Second, if the reframing above is apt (regardless of the best model for implementing it), it seems like the taxonomy being offered by the authors risks a form of "jangle fallacy," in particular by positing distinct constructs (controllability and elasticity) for processes that ultimately comprise aspects of the same process (estimation of the relationship between investment and outcome likelihood). Which of these two frames is used doesn't bear on the rigor of the approach or the strength of the findings, but it does bear on how readers will digest and draw inferences from this work. It is ultimately up to the authors which of these they choose to favor, but I think the paper would benefit from some discussion of a common-process alternative, at least to prevent too strong of inferences about separate processes/modes that may not exist. I personally think the approach and findings in this*

> *paper would also be easier to digest under a common-construct approach rather than forcing new terminology but, again, I defer to the authors on this.*

We thank the reviewer for suggesting this interesting alternative modeling approach. We agree that a Bayesian framework evaluating different continuous functions could offer advantages, particularly in its ability to generalize to other ticket quantities and prices. We will attempt to implement this as an alternative model and compare it with the current model.

We also acknowledge the importance of avoiding a potential "jangle fallacy". We entirely agree with the Reviewer that elasticity and controllability inferences are not distinct processes. Specifically, we view resource elasticity as a dimension of controllability, hence the name of our 'elastic controllability' model. In response to this and other Reviewers' comments, we now offer a formal definition of elasticity as the reduction in uncertainty about controllability due to knowing the amount of resources the agent is able and willing to invest (see further details in response to Reviewer 3 below).

With respect to how this conceptualization is expressed in the modelling, we note that the representation in our model of maximum controllability and its elasticity via different variables is analogous to how a distribution may be represented by separate mean and variance parameters. Ultimately, even in the model suggested by the Reviewer, there would need to be a dedicated variable representing elasticity, such as the probability of sloped controllability functions. A single-process account thus allows that different aspects of this process would be differently biased (e.g., one can have an accurate estimate of the mean of a distribution but overestimate its variance). Therefore, our characterization of distinct elasticity and controllability biases (or to put it more accurately, 'elasticity of controllability bias' and 'maximum controllability bias') is consistent with a common construct account.

That said, given the Reviewer's comments, we believe that some of the terminology we used may have been misleading. In our planned revision, we will modify the text to clarify that we view elasticity as a dimension of controllability that can only be estimated in conjunction with controllability.

> ***Reviewer 2:***
>
> *This research investigates how people might value different factors that contribute to controllability in a creative and thorough way. The authors use computational modeling to try to dissociate "elasticity" from "overall controllability," and find some differential associations with psychopathology. This was a convincing justification for using modeling above and beyond behavioral output and yielded interesting results. Interestingly, the authors conclude that these findings suggest that biased elasticity could distort agency beliefs via maladaptive resource allocation. Overall, this paper reveals some important findings about how people consider components of controllability.*

We appreciate the Reviewer's positive assessment of our findings and computational approach to dissociating elasticity and overall controllability.

> *The primary weakness of this research is that it is not entirely clear what is meant by "elastic" and "inelastic" and how these constructs differ from existing considerations of various factors/calculations that contribute to perceptions of and decisions about controllability. I think this weakness is primarily an issue of framing, where it's not clear whether elasticity is, in fact, theoretically dissociable from controllability. Instead, it seems that the elements that make up "elasticity" are simply some of the many calculations that contribute to controllability. In other words, an "elastic" environment is*

> *inherently more controllable than an "inelastic" one, since both environments might have the same level of predictability, but in an "elastic" environment, one can also partake in additional actions to have additional control overachieving the goal (i.e., expend effort, money, time).*

We thank the reviewer for highlighting the lack of clarity in our concept of elasticity. We first clarify that elasticity cannot be entirely dissociated from controllability because it is a dimension of controllability. If no controllability is afforded, then there cannot be elasticity or inelasticity. This is why in describing the experimental environments, we only label high-controllability, but not low-controllability, environments as 'elastic' or 'inelastic'. For further details on this conceptualization of elasticity, and a planned revision of the text, see our response above to Reviewer 1.

Second, we now clarify that controllability can also be computed without knowing the amount of resources the agent is able and willing to invest, for instance by assuming infinite resources available or a particular distribution of resource availabilities. However, knowing the agent's available resources often reduces uncertainty concerning controllability. This reduction in uncertainty is what we define as elasticity. Since any action requires some resources, this means that no controllable environment is entirely inelastic if we also consider agents that do not have enough resources to commit any action. However, even in this case environments can differ in the degree to which they are elastic. For further details on this formal definition, see our response to Reviewer 3 below. We will make these necessary clarifications in the revised manuscript.

Importantly, whether an environment is more or less elastic does not determine whether it is more or less controllable. In particular, environments can be more controllable yet less elastic. This is true even if we allow that investing different levels of resources (i.e., purchasing 0, 1, 2, or 3 tickets) constitute different actions, in conjunction with participants' vehicle choices. Below, we show this using two existing definitions of controllability.

Definition 1, reward-based controllability[1]: If control is defined as the fraction of available reward that is controllably achievable, and we assume all participants are in principle willing and able to invest 3 tickets, controllability can be computed in the present task as:

$$\chi = \mathrm{Max}_{A,C}\, P(S' = \mathrm{goal}|S, A, C) - \mathrm{Min}_{A,C}\, P(S' = \mathrm{goal}|S, A, C)$$

where P(S' = goal | S, A, C) is the probability of reaching the treasure from present state S when taking action A and investing C resources in executing the action. In any of the task environments, the probability of reaching the goal is maximized by purchasing 3 tickets (C = 3) and choosing the vehicle that leads to the goal (A = correct vehicle). Conversely, the probability of reaching the goal is minimized by purchasing 3 tickets (C = 3) and choosing the vehicle that does not lead to the goal (A = wrong vehicle). This calculation is thus entirely independent of elasticity, since it only considers what would be achieved by maximal resource investment, whereas elasticity consists of the reduction in controllability that would arise if the maximal available C is reduced. Consequently, any environment where the maximum available control is higher yet varies less with resource investment would be more controllable and less elastic.

Note that if we also account for ticket costs in calculating reward, this will only reduce the fraction of achievable reward and thus the calculated control in elastic environments.

Definition 2, information-theoretic controllability[2]: Here controllability is defined as the reduction in outcome entropy due to knowing which action is taken:

$$I(S'; A, C \mid S) = H(S'|S) - H(S'|S, A, C)$$

where $H(S'|S)$ is the conditional entropy of the distribution of outcomes S' given the present state S, and $H(S'|S, A, C)$ is the conditional entropy of the outcome given the present state, action, and resource investment.

To compare controllability, we consider two environments with the same maximum control:

• Inelastic environment: If the correct vehicle is chosen, there is a 100% chance of reaching the goal state with 1, 2, or 3 tickets. Thus, out of 7 possible action-resource investment combinations, three deterministically lead to the goal state (≥1 tickets and correct vehicle choice), three never lead to it (≥1 tickets and wrong vehicle choice), and one (0 tickets) leads to it 20% of the time (since walking leads to the treasure on 20% of trials).

• Elastic Environment: If the correct vehicle is chosen, the probability of boarding it is 0% with 1 ticket, 50% with 2 tickets, and 100% with 3 tickets. Thus, out of 7 possible actionresource investment combinations, one deterministically leads to the goal state (3 tickets and correct vehicle choice), one never leads to it (3 tickets and wrong vehicle choice), one leads to it 60% of the time (2 tickets and correct vehicle choice: 50% boarding + 50% × 20% when failing to board), one leads to it 10% of time (2 ticket and wrong vehicle choice), and three lead to it 20% of time (0-1 tickets).

Here we assume a uniform prior over actions, which renders the information-theoretic definition of controllability equal to another definition termed 'instrumental divergence'3,4. We note that changing the uniform prior assumption would change the results for the two environments, but that would not change the general conclusion that there can be environments that are more controllable yet less elastic.

Step 1: Calculating $H(S'|S)$

For the inelastic environment:

P(goal) = (3 × 100% + 3 × 0% + 1 × 20%)/7 = .46, P(non-goal) = .54 $H(S'|S) = -[.46 \times \log_2(.46) + .54 \times \log_2(.54)] = 1$ bit

For the elastic environment:

P(goal) = (1 × 100% + 1 × 0% + 1 × 60% + 1 × 10% + 3 × 20%)/7 = .33, P(non-goal) = .67 $H(S'|S) = -[.33 \times \log_2(.33) + .67 \times \log_2(.67)] = .91$ bits

Step 2: Calculating $H(S'|S, A, C)$

Inelastic environment: Six action-resource investment combinations have deterministic outcomes entailing zero entropy, whereas investing 0 tickets has a probabilistic outcome (20%). The entropy for 0 tickets is: $H(S'|C = 0) = -[.2 \times \log_2(.2) + 0.8 \times \log_2(.8)] = .72$ bits. Since this actionresource investment combination is chosen with probability 1/7, the total conditional entropy is approximately .10 bits

Elastic environment: 2 actions have deterministic outcomes (3 tickets with correct/wrong vehicle), whereas the other 5 actions have probabilistic outcomes:

2 tickets and correct vehicle (60% success):

$H(S'|A = \text{correct}, C = 2) = -[.6 \times \log_2(.6) + .4 \times \log_2(.4)] = .97$ bits 2 tickets and wrong vehicle (10% success):

$H(S'|A = \text{wrong}, C = 2) = -[.1 \times {}_2(.1) + .9 \times {}_2(.9)] = .47$ bits 0-1 tickets (20% success):

$H(S'|C = \text{0-1}) = -[.2 \times {}_2(.2) + .8 \times {}_2 .8)] = .72$ bits

Thus the total conditional entropy of the elastic environment is: $H(S'|S, A, C) = (1/7) \times .97 + (1/7) \times .47 + (3/7) \times .72 = .52$ bits

Step 3: Calculating $I(S' \mid A, S)$

Inelastic environment: $I(S'; A, C \mid S) = H(S'|S) - H(S'|S, A, C) = 1 - 0.1 = .9$ bits

Elastic environment: $I(S'; A, C \mid S) = H(S'|S) - H(S'|S, A, C) = .91 - .52 = .39$ bits

Thus, the inelastic environment offers higher information-theoretic controllability (.9 bits) compared to the elastic environment (.39 bits).

Of note, even if each combination of cost and goal reaching is defined as a distinct outcome, then information-theoretic controllability is higher for the inelastic (2.81 bits) than for the elastic (2.30 bits) environment.

In sum, for both definitions of controllability, we see that environments can be more elastic yet less controllable. We will amend the manuscript to clarify this distinction between controllability and its elasticity.

> ***Reviewer 3:***
>
> *A bias in how people infer the amount of control they have over their environment is widely believed to be a key component of several mental illnesses including depression, anxiety, and addiction. Accordingly, this bias has been a major focus in computational models of those disorders. However, all of these models treat control as a unidimensional property, roughly, how strongly outcomes depend on action. This paper proposes---correctly, I think---that the intuitive notion of "control" captures multiple dimensions in the relationship between action and outcome is multi-dimensional. In particular, the authors propose that the degree to which outcome depends on how much *effort* we exert, calling this dimension the "elasticity of control". They additionally propose that this dimension (rather than the more holistic notion of controllability) may be specifically impaired in certain types of psychopathology. This idea thus has the potential to change how we think about mental disorders in a substantial way, and could even help us better understand how healthy people navigate challenging decision-making problems.*
>
> *Unfortunately, my view is that neither the theoretical nor empirical aspects of the paper really deliver on that promise. In particular, most (perhaps all) of the interesting claims in the paper have weak empirical support.*

We appreciate the Reviewer's thoughtful engagement with our research and recognition of the potential significance of distinguishing between different dimensions of control in understanding psychopathology. We believe that all the Reviewer’s comments can be addressed with clarifications or additional analyses, as detailed below.

> *Starting with theory, the elasticity idea does not truly "extend" the standard control model in the way the authors suggest. The reason is that effort is simply one dimension of action. Thus, the proposed model ultimately grounds out in how strongly our outcomes depend on our actions (as in the standard model). Contrary to the authors' claims, the elasticity of control is still a fixed property of the environment. Consistent with this, the computational model proposed here is a learning model of this fixed environmental property. The idea is still valuable, however, because it identifies a key dimension of action (namely, effort) that is particularly relevant to the notion of perceived control. Expressing the elasticity idea in this way might support a more general theoretical formulation of the idea that could be applied in other contexts. See Huys &*

> *Dayan (2009), Zorowitz, Momennejad, & Daw (2018), and Gagne & Dayan (2022) for examples of generalizable formulations of perceived control.*

We thank the Reviewer for the suggestion that we formalize our concept of elasticity to resource investment, which we agree is a dimension of action. We first note that we have not argued against the claim that elasticity is a fixed property of the environment. We surmise the Reviewer might have misread our statement that “controllability is not a fixed property of the environment”. The latter statement is motivated by the observation that controllability is often higher for agents that can invest more resources (e.g., a richer person can buy more things). We will clarify this in our revision of the manuscript.

To formalize elasticity, we build on Huys & Dayan’s definition of controllability(1) as the fraction of reward that is controllably achievable, χ (though using information-theoretic definitions(2,3) would work as well). To the extent that this fraction depends on the amount of resources the agent is able and willing to invest (max C), this formulation can be probabilistically computed without information about the particular agent involved, specifically, by assuming a certain distribution of agents with different amounts of available resources. This would result in a probability distribution over χ. Elasticity can thus be defined as the amount of information obtained about controllability due to knowing the amount of resources available to the agent: $\mathrm{I}(\chi; \max C)$. We will add this formal definition to the manuscript.

> *Turning to experiment, the authors make two key claims: (1) people infer the elasticity of control, and (2) individual differences in how people make this inference are importantly related to psychopathology. Starting with claim 1, there are three sub-claims here; implicitly, the authors make all three. (1A) People's behavior is sensitive to differences in elasticity, (1B) people actually represent/track something like elasticity, and (1C) people do so naturally as they go about their daily lives. The results clearly support 1A. However, 1B and 1C are not supported. Starting with 1B, the experiment cannot support the claim that people represent or track elasticity because the effort is the only dimension over which participants can engage in any meaningful decision-making (the other dimension, selecting which destination to visit, simply amounts to selecting the location where you were just told the treasure lies). Thus, any adaptive behavior will necessarily come out in a sensitivity to how outcomes depend on effort. More concretely, any model that captures the fact that you are more likely to succeed in two attempts than one will produce the observed behavior. The null models do not make this basic assumption and thus do not provide a useful comparison.*

We appreciate the reviewer's critical analysis of our claims regarding elasticity inference, which as detailed below, has led to an important new analysis that strengthens the study’s conclusions. However, we respectfully disagree with two of the Reviewer’s arguments. First, resource investment was not the only meaningful decision dimension in our task, since participant also needed to choose the correct vehicle to get to the right destination. That this was not trivial is evidenced by our exclusion of over 8% of participants who made incorrect vehicle choices more than 10% of the time. Included participants also occasionally erred in this choice (mean error rate = 3%, range [0-10%]).

Second, the experimental task cannot be solved well by a model that simply tracks how outcomes depend on effort because 20% of the time participants reached the treasure despite failing to board their vehicle of choice. In such cases, reward outcomes and control were decoupled. Participants could identify when this was the case by observing the starting location, which was revealed together with the outcome (since depending on the starting location, the treasure location was automatically reached by walking). To determine whether participants distinguished between control-related and non-control-related reward, we have now fitted a variant of our model to the data that allows learning from each of these kinds of outcomes by means of a different free parameter. The results show that participants learned considerably more from control-related outcomes. They were thus not merely tracking outcomes, but specifically inferred when outcomes can be attributed to control. We will include this new analysis in the revised manuscript.

Controllability inference by itself, however, still does not suffice to explain the observed behavior. This is shown by our 'controllability' model, which learns to invest more resources to improve control, yet still fails to capture key features of participants' behavior, as detailed in the manuscript. This means that explaining participants' behavior requires a model that not only infers controllability—beyond merely outcome probability—but also assumes a priori that increased effort could enhance control. Building these a priori assumption into the model amounts to embedding within it an understanding of elasticity – the idea that control over the environment may be increased by greater resource investment.

That being said, we acknowledge the value in considering alternative computational formulations of adaptation to elasticity. Thus, in our revision of the manuscript, we will add a discussion concerning possible alternative models.

> *For 1C, the claim that people infer elasticity outside of the experimental task cannot be supported because the authors explicitly tell people about the two notions of control as part of the training phase: "To reinforce participants' understanding of how elasticity and controllability were manifested in each planet, [participants] were informed of the planet type they had visited after every 15 trips." (line 384).*

We thank the reviewer for highlighting this point. We agree that our experimental design does not test whether people infer elasticity spontaneously. Our research question was whether people can distinguish between elastic and inelastic controllability. The results strongly support that they can, and this does have potential implications for behavior outside of the experimental task. Specifically, to the extent that people are aware that in some contexts additional resource investment improve control, whereas in other contexts it does not, then our results indicate that they would be able to distinguish between these two kinds of contexts through trial-and-error learning. That said, we agree that investigating whether and how people spontaneously infer elasticity is an interesting direction for future work. We will clarify the scope of the present conclusions in the revised manuscript.

> *Finally, I turn to claim 2, that individual differences in how people infer elasticity are importantly related to psychopathology. There is much to say about the decision to treat psychopathology as a unidimensional construct. However, I will keep it concrete and simply note that CCA (by design) obscures the relationship between any two variables. Thus, as suggestive as Figure 6B is, we cannot conclude that there is a strong relationship between Sense of Agency and the elasticity bias---this result is consistent with any possible relationship (even a negative one). The fact that the direct relationship between these two variables is not shown or reported leads me to infer that they do not have a significant or strong relationship in the data.*

We agree that CCA is not designed to reveal the relationship between any two variables. However, the advantage of this analysis is that it pulls together information from multiple variables. Doing so does not treat psychopathology as unidimensional. Rather, it seeks a particular dimension that most strongly correlates with different aspects of task performance. This is especially useful for multidimensional psychopathology data because such data are often dominated by strong correlations between dimensions, whereas the research seeks to explain the distinctions between the dimensions. Similar considerations hold for the multidimensional task parameters, which although less correlated, may still jointly predict the relevant psychopathological profile better than each parameter does in isolation. Thus, the CCA enabled us to identify a general relationship between task performance and psychopathology that accounts for different symptom measures and aspects of controllability inference.

Using CCA can thus reveal relationships that do not readily show up in two-variable analyses. Indeed, the direct correlation between Sense of Agency (SOA) and elasticity bias was not significant – a result that, for completeness, we will now report in the supplementary materials along with all other direct correlations. We note, however, that the CCA analysis was preregistered and its results were replicated. Furthermore, an auxiliary analysis specifically confirmed the contributions of both elasticity bias (Figure 6D, bottom plot) and, although not reported in the original paper, of the Sense of Agency score (SOA; p=.03 permutation test) to the observed canonical correlation. Participants scoring higher on the psychopathology profile also overinvested resources in inelastic environments but did not futilely invest in uncontrollable environments (Figure 6A), providing external validation to the conclusion that the CCA captured meaningful variance specific to elasticity inference. The results thus enable us to safely conclude that differences in elasticity inferences are significantly associated with a profile of controlrelated psychopathology to which SOA contributed significantly.

Finally, whereas interpretation of individual CCA loadings that were not specifically tested remains speculative, we note that the pattern of loadings largely replicated across the initial and replication studies (see Figure 6B), and aligns with prior findings. For instance, the positive loadings of SOA and OCD match prior suggestions that a lower sense of control leads to greater compensatory effort(7), whereas the negative loading for depression scores matches prior work showing reduced resource investment in depression(5-6).

We will revise the text to better clarify the advantageous and disadvantageous of our analytical approach, and the conclusions that can and cannot be drawn from it.

> *There is also a feature of the task that limits our ability to draw strong conclusions about individual differences in elasticity inference. As the authors clearly acknowledge, the task was designed "to be especially sensitive to overestimation of elasticity" (line 287). A straightforward consequence of this is that the resulting *empirical* estimate of estimation bias (i.e., the gamma_elasticity parameter) is itself biased. This immediately undermines any claim that references the directionality of the elasticity bias (e.g. in the abstract). Concretely, an undirected deficit such as slower learning of elasticity would appear as a directed overestimation bias. When we further consider that elasticity inference is the only meaningful learning/decisionmaking problem in the task (argued above), the situation becomes much worse. Many general deficits in learning or decision-making would be captured by the elasticity bias parameter. Thus, a conservative interpretation of the results is simply that psychopathology is associated with impaired learning and decision-making.*

We apologize for our imprecise statement that the task was 'especially sensitive to overestimation of elasticity', which justifiably led to Reviewer's concern that slower elasticity learning can be mistaken for elasticity bias. To make sure this was not the case, we made use of the fact that our computational model explicitly separates bias direction (λ) from the rate of learning $(\frac{1}{\epsilon})$ through two distinct parameters, which initialize the prior concentration and mean of the model's initial beliefs concerning elasticity (see Methods pg. 22). The higher the concentration of the initial beliefs (ϵ), the slower the learning. Parameter recovery tests confirmed that our task enables acceptable recovery of both the bias $\lambda_{elasticity}$ (r=.81) and the concentration $\epsilon_{elasticity}$ (r=.59) parameters. And importantly, the level of confusion between the parameters was low (confusion of 0.15 for $\epsilon_{elasticity} \rightarrow \lambda_{elasticity}$ and 0.04 for $\lambda_{elasticity} \rightarrow \epsilon_{elasticity}$). This result confirms that our task enables dissociating elasticity biases from the rate of elasticity learning.

Moreover, to validate that the minimal level of confusion existing between bias and the rate of learning did not drive our psychopathology results, we re-ran the CCA while separating concentration from bias parameters. The results (Author response image 1) demonstrate that differences in learning rate (ϵ) had virtually no contribution to our CCA results, whereas the contribution of the pure bias (λ) was preserved.

We will incorporate these clarifications and additional analysis in our revised manuscript.

Author response image 1.

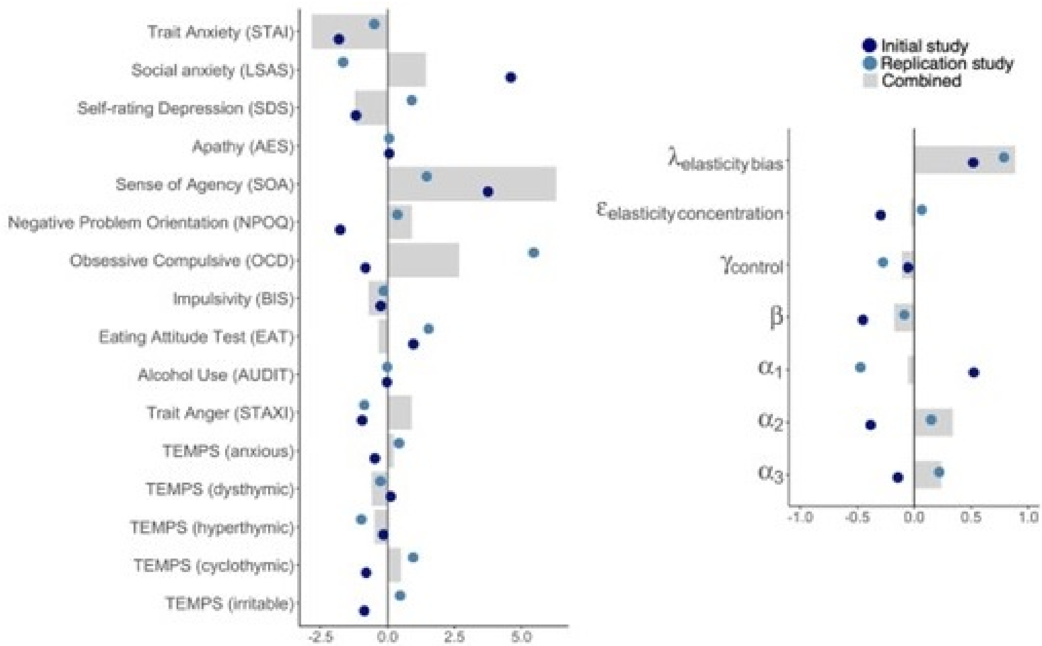

> *Showing that a model parameter correlates with the data it was fit to does not provide any new information, and cannot support claims like "a prior assumption that control is likely available was reflected in a futile investment of resources in uncontrollable environments." To make that claim, one must collect independent measures of the assumption and the investment.*

We apologize if this and related statements seemed to be describing independent findings. They were merely meant to describe the relationship between model parameters and modelindependent measures of task performance. It is inaccurate, though, to say that they provide no new information, since results could have been otherwise. For instance, instead of a higher controllability bias primarily associating with futile investment of resources in uncontrollable environments, it could have been primarily associated with more proper investment of resources in high-controllability environments. Additionally, we believe these analyses are of value to readers who seek to understand the role of different parameters in the model. In our planned revision, we will clarify that the relevant analyses are merely descriptive.

> *Did participants always make two attempts when purchasing tickets? This seems to violate the intuitive model, in which you would sometimes succeed on the first jump. If so, why was this choice made? Relatedly, it is not clear to me after a close reading how the outcome of each trial was actually determined.*

We thank the reviewer for highlighting the need to clarify these aspects of the task in the revised manuscript.

When participants purchased two extra tickets, they attempted both jumps, and were never informed about whether either of them succeeded. Instead, after choosing a vehicle and attempting both jumps, participants were notified where they arrived at. This outcome was determined based on the cumulative probability of either of the two jumps succeeding. Success meant that participants arrived at where their chosen vehicle goes, whereas failure meant they walked to the nearest location (as determined by where they started from).

Though it is unintuitive to attempt a second jump before seeing whether the first succeed, this design choice ensured two key objectives. First, that participants would consistently need to invest not only more money but also more effort and time in planets with high elastic controllability. Second, that the task could potentially generalize to the many real-world situations where the amount of invested effort has to be determined prior to seeing any outcome, for instance, preparing for an exam or a job interview.

> *It should be noted that the model is heuristically defined and does not reflect Bayesian updating. In particular, it overestimates control by not using losses with less than 3 tickets (intuitively, the inference here depends on your beliefs about elasticity). I wonder if the forced three-ticket trials in the task might be historically related to this modeling choice.*

We apologize for not making this clear, but in fact losing with less than 3 tickets does reduce the model's estimate of available control. It does so by increasing the elasticity estimates

($a_{elastic\geq1}$, $a_{elastic2}$ parameters), signifying that more tickets are needed to obtain the maximum available level of control, thereby reducing the average controllability estimate across ticket investment options.

It would be interesting to further develop the model such that losing with less than 3 tickets would also impact inferences concerning the maximum available control, depending on present beliefs concerning elasticity, but the forced three-ticket purchases already expose participants to the maximum available control, and thus, the present data may not be best suited to test such a model. These trials were implemented to minimize individual differences concerning inferences of maximum available control, thereby focusing differences on elasticity inferences. We will discuss the Reviewer's suggestion for a potentially more accurate model in the revised manuscript.

References

(1) Huys, Q. J. M., & Dayan, P. (2009). A Bayesian formulation of behavioral control. Cognition, 113(3), 314– 328.

(2) Ligneul, R. (2021). Prediction or causation? Towards a redefinition of task controllability. Trends in Cognitive Sciences, 25(6), 431–433.

(3) Mistry, P., & Liljeholm, M. (2016). Instrumental divergence and the value of control. Scientific Reports, 6, 36295.

(4) Lin, J. (1991). Divergence measures based on the Shannon entropy. IEEE Transactions on Information Theory, 37(1), 145–151

(5) Cohen RM, Weingartner H, Smallberg SA, Pickar D, Murphy DL. Effort and cognition in depression. Arch Gen Psychiatry. 1982 May;39(5):593-7. doi: 10.1001/archpsyc.1982.04290050061012. PMID: 7092490.

(6) Bi R, Dong W, Zheng Z, Li S, Zhang D. Altered motivation of effortful decision-making for self and others in subthreshold depression. Depress Anxiety. 2022 Aug;39(8-9):633-645. doi: 10.1002/da.23267. Epub 2022 Jun 3. PMID: 35657301; PMCID: PMC9543190.

(7) Tapal, A., Oren, E., Dar, R., & Eitam, B. (2017). The Sense of Agency Scale: A measure of consciously perceived control over one's mind, body, and the immediate environment. Frontiers in Psychology, 8, 1552

https://doi.org/10.7554/eLife.105507.1.sa0