Attention Allocation under Scarcity*

Patrick Sewell

November 1, 2026

Click here for the latest version.

Abstract

Scarcity is central to economics because limited resources force trade-offs. But can the mere *feeling* of scarcity affect economic behavior beyond these constraints? I study this question using theory, experimental evidence, and administrative data. I develop a model that defines *perceived scarcity* as resources falling below subjective needs and shows how it leads individuals to optimally reallocate limited attention toward the scarce domain and away from competing tasks. The model predicts when these shifts improve, worsen, or do not affect performance, depending on how attention translates into outcomes in each task. I test the model in a preregistered experiment in which 501 participants allocate attention between a consumption and a production task. To induce differences in perceived scarcity, I randomly assign participants to high or low consumption budgets before giving them the same intermediate budget. Despite facing identical constraints, participants accustomed to higher budgets are substantially more likely to feel scarce (77% vs. 45%). They devote 10 percentage points more attention to consumption (measured through mouse-tracking) and search more extensively, yet do not make significantly better consumption decisions. Meanwhile, their performance in the production task declines, and their total earnings are 15% lower. Exploiting a reform in Hungarian unemployment benefits, I examine the model's implications for job search. I find a pattern consistent with the model: despite facing *identical* current benefits, workers who experienced a larger benefit drop have a 56% higher job-finding rate but 9% lower reemployment wages.

*Universitat Pompeu Fabra and BSE; patrick.sewell@upf.edu. I am very grateful to my advisor Jose Apesteguia for his support and guidance. I also thank Chiara Aina, Larbi Alaoui, Amil Camilo, Colin Camerer, Brenden Eum, Charlotte Hervy, Shachar Kariv, Ian Krajbich, Gaël Le Mens, Ryan Oprea, Rosemarie Nagel, Gautam Rao, Pedro Rey Biel, Andrea Salvanti, Charlie Sprenger, Marta Serra Garcia, Frank Schilbach, Agnieszka Tymula and Michael Woodford, as well as seminar audiences at Caltech, Krajbich Lab UCLA, UPF, University of Sydney, and participants at the Berlin PhD Conference, BSPA Harvard Kennedy School, Behavioral Spring School UCSD, EEA-ESEM Dublin, ESA Europe, FUR Alicante, Theory-Based Experiments Caltech, Choice & Welfare Queen Mary, FUR Queensland, and Cognitive Foundations of Decision Making Ghent for helpful discussions and suggestions. Funding from the UPF's REACT grant is gratefully acknowledged. The Hungarian administrative data were provided by the ELTE KRTK Databank. The study received UPF ethics approval (CIREP approval #471) and was preregistered in the AEA registry (#AEARCTR-0018956).

1 Introduction

Scarcity in the broad economic sense is ubiquitous: resources are finite and trade-offs are unavoidable. Yet some constraints are experienced as especially tight, making them cognitively salient in a way that ordinary resource limitations are not. A growing literature in economics shows that such experiences of scarcity can affect many economic outcomes including workplace productivity, school performance and consumption choice (Duquennois 2022; Fehr et al. 2022; Kaur et al. 2025; Goldin and Homonoff 2013). Importantly, these experiences need not map one-to-one into objective constraints: the same level of resources may feel more or less scarce depending on individuals' subjective needs (Bonke and Browning 2009; Zubrickas 2026). Perceptions are malleable, so the same individual may feel scarce at one point but not another, even when facing identical resources.

In this paper, I formalize this distinction by defining *perceived scarcity* as the feeling that available resources are insufficient relative to one's subjective needs, and study its implications for economic behavior using theory, experimental evidence, and administrative data. I show that, even when objective resources are fixed, the mere *perception* of scarcity reallocates attention toward tasks in the scarce domain and away from competing tasks. Since attention is an important input to decision-making and productivity, this creates a channel through which economic outcomes depend not only on available resources but also on how scarce those resources feel.

The theoretical framework builds on seminal work in psychology and economics, which argues that scarcity captures individuals' attention, leaving less bandwidth for important tasks (Mullainathan and Shafir 2013; Shah et al. 2012), and provides evidence that it can impede cognitive function (Mani et al. 2013). I formalize this intuition with a model in which agents allocate limited attention across competing tasks. Crucially, I show that the negative performance effects of scarcity need not result from individuals becoming less capable; they can arise from optimal behavior under cognitive constraints.

Distinguishing perceived scarcity from objective constraints is also important for identification, as comparing individuals facing different constraints makes it difficult to isolate the psychological effects of scarcity (tho Pesch et al. 2026). For example, a low-income person may be more likely to take out a costly payday loan than a high-income person simply because they face tighter financial constraints, not necessarily because feeling scarce changes their decision-making. To address this challenge, I develop a novel experimental design that induces perceived scarcity while holding objective resources fixed. The design allows me to directly measure attention allocation and jointly test for the potential positive and negative

effects of perceived scarcity across competing tasks. In addition, I complement the experiment with suggestive field evidence from a Hungarian unemployment insurance reform, in which workers face identical benefits after different benefit histories.

Despite mounting evidence on the behavioral effects of scarcity, recent reviews highlight a lack of clear definitions and conflicting empirical findings (De Bruijn and Antonides 2022; Haushofer and Salicath 2023). Moreover, as Park et al. (2025) note, "scarcity theory fails to make clear a priori predictions" about when its effects should be positive or negative. On one hand, there is evidence that scarcity impairs cognitive performance and outcomes at school and work (Duquennois 2022; Kaur et al. 2025; Mani et al. 2013; Sergeyev et al. 2025). On the other, some studies find that individuals facing scarcity make more careful, and sometimes better, decisions when scarce resources are at stake (Fehr et al. 2022; Goldin and Homonoff 2013; Shah et al. 2015). This paper shows that these seemingly opposing effects can coexist and arise from a common mechanism of optimal attention allocation. My model predicts that perceived scarcity draws attention toward tasks in the scarce domain and away from competing tasks, while the resulting effects on performance depend on how attention translates into outcomes in each task.

In my model, an agent behaves optimally subject to an attention constraint, allocating attention across competing tasks according to their expected returns and the complexity of each task. The agent has subjective thresholds of need for resources such as time, money, information, or calories, which may be shaped by past experience or social comparison. Perceived scarcity arises when available resources fall below these thresholds. To illustrate the mechanism, I focus on an agent who allocates attention between two tasks—a consumption task and a production task—with independent payoffs that enter utility additively. The agent intuitively expects more attention to improve performance in the production task, while in the consumption task attention expands the consideration set, allowing for better choices. When the available consumption budget falls short of the agent's perceived needs, the marginal value of devoting attention to consumption increases. The agent therefore optimally reallocates attention toward consumption and away from the production task.

Importantly, the result on attention reallocation depends on agents *expecting* that more attention translates into better performance; however, the actual mapping from attention to performance may be flat or even non-monotonic depending on the task (Camerer and Hogarth 1999; Kahneman et al. 1968).[1] For instance, in the consumption task, a larger consideration set also makes identifying the best bundle more computationally demanding (Murawski and Bossaerts 2016), so additional attention need not improve realized deci-

[1]See also Alaoui and Penta (2022), who stress the importance of separating reasoning from performance.

sion quality. Crucially, it is this actual mapping that determines whether the reallocation translates into the corresponding performance effects. When performance increases with attention, the predicted reallocation leads to better consumption decisions and lower productivity. If this relationship does not hold for a given task, attention may still shift without the predicted performance effect in that task. This distinction highlights the importance of directly measuring attention. If a researcher only measures task performance and observes a null effect, they may conclude that the psychological channel is inactive, when in reality it is shifting attention, which can be welfare-relevant in itself and may also generate spillovers on other, unmeasured economic outcomes.[2]

Guided by the model, I develop a dual-task experiment in which participants perform a consumption task and a production task. In each round, participants can freely allocate their attention between the two tasks within a fixed time limit. The *production task* is a modified version of the standard slider task, commonly used as a pure-effort task, adapted in this experiment to also require cognitive processing. Participants are shown a series of sliders, but rather than being provided the target value, they need to sum three integers and then drag the slider to match the computed value. In the *consumption task*, participants receive a budget and are incentivized to select the highest-valued affordable bundle from a set of 18 items, each with a price and value. This allows me to objectively measure decision quality relative to the optimal solution. All consumption items are initially covered, and participants have to hover the mouse over each one to reveal it. Once revealed, the item remains visible. The number of revealed items provides an empirical proxy for the size of each participant's consideration set. Attention is directly measured via mouse-tracking, which records the time spent on each task and switching behavior.

Perceived scarcity is induced by randomly assigning participants to a *High-Budget Adaptation* (HBA) or *Low-Budget Adaptation* (LBA) condition. During three adaptation rounds, HBA participants receive high budgets in the consumption task, habituating them to higher levels of consumption, while LBA participants receive low budgets. In the fourth (test) round, both groups receive the same intermediate budget, such that HBA experiences a budget drop while LBA experiences an increase. Thus, the groups face identical resource constraints and incentives in the test round, but enter it with different budget histories.

In a preregistered online experiment with 501 participants recruited through Prolific, I find that these different budget histories lead to sharply different perceptions of scarcity: 77% of HBA participants report feeling scarce, compared with 45% of LBA participants

[2]See, for example, Quispe-Torreblanca et al. (2026), Bolte and Raymond (2026), and Loewenstein and Wojtowicz (2025).

($p < 0.001$). Consistent with the model, HBA participants allocate 10 percentage points more of their attention to the consumption task than the LBA group (69.7% vs. 59.7%; $p < 0.001$). This shift is accompanied by more extensive search in the consumption task: HBA participants form larger consideration sets ($\Delta = 0.87$ items; $p = 0.03$) and revise their selected bundles more frequently ($\Delta = 0.66$ revisions; $p < 0.001$).

Yet, despite allocating more attention and searching more extensively, HBA participants do not make significantly better consumption decisions, achieving a similar share of the optimal bundle's value (0.704 vs. 0.722; $p = 0.37$). Additional analysis suggests why: decision quality rises with attention up to about 30 seconds but is essentially flat thereafter, and most participant-rounds lie in this flat region. By contrast, performance in the production task increases monotonically with attention. HBA participants, who devote less attention to production, complete 23.7% fewer correct sliders ($p < 0.001$). The overall consequence is a substantial reduction in earnings: HBA participants earn approximately 15% less across both tasks (£2.14 vs. £2.53; $p < 0.001$). Thus, in this setting, perceived scarcity reallocates attention toward consumption without improving consumption decisions, while reducing productivity and overall earnings.

I also examine whether a similar trade-off emerges outside the laboratory, using administrative data from an unemployment insurance reform in Hungary (DellaVigna et al. 2017; Lindner and Reizer 2020). The reform provides a real-world analogue to my experiment and a natural setting to examine the model's predictions. In November 2005, Hungary replaced its unemployment insurance schedule with a two-step schedule that front-loaded benefits while leaving total benefits roughly unchanged. As a result, workers entering unemployment before and after the reform had different benefit histories but eventually faced identical benefits after day 270, which is the period I focus on in my analysis. Importantly, those entering *before* the reform experienced a larger drop from the benefit level they were used to, much like HBA participants in my experiment.

In an extension of my model, job search involves two competing uses of attention: sending applications, which increases the chances of finding a job quickly, and evaluating vacancies, which helps workers find jobs that better match their skills. Perceived scarcity increases the value of finding a job quickly, shifting attention toward applications and away from evaluating vacancies.

I find patterns consistent with this prediction. In the six weeks after day 270, workers who entered unemployment before the reform exit unemployment faster, with a 56% higher exit rate than those entering after the reform ($p < 0.001$), reproducing the findings of DellaVigna et al. (2017). Yet these faster exits come with lower reemployment wages:

relative to their previous earnings, their new wages are approximately 9% lower than those of post-reform workers ($p < 0.01$). This wage gap is absent before day 270 and disappears after day 315, consistent with subjective needs gradually adjusting to the new benefit level. While DellaVigna et al. (2017) focus on search effort, my model provides a complementary perspective by also considering how search affects match quality, predicting effects on both the speed of reemployment and the wages workers obtain. Although attention is not directly observed in the administrative data, the results provide suggestive evidence consistent with the model's predicted dual effect: faster reemployment but lower wages while facing identical unemployment benefits.[3]

This paper contributes to three strands of literature. First, it contributes to the scarcity literature by providing a formal framework that defines perceived scarcity and distinguishes it from objective resource constraints. The framework helps explain why scarcity can have positive, negative, or null effects on task performance, depending on the task domain and how attention translates into outcomes. The experiment introduces a methodology to isolate the behavioral effects of perceived scarcity by manipulating prior resource experiences while holding current constraints fixed. Using mouse-tracking to directly measure attention, I provide evidence that perceived scarcity shifts attention toward the scarce domain and away from competing tasks.

Second, it contributes to the economics of attention by modeling how agents allocate limited attention across competing tasks according to expected returns and task complexity. While models of rational inattention emphasize constraints on information processing (Sims 2003), my framework focuses on the opportunity cost of allocating attention across competing tasks: more attention to one task leaves less available for others. This connects to a rich literature in psychology on the opportunity costs of cognitive effort (Kurzban et al. 2013), a perspective I apply to attention allocation across economic tasks.

Third, it contributes to the job-search literature by providing a framework in which different aspects of job search compete for limited attention, offering a complementary explanation for why unemployment insurance reforms may lead to faster reemployment but lower wages. This relates to evidence on how unemployment benefits affect reemployment wages and job quality (Schmieder et al. 2016; Nekoei and Weber 2017). I find suggestive evidence consistent with this mechanism using administrative data from Hungary.

These findings have implications for policy design. Many policies target individuals based on objective constraints such as income, liquidity, or wealth. My results suggest that

[3]Some patterns in the data fit the attention mechanism particularly well: pre-reform workers do not take more of the lowest-paid jobs, but they find fewer of the well-paid ones. This is harder to explain by workers simply lowering their reservation wage (Section 6).

subjective needs also matter: individuals facing identical resources may allocate attention, make decisions, and perform differently depending on how scarce those resources feel. Such perceptions can be measured, and several large surveys already ask questions about subjective financial strain.[4] Collecting these measures alongside objective constraints could help identify individuals who perceive scarcity but might otherwise be overlooked by standard measures. The design of benefit schedules may also matter beyond their total generosity. When benefits drop sharply or temporary support expires, individuals whose needs have adapted to higher support may perceive scarcity, generating attentional effects beyond the direct effects of reduced resources. Because needs appear to adjust over time, phasing out support gradually could help reduce these effects. More broadly, policies may affect behavior beyond the domain they target, as changes in perceived scarcity can shift attention away from other important tasks.

2 Related Literature

Evidence on Scarcity

A large body of work links scarcity to three main behavioral effects. First, scarcity can increase trade-off thinking (Spiller 2011; Shah et al. 2015), making individuals more likely to consider the opportunity costs of their choices and potentially leading to more consistent decisions. Second, scarcity can lead to tunneling (Shah et al. 2012; Shah et al. 2019; Zhao and Tomm 2017; Lichand and Mani 2020), an intensified focus on scarcity-related tasks and information, accompanied by neglect of other tasks and considerations. Third, scarcity may impose a bandwidth tax (Duquennois 2022; Mani et al. 2013; Ong et al. 2019; Lichand and Mani 2020), reducing the cognitive resources available for other tasks and potentially undermining productivity and decision quality. These psychological effects are of particular interest for economists because they provide mechanisms through which scarcity can affect economic outcomes beyond its direct effect on the feasible set.

Evidence from field studies and natural experiments shows that scarcity can have both positive and negative effects on economic behavior. Fehr et al. (2022) find that financially constrained households are less susceptible to the endowment effect and make more consistent

[4]The EU Statistics on Income and Living Conditions (Eurostat 2011) asks households how easily they are able to make ends meet and has also asked for the lowest monthly income they would need to do so, which closely corresponds to the subjective threshold in my framework (Goedhart et al. 1977). The European Social Survey (2002) asks whether respondents are living comfortably, coping, or finding it difficult on their present income, and the Financial Well-Being Scale of the U.S. Consumer Financial Protection Bureau (2015) asks whether respondents are "just getting by financially."

decisions, without finding evidence of impaired cognitive performance. In contrast, Kaur et al. (2025) show that financial concerns reduce worker productivity. By paying some workers earlier, they reduce financial concerns and increase productivity by 7.1%. Similarly, Duquennois (2022) finds that disadvantaged students perform worse on financially salient math questions, while Lichand and Mani (2020) show that rainfall uncertainty among farmers increases attention to scarce-resource concerns while impairing cognitive performance.

Despite this evidence, identifying the psychological effects of scarcity remains challenging. Much of the empirical literature studies settings where scarcity is intertwined with objective resource constraints, making it difficult to separate their direct and psychological effects. Laboratory experiments offer greater control, but some findings have been sensitive to experimental design. For example, a result on cognitive fatigue from Shah et al. (2012) did not replicate in a large replication exercise (Camerer et al. 2018), although subsequent replications by the original authors supported several other findings (Shah et al. 2019). More recently, tho Pesch et al. (2026) argue that some of the behavioral differences previously attributed to tunneling can instead be explained by the structure of the experimental task.

A related issue concerns how scarcity is defined and measured. De Bruijn and Antonides (2022) highlight a mismatch between the subjective definition of scarcity and the objective measures often used in empirical studies. This paper addresses these issues by formally defining perceived scarcity in terms of subjective needs and deriving predictions for attention allocation and performance. The experiment then tests these predictions by inducing differences in perceived scarcity while holding current resource constraints and incentives fixed, and directly measuring how participants allocate their attention. The dual-task design also allows us to test for both potential benefits and costs of scarcity within the same experiment.

Attention Allocation and Performance

A large literature in economics studies how individuals allocate limited attention when making decisions. Rational inattention models costly information acquisition (Sims 2003), while sparse maximization and directed cognition study how individuals selectively process information or evaluate alternatives (Gabaix 2014; Gabaix and Laibson 2005). Related models study choices made from limited consideration sets (Masatlioglu et al. 2012; Manzini and Mariotti 2014; Cattaneo et al. 2020). While these approaches mainly focus on how attention is used within a decision, this paper studies how limited attention is allocated across competing tasks. The closest papers to my theoretical framework are Alonso et al. (2014) and Altmann et al. (2022), who model how limited cognitive resources are allocated across competing tasks. Like Altmann et al. (2022), I abstract from the asymmetric information

across brain regions considered by Alonso et al. (2014). My model differs by incorporating task complexity directly into the attention constraint and allowing attention to affect consideration sets and productive output, rather than only the probability of task success. Finally, I study how perceived scarcity changes attention allocation even when objective resources and incentives remain unchanged.

A related literature studies how cognitive effort translates into task performance. Although greater attention and effort are generally expected to improve performance, this may not always be the case. In a review of 74 experiments, Camerer and Hogarth (1999) emphasize that the effects of stronger incentives depend on how additional effort translates into performance. For example, Kahneman et al. (1968) find that higher incentives increase pupil dilation, a proxy for cognitive effort, without improving accuracy in a mental arithmetic task. Other studies find that stronger incentives can increase effort without corresponding improvements in performance, or even worsen it (Libby and Lipe 1992; Arkes et al. 1986; Hogarth et al. 1991). Relatedly, Alaoui and Penta (2022) develop a cost-benefit framework of reasoning and highlight the importance of distinguishing reasoning from task performance. While these studies primarily focus on monetary incentives and reasoning effort, my experiment studies how perceived scarcity reallocates attention across competing tasks and separately measures performance in each. This allows me to distinguish changes in attention from their effects on performance.

Job Search

The application to job search in Appendix B relates to models in which workers choose their search intensity, which affects the probability of finding a job (DellaVigna and Paserman 2005; DellaVigna et al. 2017; Nekoei and Weber 2017), and to models of directed and simultaneous search in which workers choose where to apply (Chade and L. Smith 2006; Menzio 2007; Kircher 2009). More recently, Wu (2024) models costly information acquisition that allows workers to better direct their applications toward high-paying jobs. In my model, search intensity corresponds to the attention devoted to submitting applications, which increases the probability of finding a job, while search direction depends on the attention devoted to evaluating job opportunities. These two activities compete for the same limited attention: more attention to evaluating opportunities increases the chances of finding a better match, but leaves less attention for applications. This provides a mechanism through which perceived scarcity can affect both the speed and quality of reemployment.

3 Theoretical Framework

3.1 Decision Problem(s)

The core of the model is an attention allocation problem across multiple tasks. The agent's goal is to maximize his overall performance on tasks given the attention constraints. I assume a dual-task problem to simplify analysis, where one task is a modified consumption choice problem and the other is a productive task. In the former, the agent chooses the utility maximizing bundle from a subset of all possible alternatives (consideration set). Increasing attention to this problem leads to bigger consideration sets and superior bundles. In the second task, the agent's goal is to maximize production, which is an increasing function of attention.[5]

3.2 Primitives

The outcome of the tasks described above are formalized with $x(a_x)$ and $y(a_y)$, where $x(a_x)$ is the chosen bundle of goods and $y(a_y)$ is the amount produced given the attention allocated pair (a_x, a_y). The agent is endowed with finite attention $A \in \mathbb{R}_+$ (Simon 1955), therefore he faces an attention constraint: $k_x a_x + k_y a_y \leq A$ where $k_j \in \mathbb{R}_{++}$ denotes the complexity[6] of task j. One can observe that the relationship between k_j and A is analogous to the price of a good and the budget constraint in a standard consumption problem.

Following the allocation problem, the agent performs two tasks. In the consumption problem, the agent is endowed with a budget $b \in \mathbb{R}_+$ and spends it on the bundle that maximizes his utility. The bundle is represented by a vector $x = (x_1, x_2, \ldots x_n) \in X \subset \mathbb{R}^n_+$, where x_i represents the quantity of good i. Denote the price of each good by $p_i \in \mathbb{R}_+$ and $p \in \mathbb{R}^n_+$, therefore the agent's budget set is given by $B(p, b) = \{x \in X | p \cdot x \leq b\}$. Individuals evaluate their consumption relative to a subjective level of need. Let $c(x)$ denote the aggregate level of consumption and η the minimum level of consumption that the agent perceives as sufficient.[7] Depending on the domain, this threshold may reflect physical perceived needs, such as caloric intake, or social norms, as illustrated by Adam

[5] The agent allocates attention according to how he perceives the returns to attention in each task, and the primitives below describe this perception. In practice, the actual mapping from attention to performance may be flat or even non-monotonic in a given task. I return to the implications of this distinction for performance in Corollary 1.

[6] Oprea (2026) provides an extensive discussion of task complexity, which I connect to the parameters k_x, k_y

[7] Notice that this reduces to a standard consumption problem if $\eta = 0$.

Smith's example of the linen shirt: although not necessary for survival, a respectable day-labourer would be ashamed to appear in public without one (A. Smith 1776). Crucially, η is subjective: individuals facing identical resources and prices may have different perceptions of what constitutes enough. Section 3.5 discusses how these perceptions may be formed by past consumption or social comparison.[8]

Definition 3.1 (Perceived Scarcity)**.** Scarcity is perceived whenever the level of consumption is below a subjective need,[9] formally: $c(x) < \eta$.

The key object in this framework is therefore not objective scarcity but perceived scarcity, namely whether consumption falls below the subjective need. The idea of need is captured in a kinked-utility function, where the marginal utility of consumption is higher when the level of consumption is below η than when the need is satisfied.[10]

Definition 3.2 (Kink Utility)**.** A utility function representing scarcity has the form:

$$v(c;\eta) = \begin{cases} v_1(c) & \text{if } c < \eta \\ v_2(c) + \Delta(\eta) & \text{if } c \geq \eta \end{cases} \tag{1}$$

where crucially, $\frac{dv_1}{dc} > \frac{dv_2}{dc} > 0$; and $\Delta(\eta) = v_1(\eta) - v_2(\eta)$ just ensures that there are no discontinuities in the utility function. Furthermore, let $v_1(0) = v_2(0) = 0$.

The level of consumption aggregates all goods in a standard CES form[11]: $c(x) = \left[\sum_i \alpha_i x_i^r\right]^{\frac{1}{r}}$, where $\alpha_i > 0$ captures how much good i contributes to the level of consumption and $r \in (-\infty, 1)$ reflects the degree of substitutability between goods. The utility of a bundle x is then $u(x;\eta) := v(c(x);\eta)$. Therefore, this utility function is (i) continuous, (ii) monotonic, and (iii) strictly quasiconcave in x, so that preferences are strictly convex (proofs in Appendix A).

[8] The consumption level and subjective need can be defined broadly or for a specific consumption domain. For instance, one could model separate consumption tasks for food, clothing, or transportation, each with its own budget and subjective need threshold, as when households keep separate mental budgets for different categories of spending (Thaler 1999; Heath and Soll 1996). Each domain would then be an additional task competing for attention, and the same logic would apply to each of them. Here, I adopt the broader interpretation for simplicity.

[9] This can be extended to other dimensions, such as time.

[10] Unlike the Stone–Geary utility function (Geary 1950; Stone 1954), where needs are objective subsistence quantities of each good below which utility is not defined, here the need is subjective, applies to the overall level of consumption, and can be unmet: consumption below η is feasible but carries a higher marginal utility.

[11] Other aggregating functions $c : \mathbb{R}^n_+ \to \mathbb{R}_+$ also work, as long as they are continuous, increasing in each good and strictly quasiconcave.

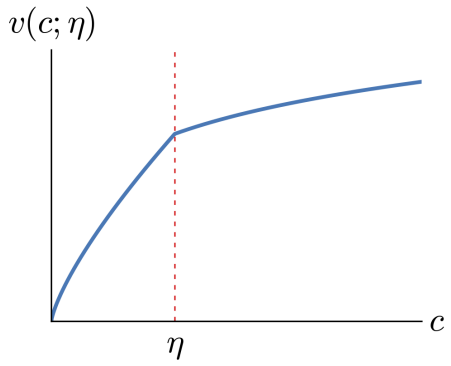


Figure 1: Example of kink utility with $\eta = 5$

In the second task, the agent's production output is a function of the attention allocated to it $y(a_y)$, which is motivated by the findings in Kaur et al. (2025). The only restrictions imposed on $y(\cdot)$ are that it is continuous, strictly increasing, strictly concave, and $y(a_y) \in Y \subset \mathbb{R}_+$ is bounded above $Y = [0, \bar{y}]$, where $y(a_y = \infty) = \bar{y}$. To simplify the analysis, I also assume an Inada type condition $\lim_{a_y \to 0} \frac{dy}{da_y} = \infty$, which is satisfied by popular strictly concave functions. The utility from production $\omega(y(a_y))$ is continuous, increasing and concave[12]. This production process can be modeled differently, for instance incorporating effort costs; however, this would not bring additional insights to the main channel of scarcity that I am interested in.

Lastly, the agent's final utility $U : \Delta(X) \times Y \to \mathbb{R}_+$ aggregates the utility of the outcomes in all tasks. Since the consumption choice $x(a_x)$ is stochastic ex-ante, he will have to allocate the attention by maximizing his expected utility. For simplicity, I assume that the two tasks are independent and that $U(\cdot)$ is additively separable[13]:

$$\mathbb{E}U\Big(x(a_x), y(a_y)\Big) = \mathbb{E}\Big[u(x(a_x); \eta)\Big] + \omega(y(a_y))$$

3.3 Leading Example

Before solving the model, to fix ideas, consider the following example:

A financially constrained worker has to decide how much time and attention he should allocate between shopping and a productive task. His budget is insufficient to afford all

[12] Any function that preserves the properties of $y(\cdot)$ works, the simplest example is the identity function.

[13] Here, production in the productive task does not affect the agent's budget in the consumption task; this separation is built into the experimental design for simplicity. The assumption is also reasonable for many types of jobs with stable income (e.g., tenured professor), where producing more or less today does not necessarily impact salaries. Another interpretation is that the agent is myopic with respect to these interdependencies. In section C.1 of the Appendix, I relax this assumption and analyze the case in which the budget in the consumption problem depends on performance in the productive task.

he feels he needs, triggering the scarcity mindset. Scarcity pulls his attention towards the shopping task since any waste or inefficient decision with his little money is very costly. He carefully considers the trade-offs between picking each item and ends up with a more efficient bundle than usual. However, since attention is finite, he ends up with less of it for work, jeopardizing his productivity.

This example illustrates the case where the worker can only attend to one task or another; however, the same intuition applies for a worker who is trying to engage in a productive task but his mind is wandering about another problem. This would be more in line with the settings in Kaur et al. (2025) and Duquennois (2022). Another example is an academic trying to read a paper but has his attention captured by urgent deadlines (time scarcity).

3.4 Attention Allocation

In this model, the main interest is in the attention allocation problem. This is solved in two stages:

1. **Global attention allocation problem**

$$\max_{a_x, a_y} \mathbb{E}U\Big(x(a_x), y(a_y)\Big) \tag{2}$$

$$\text{s.t. } k_x a_x + k_y a_y \leq A \tag{3}$$

2. **Given attention, perform tasks**

   (i) The bundle $x(a_x)$ is chosen optimally given his consideration set, which probabilistically increases in size with a_x.

   (ii) $y(a_y)$ is produced deterministically as an increasing function of a_y.

To solve the global attention allocation problem, I work backwards to understand the expected utility given all attention allocation combinations. Since the production problem is straightforward, I start with the consumption problem.

Consumption problem with Random Attention

The consumption problem depends only on the attention allocated to it. I therefore describe it for a generic attention level $a \geq 0$; in the allocation problem, it is evaluated at $a = a_x$.

$$\max_x \; u(x;\eta)$$
$$\text{s.t. } \sum p_i x_i \leq b \text{ and } x_i = 0 \;\; \forall i \notin T(a)$$

The agent can choose quantities of multiple goods $x_i \in \mathbb{R}_+$, as long as they satisfy his budget constraint $x \in B(p, b)$. His final choice will be a bundle $x = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n_+$. However, the agent can only observe and compare the goods if he pays attention to them. Therefore, the more attention allocated, the more goods he will consider before making the choice. Finally, among the considered goods, the agent chooses the utility maximizing bundle subject to the budget constraint. This is formalized in the same spirit as the Random Attention Model (RAM) in Cattaneo et al. (2020).

Let $S = \{1, \ldots, n\}$ denote the set of all goods. A consideration set $T \subseteq S$ is the nonempty set of goods that the agent considers, and he can only buy goods in T. The attention rule $\mu(T|a)$ expresses the probability of consideration set T given the level of attention a allocated to the problem, where $\mu(T|a) \geq 0$ and $\sum_{T \subseteq S} \mu(T|a) = 1$, with the sum taken over nonempty subsets, as are all sums over T below. I let these probabilities depend only on the size of T and on attention. With some abuse of notation, denote $\mu(T|a) = \mu(|T|, a)$, which means μ is a function of a and the size of T, and $\mu(T'|a) = \mu(T|a)$ if $|T'| = |T|$. Intuitively, if someone makes a choice with almost no attention, they will consider very few alternatives and their decision will be seemingly random. On the contrary, with full attention, all alternatives are considered. This is captured in three assumptions on $\mu(\cdot)$:

Assumption 1 (Singleton)**.** *For* $|T| = 1$*:* $\frac{\partial \mu(T|a)}{\partial a} < 0$

Assumption 2 (Full attention)**.** $\lim_{a \to \infty} \mu(S|a) = 1$

Assumption 3 (Supermodularity)**.** *For any* $a' > a$*,* $|T'| > |T|$*:*

$$\mu(T'|a') - \mu(T|a') \geq \mu(T'|a) - \mu(T|a)$$

Supermodularity is the discrete analogue of $\frac{\partial^2 \mu(T|a)}{\partial a \partial |T|} \geq 0$: more attention shifts probability from smaller toward larger consideration sets. To guarantee interior solutions, I also assume an Inada type condition on μ. Formally, $\forall T$ s.t. $|T| = 1$, $\lim_{a \to 0} \frac{\partial \mu(T|a)}{\partial a} = -\infty$.

Lemma 3.1. *If Assumptions 1 and 3 hold, then for every* a *there exists a threshold* $|T_0|$*, which may depend on* a*, such that* $\frac{\partial \mu(T|a)}{\partial a} \leq 0 \ \forall |T| \leq |T_0|$*, and* $\frac{\partial \mu(T|a)}{\partial a} > 0 \ \forall |T| > |T_0|$*.*

Finally, for a given consideration set T, the agent picks the best affordable bundle that uses only goods in T,

$$x_T := \arg\max\{u(x; \eta) : x \in B(p, b), \ x_i = 0 \ \forall i \notin T\}.$$

Therefore, the agent first forms a consideration set T and from it picks the bundle that maximizes his utility given prices and his budget, so that bundle x is chosen with probability

$\pi(x|a) = \sum_{T \subseteq S} \mu(T|a)\, \mathcal{I}\{x = x_T\}$. This is in line with evidence from Reutskaja et al. (2011), who find that individuals are good at optimizing within the set of items they consider during the search process. Since preferences are strictly convex and, for any consideration set, the set of affordable bundles using only the considered goods is convex, the optimal bundle is unique.[14]

The formalization of choices from consideration sets has become popular among decision theorists as can be seen in Abaluck and Adams-Prassl (2021), Cattaneo et al. (2020), Manzini and Mariotti (2014) and Masatlioglu et al. (2012). The main departure from these models is that I introduce an endogenous level of attention that affects the consideration set probabilities.

The agent then maximizes his expected utility given a choice problem S, price $p \in \mathbb{R}^n_+$, budget $b \in \mathbb{R}_+$, and level of attention $a \in \mathbb{R}_+$. The expected utility can be written as

$$\mathbb{E}[u(x(a);\eta)] = \sum_x \pi(x|a)\, u(x;\eta) = \sum_{T \subseteq S} \mu(T|a)\, u(x_T;\eta).$$

Lemma 3.2. *The higher the allocated attention to consumption, the greater the expected consumption utility.*

Throughout the theoretical framework, the attention rule μ, the choice of the best bundle within the consideration set, and the production function $y(\cdot)$ describe how the agent *perceives* the returns to attention in each task. The attention allocation problem is formulated entirely in terms of these perceived returns. Whether actual performance responds to attention in the same way is discussed after Proposition 1.

3.4.1 Benchmark with Unlimited Attention

In order to understand the mechanism of attention allocation under limited attention, I first solve (2) for the case where $A = \infty$. Due to unlimited attention, the constraint $k_x a_x + k_y a_y \leq A$ never binds, so the agent can pay arbitrarily large attention to both tasks; I consider the limit $a_x, a_y \to \infty$. The solution is straightforward and matches that of a fully rational agent with no effort costs and that maximizes his choice over all available goods.

In the consumption problem, Assumption 2 implies that $\mu(S|a) \to 1$ as $a \to \infty$, so that

$$\lim_{a \to \infty} \mathbb{E}[u(x(a);\eta)] = u(x_S;\eta) = u(x^*;\eta),$$

[14] In the experiment, goods are indivisible, so uniqueness is not guaranteed. However, the experimental design ensures that all rounds have a unique optimal bundle for each condition.

where $x^* = \arg\max_{x \in B(p,b)} u(x;\eta)$ is the optimal bundle when all goods are considered.

In the production problem, since output is increasing in attention, it approaches its upper bound: $y(a_y) \to \bar{y}$ as $a_y \to \infty$.

3.4.2 Solution with Limited Attention

I now solve the attention allocation problem (2) when $A < \infty$. Before getting to the interesting cases, note that when the attention constraint does not bind, the agent will obtain the same outcomes as in the benchmark with unlimited attention. Hereafter, I focus on cases when either k_j are high enough or A small such that the attention constraint has some bite. I proceed by defining the Lagrangian and finding the optimal solutions.

$$\begin{aligned}
\mathcal{L} =& \mathbb{E}\Big[u(x(a_x);\eta)\Big] + \omega(y(a_y)) - \phi[k_x a_x + k_y a_y - A] \\
&\frac{\partial \mathcal{L}}{\partial a_y} = \frac{\partial \omega(y(a_y))}{\partial y}\frac{\partial y}{\partial a_y} - \phi k_y = 0 \qquad (4) \\
&\frac{\partial \mathcal{L}}{\partial a_x} = \frac{\partial \mathbb{E}\Big[u(x(a_x);\eta)\Big]}{\partial a_x} - \phi k_x = 0 \qquad (5) \\
&\phi(k_x a_x + k_y a_y - A) = 0 \qquad (6)
\end{aligned}$$

The case $\phi = 0$ was solved above, so I focus on cases where $\phi > 0$; by (6) the constraint will bind in optimality. By combining the first two expressions, one obtains the results analogous to a standard consumption choice problem:

$$\frac{k_x}{k_y} = \frac{MU_{a_x}}{MU_{a_y}} = MRS_{xy} \quad \text{and} \quad k_x a_x^* + k_y a_y^* = A \qquad (7)$$

Thanks to the Inada assumption on both $y(\cdot)$ and $\mu(\cdot)$, corner solutions are ruled out, which simplifies the analysis. Thus, at any optimum, which is interior, the marginal rate of substitution of allocating attention to task x instead of task y is equal to the ratio of their complexity.[15] From this optimality condition, one can perform comparative statics to obtain predictions. In particular, the comparative static on the threshold of need η gives rise to the main result: holding the budget and attention allocation problem fixed, differences in the subjective need η can generate differences in the marginal value of attention to consumption. Therefore, two individuals facing identical objective constraints may allocate attention differently because one perceives scarcity while the other does not.

[15] The main results do not require the optimal allocation to be unique. Uniqueness is guaranteed if, for every m, the probability of considering at least m goods is concave in attention, so that attention has diminishing returns in the consumption task, as in the productive task.

Before stating this result, I first show that (i) when the agent's budget is below a threshold determined by η, he perceives scarcity, and (ii) perceived scarcity increases the marginal value of allocating attention to the consumption task.

Definition 3.3 (Need Budget)**.** $b_\eta := \min_x\{p \cdot x \mid c(x) \geq \eta\}$ represents the minimum amount of budget needed to meet the perceived need η.

Lemma 3.3. *The agent will perceive scarcity if $b < b_\eta$. Moreover, when he considers all goods, he perceives scarcity if and only if $b < b_\eta$.*

The threshold b_η highlights two ways perceived scarcity can arise: through a decline in purchasing power, such as a negative income shock or rising prices, or through an increase in subjective needs. Importantly, perceived scarcity can therefore change even without a change in objective resources.

Lemma 3.4. *Holding objective resources fixed, the marginal value of attention to consumption is greater for an individual who perceives scarcity than for an otherwise identical individual who does not, at every level of attention. Formally, if $b < b_{\eta^H}$ and $b > b_{\eta^L}$, then $\frac{\partial \mathbb{E}[u(x(a);\eta^H)]}{\partial a} > \frac{\partial \mathbb{E}[u(x(a);\eta^L)]}{\partial a}$ for every $a > 0$.*

Proposition 1 (Attention Allocation)**.** *Consider two otherwise identical individuals with the same budget b, but different subjective needs η^H and η^L, such that $b < b_{\eta^H}$ and $b > b_{\eta^L}$. Then, for any optimal allocation of attention:*

$$a_x^*(b, \eta^H) > a_x^*(b, \eta^L) \text{ and } a_y^*(b, \eta^H) < a_y^*(b, \eta^L).$$

Therefore, the individual perceiving scarcity allocates more attention to consumption at the expense of reduced attention in the productive task.

This theoretical result aligns with Herbert Simon's idea that the human cognitive system incorporates an interruption mechanism that reallocates attention toward addressing urgent internal needs (Simon 1967). In this case, the perception of scarcity draws attentional resources to the task in the scarce domain.

Proposition 1 isolates the effect of perceived scarcity by comparing individuals with the same objective resources but different subjective needs. A natural question is whether tighter resource constraints alone, with smooth, concave utility and no subjective need, would generate the same reallocation of attention. After all, individuals with fewer resources have higher marginal utility of consumption, which might make allocating attention to consumption more attractive.

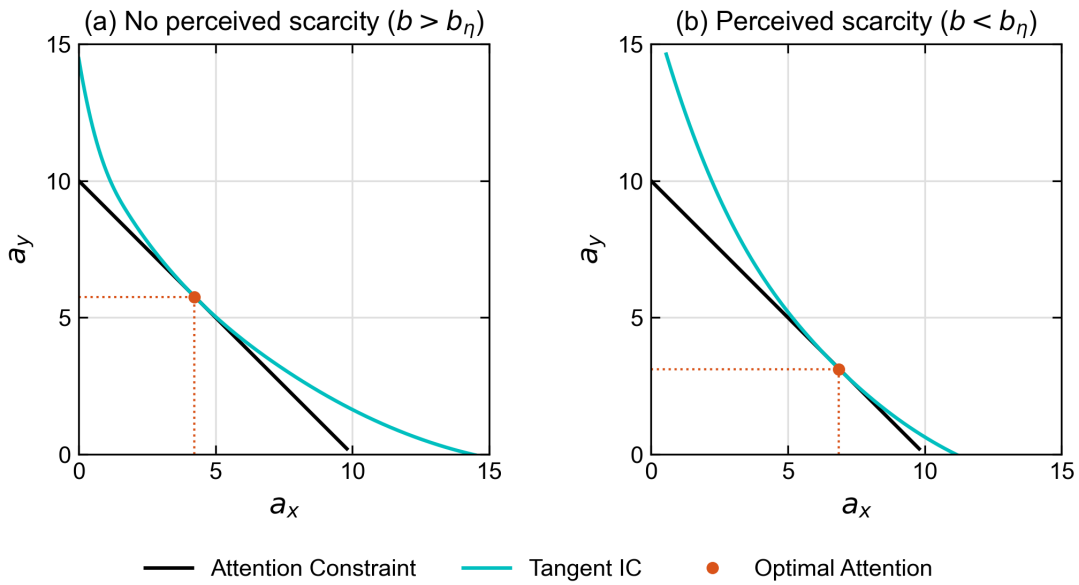


Figure 2: Optimal attention allocation without (a) and with (b) perceived scarcity. The attention constraint is the same in both panels. Perceived scarcity makes the indifference curve steeper, so the tangency moves toward attention to consumption (a_x) and away from the productive task (a_y).

Remark 1 (Objective versus perceived scarcity)**.** A lower budget does not necessarily increase attention to consumption. Without incorporating subjective needs in the model, a lower budget has two opposing effects: higher marginal utility makes consumption gains more valuable, but a smaller budget reduces the consumption gains attainable from better choices.[16] Which effect dominates depends on the elasticity of marginal utility, $\varepsilon(c) = -cv''(c)/v'(c)$. With substitutable goods, a lower budget increases attention to consumption if $\varepsilon(c) > 1$ for all c, decreases it if $\varepsilon(c) < 1$ for all c, and has no effect on attention under logarithmic utility. The derivation is provided in Appendix A.

Thus, observing greater focus on a domain and neglect elsewhere among individuals with fewer resources does not by itself identify an effect of perceived scarcity, and its absence does not rule one out. This distinction motivates the experiment, which induces differences in perceived scarcity while holding current resource constraints and incentives fixed.

The model also suggests that task complexity can affect the size of spillovers. Since complexity enters the attention constraint, increasing the complexity of the consumption

[16] With a subjective need, a lower budget can also push consumption below the need threshold (Lemma 3.3), introducing an additional effect on the value of attention. However, the overall effect remains ambiguous because the opposing forces described in the remark still apply. The remark therefore abstracts from subjective needs to isolate the effect of objective resource constraints.

task makes each unit of attention allocated to consumption more costly in terms of foregone production attention. This can amplify the negative spillover on production, provided that the scarcity-induced increase in consumption attention does not decline too quickly as the task becomes more complex. Thus, complexity not only affects task difficulty, but also how costly it is to reallocate attention across tasks.

Proposition 1 concerns how attention is allocated, which depends on the returns to attention as perceived by the agent. Whether this reallocation also changes performance depends on the actual relationship between attention and performance in each task.[17] More attention is generally expected to improve performance, but the relationship may be flat or non-monotonic (Camerer and Hogarth 1999; Kahneman et al. 1968; Libby and Lipe 1992). In the consumption task, the model assumes that the agent selects the optimal bundle from his consideration set, so that more attention always weakly improves the chosen bundle. In practice, however, as the consideration set grows, the number of feasible bundles increases rapidly, making it harder to select the best one (Murawski and Bossaerts 2016). This is particularly relevant when goods are indivisible, as in the experiment. If the agent does not anticipate this difficulty, he may allocate more attention expecting better choices, even though actual decision quality may improve more slowly or even decline. The same applies to the productive task, where actual output may not respond to attention as the agent expects. Therefore, greater attention need not translate into better performance.

Corollary 1 (Performance)**.** *Consider the two individuals in Proposition 1. The individual perceiving scarcity (i) makes better consumption decisions if actual decision quality is increasing in attention to consumption, and (ii) performs worse in the productive task if actual output is increasing in attention to production.*[18]

The distinction between perceived and actual returns to attention highlights why process-tracing data are particularly useful for studying scarcity. If tasks have an increasing relationship between attention and performance, the model predicts that perceived scarcity improves decision quality in the scarce domain while reducing performance in other tasks. However, if the relationship is flat or non-monotonic in a given task, changes in attention need not translate into improvements in performance. A researcher who only measures task performance and observes a null effect may then conclude that this psychological channel is inactive, even when attention has shifted. Such changes in attention can be welfare-relevant

[17] Alaoui and Penta (2022) similarly emphasize the importance of separating reasoning from performance.

[18] The increasing relationship between attention and performance only needs to hold over the range of attention levels chosen with and without perceived scarcity, rather than globally.

in themselves (Quispe-Torreblanca et al. 2026; Bolte and Raymond 2026; Loewenstein and Wojtowicz 2025) and may also generate spillovers on other, unmeasured economic outcomes.

The theoretical results offer a unified framework for understanding the positive, negative, and null effects of scarcity documented in the literature. These effects can arise even when individuals optimally allocate their limited attention based on expected returns. Perceived scarcity redirects attention toward the scarce domain and away from competing tasks, while the resulting effects on performance depend on how attention translates into performance in each task.

3.5 Endogenous Subjective Need

The previous section treats the subjective need threshold η as given. This is useful for characterizing how perceived scarcity affects attention allocation, but in reality subjective needs are formed based on experience. What an individual perceives as enough may depend on past consumption, the consumption of others, or the uncertainty they face about future requirements. These sources can be formalized as different determinants of η. In the main text, I focus on history dependence, which is the source of subjective need used in the experimental design. Additional examples based on social comparison and environmental uncertainty are discussed in Appendix C.2.

The history-dependence formulation is closely related to backward-looking reference points and habit formation (Abel 1990; Bowman et al. 1999; Carroll et al. 2000; DellaVigna et al. 2017). The basic idea is that individuals adapt to previously available resources. A higher past consumption level raises the amount that feels necessary in the present, while a lower past consumption level lowers this benchmark. Thus, two agents with the same current budget may differ in whether they perceive scarcity because they have adapted to different consumption histories.

Formally, consider that over t periods, the agent has accumulated a history of consumption levels, denoted by $\mathcal{H}_{t-1} = \{c_0, c_1, \ldots, c_{t-1}\}$, where $c_s = c(x_s)$ is the level of consumption in period s. The perceived need at time t, η_t, evolves according to the recursive relation:

$$\eta_t = \rho\eta_{t-1} + (1-\rho)c_{t-1}$$

where $\rho \in [0, 1)$ captures the persistence of past needs relative to recent consumption. A higher ρ implies that past needs weigh more heavily, while a lower ρ indicates that recent

consumption has a stronger influence. Assuming $\eta_0 = 0$, by iteration this yields:

$$\eta_t = (1 - \rho) \sum_{j=0}^{t-1} \rho^j c_{t-1-j}$$

Thus, η is shaped by past consumption, with recent experiences receiving greater weight. This formulation provides the basis for the experimental manipulation. By exposing participants to different budget histories, the experiment generates variation in subjective needs, so that the same objective resources may feel sufficient to some participants and insufficient to others.

4 Experiment

4.1 Experimental Design

Motivated by the theoretical framework, the experiment examines how perceived scarcity affects attention allocation, decision-making, and productivity. Participants complete two tasks simultaneously: (i) a productivity task based on cognitive sliders, and (ii) a consumption-choice task modeled as a knapsack problem. Unlike Avoyan and Schotter (2020), who ask participants to commit to a time allocation *ex-ante* across strategic games and only see one task at a time; in this experiment, participants can freely switch between tasks at any point. Attention is directly measured using Mouselab (Johnson et al. 1989), which records the time spent attending to each task and transitions between them. Participants complete multiple rounds of this dual-task and after each round we elicit their perception of resources, their confidence on how close to the optimal bundle they were, and how hard they felt each task was.

Cognitive Sliders. The cognitive slider task serves as the productivity task and captures output that depends on the attention and effort participants devote to it. This is a variant of the classic slider task, which is widely used in experimental economics (Gill and Prowse 2019; Goerg et al. 2019), including by Dalton et al. (2017) to study the effect of poverty exposure on productivity. It is particularly appropriate here because it requires sustained focus, creating a natural attentional trade-off when participants perform it simultaneously with the consumption task.

I introduce a modified version of the slider task that makes performance depend more

directly on cognitive processing.[19] Participants are shown a series of sliders, but instead of observing the target value directly, they are given a mathematical expression that they must solve first. They then move the slider to match the computed value, as illustrated in Figure 3. Complexity can be increased by increasing the number of arithmetic operations involved as in Caplin et al. (2011). The task is identical across participants.

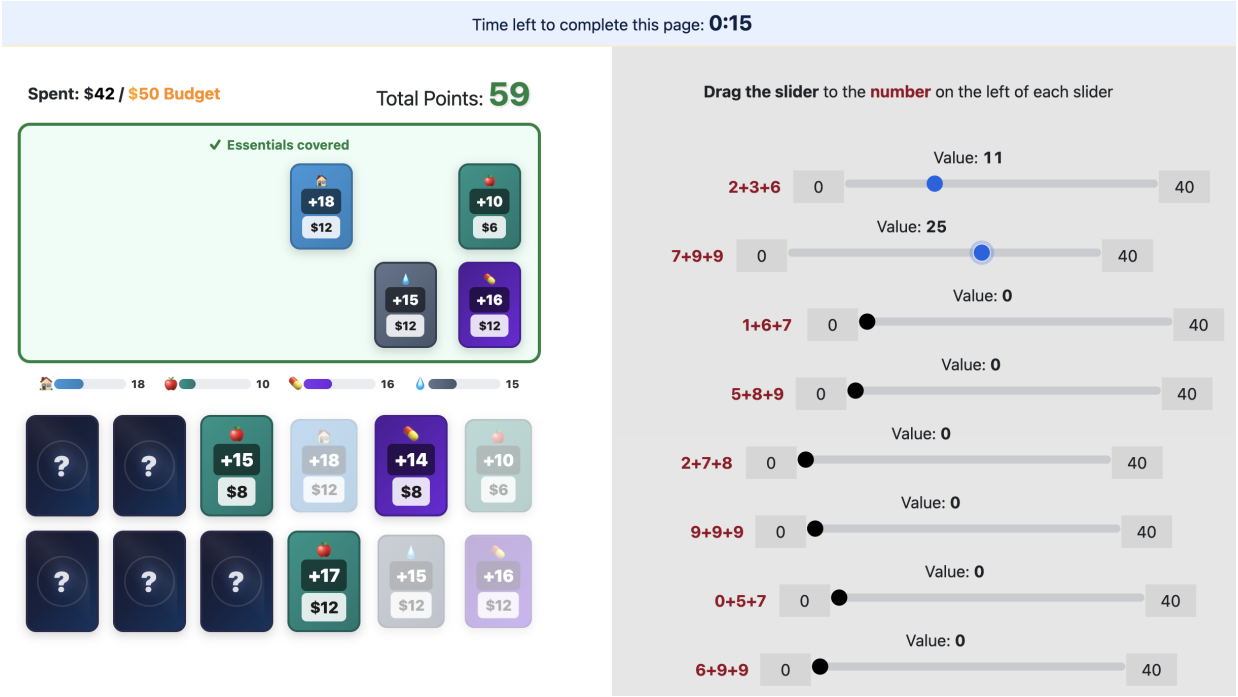


Figure 3: Dual-task interface: Consumption on the left and Cognitive Sliders on the right. The gray shaded area indicates which task was being covered, which for illustration is much less opaque than in the actual experiment.

Consumption Task. The consumption task is implemented experimentally as a knapsack problem, which involves selecting a bundle of goods subject to a budget constraint (Figure 3). Standard economic models assume that goods are perfectly divisible and that individuals can choose optimally from the consumption set. In reality, goods are often indivisible, which makes the consumption problem computationally difficult (Gilboa et al. 2021). The knapsack problem captures this feature, and previous work shows that performance improves with the attention devoted to the task (Murawski and Bossaerts 2016).

To make the task more realistic, goods are divided into four essential categories: rent, food, medicine, and utilities. Participants must select one item from each category in order

[19] This version was inspired by the sums task in Caplin et al. (2011). Since the model's predictions concern cognitive resources, the task is designed to require not only effort but also mental computation.

to receive any payment in this task. Conditional on meeting this requirement, payment is increasing in the total number of points collected.

Formally, let $G = \{\text{rent}, \text{food}, \text{medicine}, \text{utilities}\}$ denote the set of essential categories, and let $I_g \subseteq \{1, \ldots, n\}$ be the set of items belonging to category $g \in G$. Participants choose a bundle of goods $i = \{1, \ldots, n\}$ with value v_i and price p_i, subject to a budget b:

$$\begin{aligned} \max_{x \in \{0,1\}^n} \quad & \sum_{i=1}^{n} v_i x_i \\ \text{s.t.} \quad & \sum_{i=1}^{n} p_i x_i \leq b, \\ & \sum_{i \in I_g} x_i = 1 \qquad \forall g \in G. \end{aligned}$$

Motivated by the theoretical framework, items are initially hidden and can only be revealed by hovering the mouse over them for one second. This feature creates a direct link between attention and consideration set: the more attention a participant allocates to the task, the more items they uncover and can evaluate. It also allows the set of revealed items to serve as an empirical proxy for the participant's consideration set.

All participants face the same set of 18 items in the same order across all rounds, with only the budget varying by treatment condition. Therefore task complexity does not vary neither across rounds nor across participants. To reduce learning or carryover effects, the test round uses a new decision environment: the optimal solution does not overlap with the optimal solutions from the earlier rounds in neither treatment condition.

Incentives. The earnings in each task were independent from each other. Importantly, the relative incentives between both tasks were held constant across treatment and control groups. Participants received a piece-rate payment of £0.10 for every slider completed correctly. In each round there were 30 sliders, so they could earn up to £3.00. To qualify for payment in the consumption task, participants were required to select one item from each essential category. Conditional on meeting that requirement, they were paid an amount that was a linear increasing function of the total points collected relative to the maximum based on their budget. If they selected the optimal bundle, they received an additional £0.30 bonus. The most they could earn in this task was £3.30. For instance, if the maximum was 100pts and they collected 50pts, they would be paid £1.50. Therefore, in each round, the combined earnings could be up to £6.30. One round was randomized for payment.

4.2 Inducing Perceived Scarcity

Perceived scarcity is induced by manipulating participants' histories of consumption. Participants are randomly assigned to a High-Budget Adaptation condition (HBA) or Low-Budget Adaptation condition (LBA). During the first three adaptation rounds, the HBA group receives higher budgets ($78, $79, $80) in the consumption task than the LBA group ($30, $29, $32), habituating them to higher levels of consumption. In a subsequent test round, both groups receive the *same* budget ($39), which lies between their previous endowments. The test-round budget is substantially lower than the large adaptation budgets and slightly higher than the small adaptation budgets. This design is intended to create a large downward change for participants with a large-budget history, strengthening the perceived-scarcity manipulation, while avoiding a large upward change for participants with a small-budget history that could generate a separate feeling of abundance.[20]

Although objective resource constraints are identical in the test round, the HBA group is more likely to perceive that their budget is insufficient relative to their needs compared to the LBA group. This design allows for the study of perceived scarcity effects while holding the task environment constant across groups.

Perceived scarcity is elicited after each round through the question: "Did you feel you had enough budget?" which is answered on a scale from 1 ("I did not have enough") to 5 ("I had more than enough"). They are coded as perceiving scarcity if they answer 1 or 2. Participants also answer how confident they were about their consumption-task choice, and how difficult they found each task.

4.3 Hypotheses

The theoretical framework generates predictions for the common-budget test round, in which both treatment groups face the same objective constraint, incentives, and task environment. The treatment varies only participants' prior budget histories.[21] HBA participants are expected to enter the test round with a higher subjective need threshold than LBA participants. When both groups then face the same intermediate budget, the HBA group should be more likely to perceive the budget as insufficient.

[20] Importantly, both groups experience a large enough change in budget that they have to fully reoptimize to obtain the optimal bundle. This generates a new decision environment and reduces the risk of inertia.

[21] For this reason, the analysis only compares the two groups in the test round rather than using a difference-in-differences design. The budget history is itself the treatment, so previous rounds cannot serve as a pre-treatment baseline.

Hypothesis 1 (Perceived scarcity)**.** HBA participants are more likely than LBA participants to report perceived scarcity in the test round.

The main prediction concerns attention allocation. By Proposition 1, perceived scarcity increases the marginal value of attention in the consumption task. Since total attention is limited, this reallocates attention toward the consumption task and away from the productivity task. The central hypothesis is therefore:

Hypothesis 2 (Attention allocation)**.** HBA participants allocate a larger share of attention to the consumption task in the test round.

A secondary prediction concerns search behavior within the consumption task. In the model, attention to consumption expands the agent's consideration set: with more attention, the agent considers more alternatives before choosing. In the experiment, items in the consumption task are initially hidden and can only be revealed by hovering over them. The number of revealed items therefore provides a behavioral proxy for the size of the participant's consideration set.

Hypothesis 3 (Consumption Search)**.** HBA participants reveal more items and have a larger consideration set in the test round.

The remaining hypotheses concern task performance, which depends on the production technology. If the relationship between attention and performance is always increasing, Corollary 1 predicts that the reallocation of attention improves choices in the consumption task while reducing performance in the productivity task.

Hypothesis 4 (Consumption-task performance)**.** HBA participants choose higher-valued bundles in the consumption task in the test round.

Hypothesis 5 (Productivity performance)**.** HBA participants perform worse on the cognitive slider task in the test round.

5 Experiment Results

The study was conducted online on Prolific with 501 participants from the UK, randomly assigned to either the HBA (250 participants) or LBA (251 participants) condition. The median completion time was 18 minutes. Participants earned £4.34 on average, including

a £1.80 show-up fee, with a maximum possible payment of £8.10. The experiment was preregistered in the AEA RCT Registry (Sewell 2026).[22]

Result 1: Perceived scarcity differs under identical resource constraints. The experiment successfully generates substantial differences in perceived scarcity even when participants face identical objective resource constraints. Panel (a) of Figure 4 shows that, during the three adaptation rounds, participants' perceptions of scarcity are consistent with their assigned budgets. Crucially, when both groups face the *same* intermediate budget in the test round, 77% of HBA participants report feeling scarce, compared to only 45% of LBA participants ($p < 0.001$). This 32 percentage-point difference demonstrates that perceived scarcity depends not only on current resources but also on prior resource experience.

Result 2: Perceived scarcity redirects attention toward the scarce domain. The treatment also generates a substantial reallocation of attention toward the consumption task, as predicted by the model. In the test round, HBA participants devote 69.7% of their total task time to consumption, compared to 59.7% among LBA participants ($p < 0.001$; Figure 4, panel (b)). This 10 percentage-point difference provides direct evidence that prior resource experience can redirect attention toward the domain perceived as scarce, even when current resource constraints and incentives are identical.

How this reallocation of attention affects performance depends on how attention translates into performance in each task (Corollary 1). Figure 5 illustrates these relationships, pooling observations across all four rounds. In the productivity task, performance increases with time throughout the observed range. In contrast, in the consumption task, decision quality rises steeply up to about 30 seconds and is essentially flat thereafter.[23] Since participants endogenously choose how much time to allocate to each task, these relationships are descriptive rather than causal. Nevertheless, they suggest that reallocating attention toward consumption should reduce productivity, without necessarily improving consumption decisions.

Result 3: Perceived scarcity intensifies search without significantly improving consumption decisions. The treatment also intensified search within the consumption

[22] Concerns about completion by AI agents are reduced by the substantial mouse interaction required in both tasks, including hovering and dragging sliders. I also included the AI-detection task of Celebi et al. (2026), who find that only 17% of AI agents pass its video attention check. In my sample, 497 of the 501 participants passed this check, and all results are unchanged when excluding those who failed it.

[23] Figures 12 and 13 in Appendix E show these relationships separately for each round. Table 3 compares the slopes in the first and second half of the 70-second round. In every round, the number of correct sliders still increases significantly above 35 seconds, by 1.1 to 1.7 sliders per 10 seconds, compared with 1.9 to 2.0 below. Decision quality, by contrast, rises by about 18 percentage points per 10 seconds below 35 seconds, but changes by at most 3.8 percentage points per 10 seconds above it (-1.7 points when pooling all rounds).

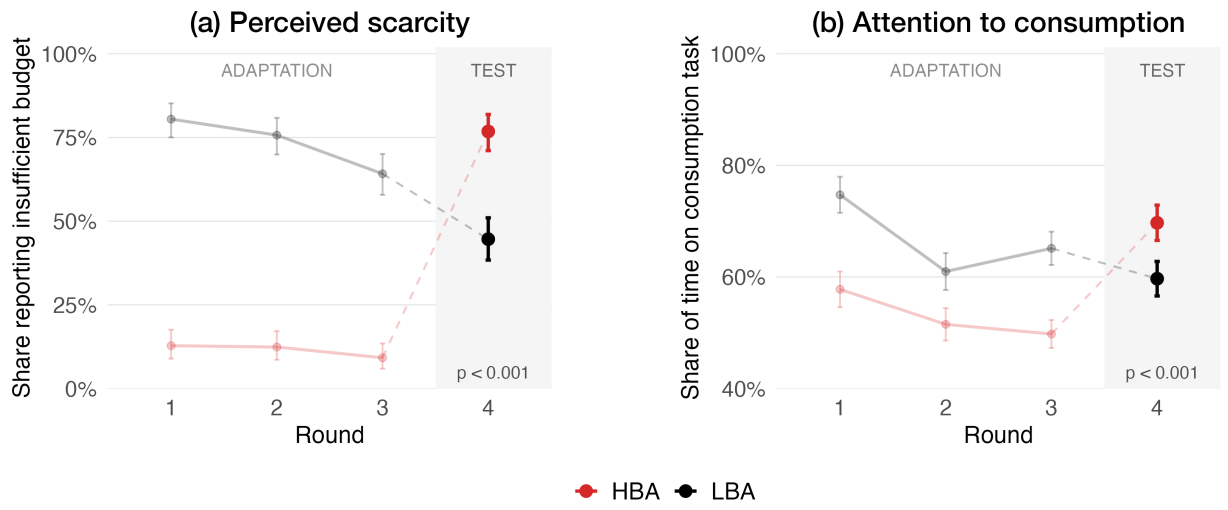


Figure 4: Perceived scarcity and attention allocation by round. Panel (a): share of participants reporting that their budget was insufficient. Panel (b): share of total task time spent on the consumption task. Error bars are 95% confidence intervals.

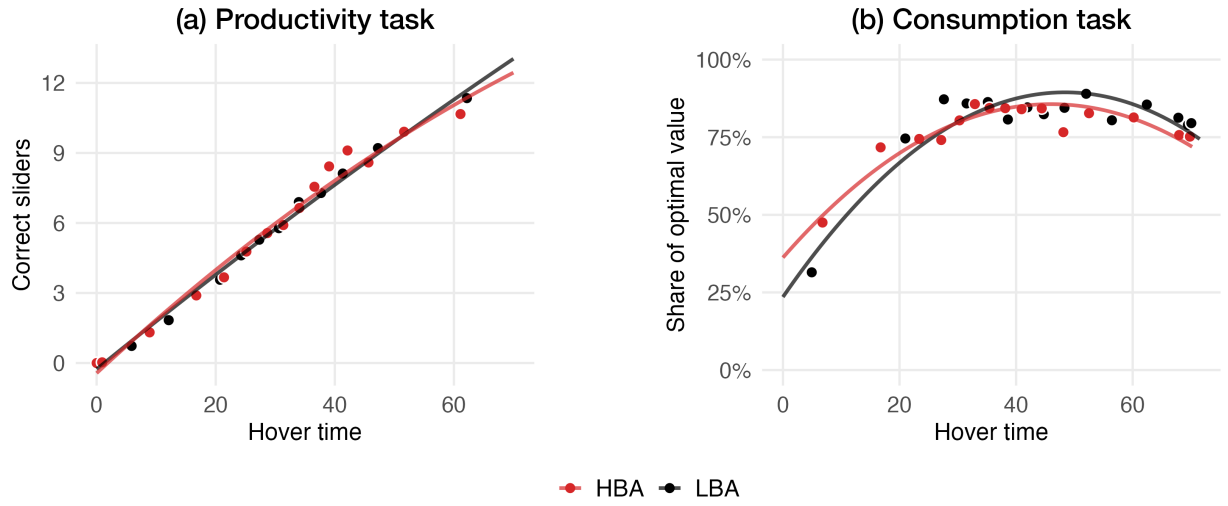


Figure 5: Attention–performance relationships. Panel (a) shows correct sliders against time spent on the productivity task. Panel (b) shows chosen-bundle value as a share of the optimum against time spent on the consumption task. All four rounds are pooled. Within each group, participant-rounds are split into 15 equally sized bins of hover time, and each point shows a bin's mean. Lines are quadratic fits to the individual observations. Because a quadratic is symmetric around its peak, the fit in panel (b) turns down after it; however, the binned means are essentially flat beyond about 30 seconds.

task, consistent with Hypothesis 3. In the test round, HBA participants considered more items, with an average consideration set of 14.96 compared to 14.09 in the LBA group (a

difference of 0.87 items, $p = 0.031$; Figure 6, panel (a)). They also revised their selected bundles more frequently, making 1.47 item-level changes on average compared to 0.81 in the LBA group ($p < 0.001$; panel (b)). Thus, the treatment affected not only how much time participants devoted to consumption, but also how they searched within the task.

Despite this more intensive search, HBA participants did not make significantly better consumption choices, contrary to Hypothesis 4. Average decision quality in the test round was 0.704 in the HBA group and 0.722 in the LBA group ($p = 0.370$; Figure 6, panel (c)). This finding illustrates how greater attention need not translate into better performance (Corollary 1), consistent with the flat relationship beyond about 30 seconds in Figure 5. Importantly, this highlights the value of collecting choice-process data. A study measuring only decision outcomes might have concluded that perceived scarcity had no effect in the consumption domain, overlooking the substantial changes in search behavior induced by the treatment.

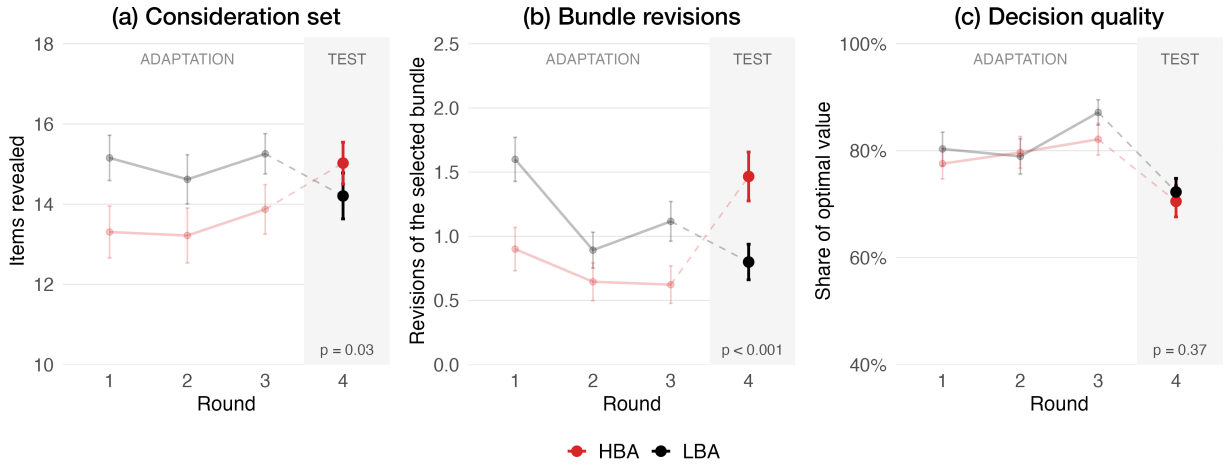


Figure 6: The consumption task by round. Panel (a): number of items revealed. Panel (b): number of times participants revised their selected bundle. Panel (c): decision quality, the value of the chosen bundle as a share of the value of the optimal bundle.

Result 4: Perceived scarcity reduces performance in the productivity task. As predicted by Hypothesis 5, HBA participants performed worse on the productivity task. In the test round, they completed 4.1 correct sliders on average, compared to 5.4 in the LBA group ($p < 0.001$; Figure 7, panel (a)). This represents a 24% reduction in productivity, consistent with the reallocation of attention away from the productivity task documented in Result 2.

Result 5: Perceived scarcity reduces total earnings. Ultimately, the treatment had substantial consequences for overall economic performance. Total earnings across both tasks

were 15% lower in the HBA group (£2.14 vs. £2.53, $p < 0.001$; Figure 7, panel (b)). Thus, the productivity losses were not offset by improvements in consumption decisions. Taken together, these findings show that perceived scarcity can have substantial economic consequences even when current resource constraints and incentives are identical.

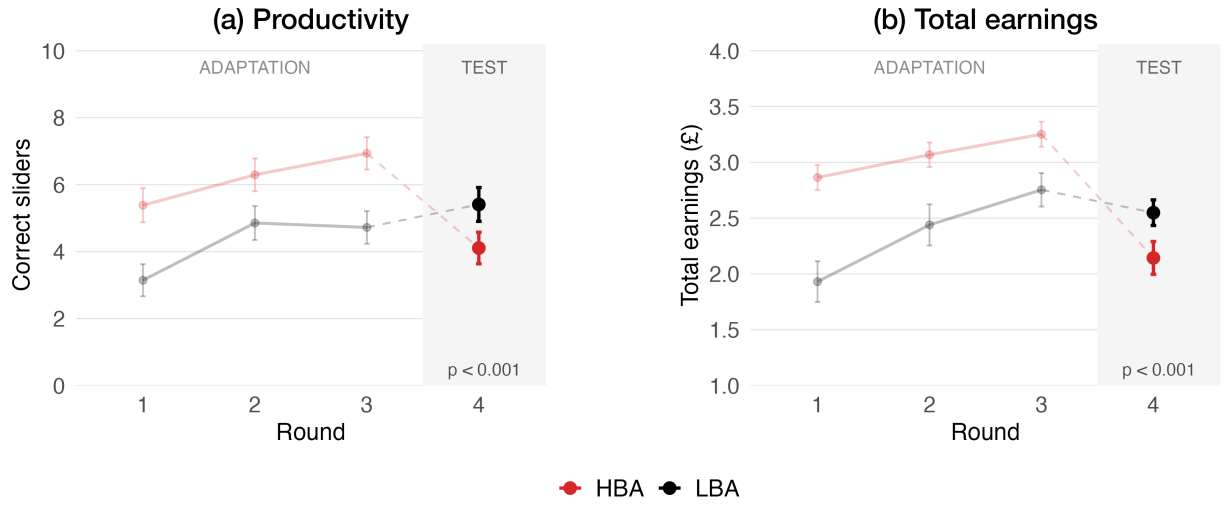


Figure 7: Productivity and total earnings by round. Panel (a): number of sliders completed correctly in the productivity task. Panel (b): total earnings from both tasks in the round.

6 Field Evidence: Job Search in Hungary

The 2005 Hungarian unemployment insurance reform provides a particularly useful setting to investigate whether past economic experiences influence job-search behavior under identical current incentives. The reform front-loaded unemployment benefits but left the benefit schedule from day 270 onward unchanged. Following DellaVigna et al. (2017), I compare the job-finding rates and reemployment wages of workers entering unemployment before and after the reform, focusing on periods in which both groups face the same current and future benefits but have experienced different benefit histories. This closely resembles the experimental manipulation, in which participants face the same current budget after experiencing different budget histories.

6.1 Setting and Data

Before November 2005, unemployment insurance in Hungary paid a constant benefit for up to 270 days. The reform split this period into two steps: a higher benefit during the first 90

days and a lower benefit during the following 180 days, while keeping the total amount paid roughly unchanged (Lindner and Reizer 2020). After day 270, both groups received the same, much lower unemployment assistance, which ended at day 360. Thus, pre-reform workers reached day 270 after receiving a constant benefit, while post-reform workers had already experienced six months of reduced benefits. The pre-reform group therefore resembles the HBA group in the experiment: both groups face the same current resources, but the pre-reform group experiences a larger drop relative to its recent benefit history. Figure 8 shows the two benefit schedules.

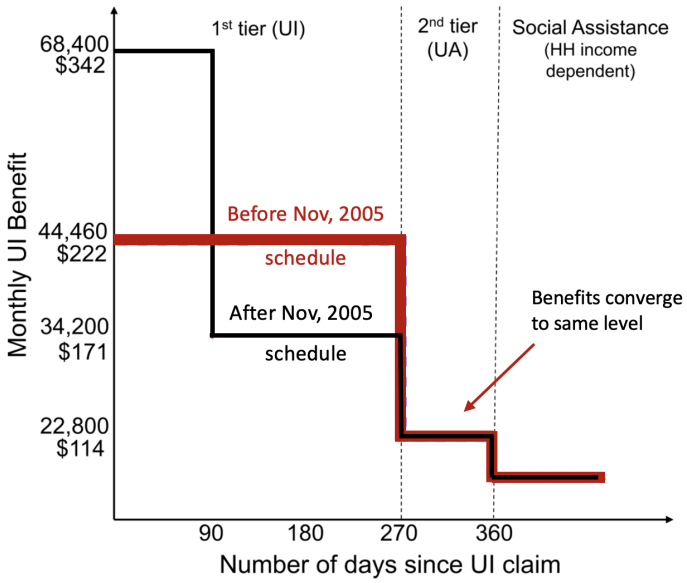


Figure 8: Unemployment benefit schedules before and after the November 2005 reform, for claimants at the maximum benefit, in Hungarian forints per month (US dollars below). From day 270 onward, both schedules pay the same benefits. Adapted from DellaVigna et al. (2017) and Lindner and Reizer (2020).

I use administrative data provided by the Databank of the ELTE Centre for Economic and Regional Studies, covering a 50% random sample of the Hungarian population between 2002 and 2008. Following DellaVigna et al. (2017), I compare workers who claimed benefits between 5 February and 15 October 2005 (pre-reform) with those who claimed benefits over the same calendar dates in 2006 (post-reform). I also follow their sample restrictions, keeping claimants aged 25 to 49 who were eligible for the full 270 days of benefits and whose pre-unemployment earnings were above the 70th percentile. The final sample contains 11,610 claimants: 6,164 pre-reform and 5,446 post-reform.

I study two main outcomes. The first is the job-finding rate, measured as the share of claimants still unemployed at the start of a 15-day period who find a job within that period.

The second is the reemployment wage, measured as the log of first-month earnings in the new job relative to the earnings base used to calculate unemployment benefits (Lindner and Reizer 2020). This provides a measure of reemployment earnings relative to previous earnings, which I use as a proxy for match quality.

I focus on days 271–315, the first 45 days after both groups start receiving the same benefits, when their different benefit histories are most recent. For comparison, I also examine the preceding 45 days (226–270) and the following 45 days (316–360). Within each window, I compare the two groups after controlling for age, sex, previous occupation and skill level, pre-unemployment earnings, time between job loss and the benefit claim, day of the month of the claim, and county.[24]

This is a comparison of entry cohorts rather than a difference-in-differences design, since the periods before day 270 already reflect different benefit histories. Because the cohorts may also differ in other ways and attention is not directly observed, I interpret the results as suggestive evidence consistent with the model, rather than causal evidence of the perceived-scarcity mechanism.

6.2 Model Predictions

Appendix B applies the model to job search. Workers do not know all available vacancies and allocate attention between evaluating job opportunities and submitting applications. Evaluation expands the worker's consideration set, increasing the chances of finding jobs that better match his skills and offer higher wages. Applications, in turn, increase the probability of receiving an offer. As in the experiment, subjective needs are assumed to depend on recent consumption. Since pre-reform claimants received higher benefits in the months preceding day 270, they are predicted to have higher subjective needs. They are therefore more likely to perceive scarcity than post-reform claimants when both groups receive the same benefits after day 270.

A higher subjective need makes finding a job more valuable, but does not change the relative value of different jobs. This holds as long as every job pays enough to meet the need.[25] Pre-reform claimants therefore shift attention toward applications and away from

[24] For job-finding rates, I pool the 15-day periods within each window, with one observation per claimant for every period in which they remain unemployed, and include 15-day period fixed effects. For reemployment wages, I include fixed effects for the 15-day period in which the worker finds a job. Standard errors are clustered at the local unemployment insurance office level.

[25] If subjective needs depend on recent benefits, the highest possible need is the pre-reform maximum benefit of 44,460 HUF per month, about a third of median previous earnings in the sample. All observed first-month reemployment earnings exceed this amount, suggesting that the assumption is reasonable in this setting.

evaluating vacancies. The model predicts that they (i) find jobs faster, but (ii) earn less when reemployed, because they consider fewer opportunities and are more likely to receive an offer from a lower-ranked vacancy.

6.3 Results

Figure 9(a) shows the probability of finding a job around day 270. Reproducing the finding in DellaVigna et al. (2017), pre-reform claimants are substantially more likely to find a job once benefits converge. In the first 15 days after day 270, 6.9% of pre-reform claimants still unemployed found a job, compared to 4.0% of post-reform claimants. Over the first 45 days, the job-finding rate is 56% higher for the pre-reform group ($p < 0.001$; Table 1). This difference persists in the following 45 days (42%, $p < 0.001$), while before day 270 it is smaller and not statistically significant.

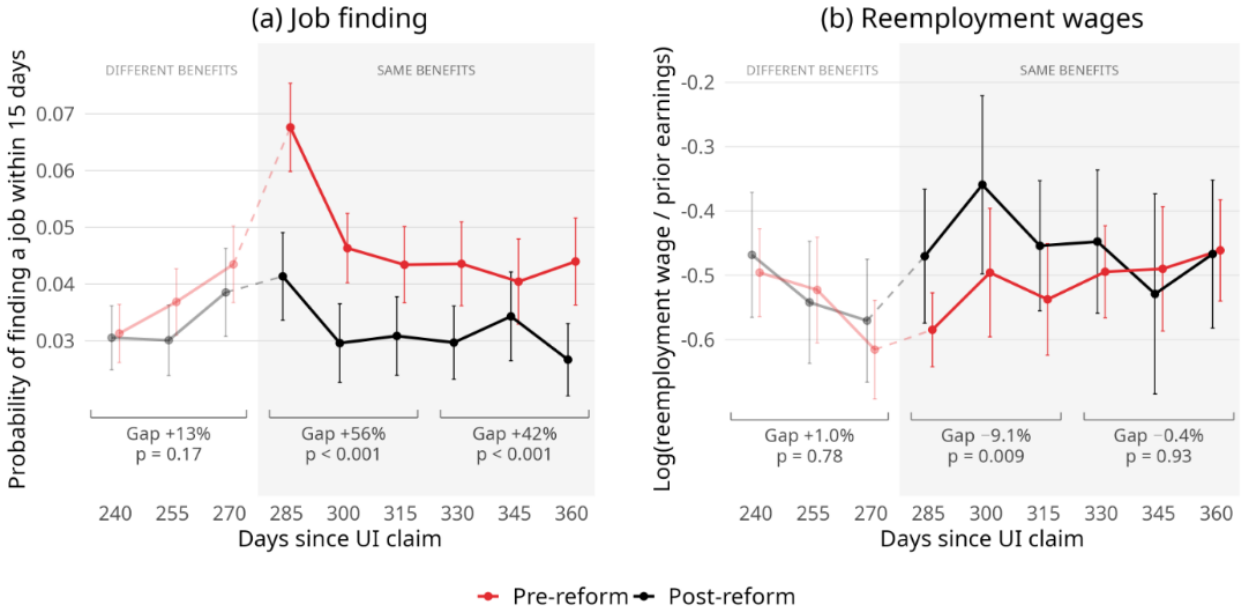


Figure 9: Job-finding rates and reemployment wages around day 270. Each point represents a group's adjusted mean in a 15-day period, identified by its final day. Panel (a) shows the probability of finding a job among claimants still unemployed; panel (b) shows reemployment wages relative to previous earnings. Means and reported differences are adjusted for the controls in Table 1. Bars are 95% CIs, with standard errors clustered by local unemployment insurance office. From day 270 onward (shaded area), both groups receive identical benefits.

These faster exits are accompanied by lower reemployment wages. Figure 9(b) shows that, in the first 45 days after day 270, pre-reform workers find jobs that pay approximately 9% less relative to their previous earnings than those of post-reform workers ($p < 0.01$;

Table 1). No significant wage differences are observed before day 270 or after day 315. Thus, the wage gap is concentrated in the period immediately after the benefit drop.

Table 1: Differences between pre-reform and post-reform claimants around day 270

	Days 226–270	Days 271–315	Days 316–360
Probability of finding a job within 15 days			
Difference	0.0041	0.0189***	0.0125***
	(0.0030)	(0.0033)	(0.0032)
Relative to post-reform	+13%	+56%	+42%
Observations	18,942	16,688	14,768
Log reemployment wage relative to previous earnings			
Difference	0.010	−0.091***	−0.004
	(0.036)	(0.035)	(0.041)
Observations	660	730	540
Change in log monthly earnings (two years after vs. before claim)			
Difference	0.053	−0.076**	0.051
	(0.040)	(0.035)	(0.040)
Observations	581	625	476

Notes: Reported differences are pre-reform minus post-reform estimates within each window, controlling for age, sex, previous occupation and skill level, pre-unemployment earnings, time between job loss and benefit claim, day of the month of the claim, county, and 15-day period fixed effects. Observations are claimant-periods for job finding and reemployed workers for wages and earnings. Standard errors clustered by local unemployment insurance office are in parentheses.
** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The wage gap is robust to additional controls and alternative wage measures. It is also present among workers who return to the same occupation, suggesting that it is not primarily driven by changes in occupation. These results are reported in Appendix D.

There is also suggestive evidence of longer-term earnings differences. To examine this, I compare workers' monthly earnings two years after their benefit claim with their earnings two years before the claim. Among workers reemployed in the first 45 days after day 270, pre-reform workers experienced an approximately 8% larger decline in earnings than post-reform workers ($p = 0.03$; Table 1). No significant differences are observed among workers reemployed in the other two periods.[26] These findings suggest that the differences observed at reemployment may have longer-term consequences for workers' earnings.

[26]Earnings two years later are only observed for workers with earnings in that year, a slightly higher share among pre-reform workers. To account for this difference, I compute bounds following Lee (2009), which suggest an earnings gap of 2–11%. However, the confidence interval around the bounds includes zero. Details are provided in Appendix D.

6.4 Interpretation and Alternative Explanations

The results are in line with the model's predictions: pre-reform workers find jobs faster after day 270, but at lower wages. However, other mechanisms could also explain these differences. In this section, I discuss alternative explanations and examine which features of the data are particularly consistent with the model's mechanism of attention allocation under perceived scarcity.

One possibility is precautionary search. Workers particularly constrained and with limited savings may search more intensively or target lower-paying jobs that are easier to obtain (Acemoglu and Shimer 1999; Chaumont and Shi 2022; Eeckhout and Sepahsalari 2024; Herkenhoff et al. 2024). However, the reform mainly changed the timing of benefits rather than the total amount received. By day 270, post-reform workers had received only about 2.4% more in total, while pre-reform workers had received higher benefits in the most recent months (DellaVigna et al. 2017; Lindner and Reizer 2020). Moreover, both groups face the same future benefits. While this does not rule out differences in savings, it makes an explanation based solely on differences in objective constraints less convincing.

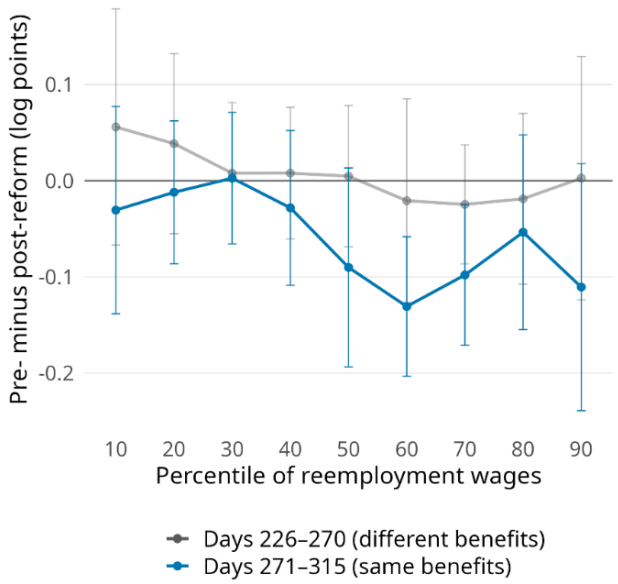


Figure 10: Differences in reemployment wages relative to previous earnings (pre-reform minus post-reform), by percentile of the wage distribution, before day 270 (grey) and during days 271–315 (blue). Negative values indicate lower wages for pre-reform workers. Bars show 95% CIs from a cluster bootstrap by local unemployment insurance office.

Another possible explanation is reference-dependent job search. Although the baseline model in DellaVigna et al. (2017) assumes fixed wages, they also consider an extension with

heterogeneous wages and a reservation-wage decision in the appendix. In this extension, a drop below the reference point can lower the minimum wage workers are willing to accept, generating faster reemployment at lower wages. However, if the wage difference were driven only by lower reservation wages, we would expect pre-reform workers to accept more low-paying jobs that they would otherwise have rejected. Figure 10 suggest a different pattern. Before day 270, the wage distributions of the two groups are similar. In the first 45 days after day 270, the lowest-paid 40% of jobs still pay similarly, but pre-reform workers earn about 9–12% less in the middle of the distribution. [27] Thus, pre-reform workers do not appear to be accepting more of the lowest-paid jobs, but rather finding fewer of the better-paid ones. Appendix D reports a more direct comparison based on the share of jobs paying above different wage thresholds.

One limitation is that reemployment wages are observed only for workers who find a job, and pre-reform workers exit unemployment faster. Differences in who finds a job may therefore also contribute to the wage patterns. Moreover, the model does not explain why the wage gap disappears after day 315 while the difference in job-finding rates persists.

Overall, the field evidence is suggestive rather than conclusive, but complements the experiment by showing that similar patterns arise in actual job-search decisions. The results suggest that the effects of perceived scarcity may extend beyond the laboratory, with potential consequences for reemployment wages and longer-term earnings.

7 Conclusion

This paper formalizes *perceived scarcity* as the state in which one's resources are insufficient relative to one's needs, and studies its cognitive and economic consequences. In the framework, perceived scarcity changes the marginal value of attention, reallocating it toward the scarce domain and away from competing tasks. This simple mechanism can improve decision quality in the scarce domain, reduce productivity elsewhere, or leave performance unchanged, depending on how attention translates into outcomes.

The paper also develops an experimental design to study these effects separately from the direct effects of current resource constraints. By inducing differences in perceived scarcity while holding current resources and incentives fixed, the experiment provides direct evidence of the reallocation of attention predicted by the model. Participants who perceive

[27] The differences are measured at the 50th, 60th, and 70th percentiles, and are statistically significant at the 60th and 70th ($p < 0.01$). Estimates are adjusted for the controls in Table 1, with confidence intervals obtained by resampling local unemployment insurance offices.

scarcity devote more attention to consumption and search more intensively, without making significantly better consumption decisions. Meanwhile, their performance in the competing productivity task declines, and their total earnings are 15% lower.

Finally, I apply the framework to job search, showing how perceived scarcity can affect both the speed and quality of reemployment. Using administrative data from Hungary, I find suggestive evidence consistent with these predictions. Following the unemployment insurance reform, workers with different benefit histories but identical current benefits show substantial differences in job-search outcomes. Those who experienced a larger drop in benefits, and are therefore expected to be more likely to perceive scarcity, find jobs faster but at lower reemployment wages. Although attention is not directly observed, these findings suggest that the mechanism may also matter for job-search behavior, with potential longer-term consequences for earnings.

Taken together, the findings suggest that scarcity matters not only because it constrains what people can do, but also because it changes what they attend to. By reallocating limited attention, perceived scarcity can affect how people make choices, their productivity, and even their job-search outcomes, despite facing identical current resource constraints.

References

Abaluck, Jason and Abi Adams-Prassl (2021). "What do consumers consider before they choose? Identification from asymmetric demand responses". In: *The Quarterly Journal of Economics* 136.3, pp. 1611–1663.

Abel, Andrew B (1990). "Asset Prices under Habit Formation and Catching up with the Joneses". In: *American Economic Review* 80.2, pp. 38–42.

Acemoglu, Daron and Robert Shimer (1999). "Efficient Unemployment Insurance". In: *Journal of Political Economy* 107.5, pp. 893–928.

Alaoui, Larbi and Antonio Penta (2022). "Cost-benefit analysis in reasoning". In: *Journal of Political Economy* 130.4, pp. 881–925.

Alonso, Ricardo, Isabelle Brocas, and Juan D Carrillo (2014). "Resource allocation in the brain". In: *Review of Economic Studies* 81.2, pp. 501–534.

Altmann, Steffen, Anita M. Glenny, Robert Mahlstedt, and Alexander Sebald (2026). "The Direct and Indirect Effects of Online Job Search Advice". In: *Review of Economic Studies*. Forthcoming.

Altmann, Steffen, Andreas Grunewald, and Jonas Radbruch (2022). "Interventions and cognitive spillovers". In: *The Review of Economic Studies* 89.5, pp. 2293–2328.

Arkes, Hal R, Robyn M Dawes, and Caryn Christensen (1986). "Factors influencing the use of a decision rule in a probabilistic task". In: *Organizational Behavior and Human Decision Processes* 37.1, pp. 93–110.

Avoyan, Ala and Andrew Schotter (2020). "Attention in games: An experimental study". In: *European Economic Review* 124, p. 103410.

Belot, Michèle, Philipp Kircher, and Paul Muller (2019). "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice". In: *Review of Economic Studies* 86.4, pp. 1411–1447.

Bolte, Lukas and Collin Raymond (2026). *Emotional inattention*. SSRN Working Paper No. 4776766.

Bonke, Jens and Martin Browning (2009). "The distribution of financial well-being and income within the household". In: *Review of Economics of the Household* 7.1, pp. 31–42.

Bowman, David, Deborah Minehart, and Matthew Rabin (1999). "Loss aversion in a consumption–savings model". In: *Journal of Economic Behavior & Organization* 38.2, pp. 155–178.

Camerer, Colin F and Robin M Hogarth (1999). "The effects of financial incentives in experiments: A review and capital-labor-production framework". In: *Journal of Risk and Uncertainty* 19.1, pp. 7–42.

Camerer, Colin F et al. (2018). "Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015". In: *Nature Human Behaviour* 2.9, pp. 637–644.

Caplin, Andrew, Mark Dean, and Daniel Martin (2011). "Search and satisficing". In: *American Economic Review* 101.7, pp. 2899–2922.

Carroll, Christopher D, Jody Overland, and David N Weil (2000). "Saving and growth with habit formation". In: *American Economic Review* 90.3, pp. 341–355.

Carvalho, Leandro S, Stephan Meier, and Stephanie W Wang (2016). "Poverty and economic decision-making: Evidence from changes in financial resources at payday". In: *American Economic Review* 106.2, pp. 260–284.

Cattaneo, Matias D, Xinwei Ma, Yusufcan Masatlioglu, and Elchin Suleymanov (2020). "A random attention model". In: *Journal of Political Economy* 128.7, pp. 2796–2836.

Celebi, Can, Christine Exley, Soren Harrs, Hannu Kivimaki, Marta Serra-Garcia, and Jeffrey Yusof (2026). *Mission possible: The collection of high-quality data*. CESifo Working Paper.

Chade, Hector and Lones Smith (2006). "Simultaneous Search". In: *Econometrica* 74.5, pp. 1293–1307.

Chaumont, Gaston and Shouyong Shi (2022). "Wealth accumulation, on-the-job search and inequality". In: *Journal of Monetary Economics* 128, pp. 51–71.

Cheremukhin, Anton, Paulina Restrepo-Echavarria, and Antonella Tutino (2020). "Targeted search in matching markets". In: *Journal of Economic Theory* 185.

Clark, Andrew E and Andrew J Oswald (1996). "Satisfaction and comparison income". In: *Journal of Public Economics* 61.3, pp. 359–381.

Consumer Financial Protection Bureau (2015). *Measuring Financial Well-Being: A Guide to Using the CFPB Financial Well-Being Scale*. Tech. rep. Washington, DC: Consumer Financial Protection Bureau.

Dalton, Patricio S, Victor H Gonzalez Jimenez, and Charles N Noussair (2017). "Exposure to poverty and productivity". In: *PLoS ONE* 12.1, e0170231.

De Bruijn, Ernst-Jan and Gerrit Antonides (2022). "Poverty and economic decision making: a review of scarcity theory". In: *Theory and Decision* 92.1, pp. 5–37.

DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F Schmieder (2017). "Reference-dependent job search: Evidence from Hungary". In: *The Quarterly Journal of Economics* 132.4, pp. 1969–2018.

DellaVigna, Stefano and M. Daniele Paserman (2005). "Job Search and Impatience". In: *Journal of Labor Economics* 23.3, pp. 527–588.

Duquennois, Claire (2022). "Fictional money, real costs: Impacts of financial salience on disadvantaged students". In: *American Economic Review* 112.3, pp. 798–826.

Eeckhout, Jan and Alireza Sepahsalari (2024). "The Effect of Wealth on Worker Productivity". In: *Review of Economic Studies* 91, pp. 1584–1633.

European Social Survey (2002). *ESS Round 1 Source Questionnaire*. Centre for Comparative Social Surveys, City University London. Variable HINCFEL.

Eurostat (2011). *Description of EU-SILC Target Variables: Cross-Sectional and Longitudinal, 2011 Operation*. European Commission, Luxembourg. Variables HS120 and HS130.

Fehr, Dietmar, Günther Fink, and B. Kelsey Jack (2022). "Poor and Rational: Decision-Making under Scarcity". In: *Journal of Political Economy* 130.11, pp. 2862–2897.

Gabaix, Xavier (2014). "A sparsity-based model of bounded rationality". In: *The Quarterly Journal of Economics* 129.4, pp. 1661–1710.

Gabaix, Xavier and David Laibson (2005). *Bounded rationality and directed cognition*. Working Paper. Harvard University.

Gali, Jordi (1994). "Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices". In: *Journal of Money, Credit and Banking* 26.1, pp. 1–8.

Geary, Roy C (1950). "A note on "A constant-utility index of the cost of living"". In: *Review of Economic Studies* 18.1, pp. 65–66.

Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler (2021). "The complexity of the consumer problem". In: *Research in Economics* 75.1, pp. 96–103.

Gill, David and Victoria Prowse (2019). "Measuring costly effort using the slider task". In: *Journal of Behavioral and Experimental Finance* 21, pp. 1–9.

Goedhart, Theo, Victor Halberstadt, Arie Kapteyn, and Bernard van Praag (1977). "The poverty line: Concept and measurement". In: *Journal of Human Resources* 12.4, pp. 503–520.

Goerg, Sebastian J, Sebastian Kube, and Jonas Radbruch (2019). "The effectiveness of incentive schemes in the presence of implicit effort costs". In: *Management Science* 65.9, pp. 4063–4078.

Goldin, Jacob and Tatiana Homonoff (2013). "Smoke gets in your eyes: cigarette tax salience and regressivity". In: *American Economic Journal: Economic Policy* 5.1, pp. 302–336.

Haushofer, Johannes and Daniel Salicath (2023). "The psychology of poverty: where do we stand?" In: *Social Philosophy and Policy* 40.1, pp. 150–184.

Heath, Chip and Jack B Soll (1996). "Mental budgeting and consumer decisions". In: *Journal of Consumer Research* 23.1, pp. 40–52.

Heckman, James J. and Bo E. Honoré (1990). "The Empirical Content of the Roy Model". In: *Econometrica* 58.5, pp. 1121–1149.

Herkenhoff, Kyle, Gordon Phillips, and Ethan Cohen-Cole (2024). "How Credit Constraints Impact Job Finding Rates, Sorting, and Aggregate Output". In: *Review of Economic Studies* 91.5, pp. 2832–2877.

Hogarth, Robin M, Brian J Gibbs, Craig R McKenzie, and Margaret A Marquis (1991). "Learning from feedback: Exactingness and incentives". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17.4, pp. 734–752.

Johnson, Eric J, John W Payne, James R Bettman, and David A Schkade (1989). *Monitoring information processing and decisions: The mouselab system*. Tech. rep. Duke University.

Kahneman, Daniel, W Scott Peavler, and Linda Onuska (1968). "Effects of verbalization and incentive on the pupil response to mental activity". In: *Canadian Journal of Psychology* 22.3, pp. 186–196.

Kaur, Supreet, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach (2025). "Do financial concerns make workers less productive?" In: *The Quarterly Journal of Economics* 140.1, pp. 635–689.

Kircher, Philipp (2009). "Efficiency of Simultaneous Search". In: *Journal of Political Economy* 117.5, pp. 861–913.

Kurzban, Robert, Angela Duckworth, Joseph W Kable, and Justus Myers (2013). "An opportunity cost model of subjective effort and task performance". In: *Behavioral and Brain Sciences* 36.6, pp. 661–679.

Lee, David S. (2009). "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects". In: *The Review of Economic Studies* 76.3, pp. 1071–1102.

Lentz, Rasmus, Jonas Maibom, and Espen R. Moen (2026). *Directedness in Search*. Discussion Paper 21073. CEPR.

Libby, Robert and Marlys Gascho Lipe (1992). "Incentives, effort, and the cognitive processes involved in accounting-related judgments". In: *Journal of Accounting Research* 30.2, pp. 249–273.

Lichand, Guilherme and Anandi Mani (2020). *Cognitive droughts*. Working Paper 341. University of Zurich, Department of Economics.

Lindner, Attila and Balázs Reizer (2020). "Front-loading the unemployment benefit: An empirical assessment". In: *American Economic Journal: Applied Economics* 12.3, pp. 140–174.

Loewenstein, George and Zachary Wojtowicz (2025). "The economics of attention". In: *Journal of Economic Literature* 63.3, pp. 1038–1089.

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao (2013). "Poverty impedes cognitive function". In: *Science* 341.6149, pp. 976–980.

Manzini, Paola and Marco Mariotti (2014). "Stochastic choice and consideration sets". In: *Econometrica* 82.3, pp. 1153–1176.

Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y Ozbay (2012). "Revealed attention". In: *American Economic Review* 102.5, pp. 2183–2205.

Menzio, Guido (2007). "A Theory of Partially Directed Search". In: *Journal of Political Economy* 115.5, pp. 748–769.

Mortensen, Dale T. (1986). "Job Search and Labor Market Analysis". In: *Handbook of Labor Economics*. Ed. by Orley Ashenfelter and Richard Layard. Vol. 2. Elsevier, pp. 849–919.

Mullainathan, Sendhil and Eldar Shafir (2013). *Scarcity: Why having too little means so much*. Macmillan.

Murawski, Carsten and Peter Bossaerts (2016). "How humans solve complex problems: The case of the knapsack problem". In: *Scientific Reports* 6.1, p. 34851.

Nekoei, Arash and Andrea Weber (2017). "Does Extending Unemployment Benefits Improve Job Quality?" In: *American Economic Review* 107.2, pp. 527–561.

Ong, Qiyan, Walter Theseira, and Irene YH Ng (2019). "Reducing debt improves psychological functioning and changes decision-making in the poor". In: *Proceedings of the National Academy of Sciences* 116.15, pp. 7244–7249.

Oprea, Ryan (2026). "Complexity and its measurement". In: *Handbook of Experimental Methods in the Social Sciences*. Edward Elgar Publishing, pp. 106–162.

Park, Ye Rang, Yuen Ho, Kristina Hallez, Supreet Kaur, Mahesh Srinivasan, and Jiaying Zhao (2025). "The psychology of poverty: current and future directions". In: *Current Directions in Psychological Science* 34.1, pp. 21–28.

Quispe-Torreblanca, Edika, John Gathergood, George Loewenstein, and Neil Stewart (2026). "Attention utility: Evidence from individual investors". In: *Review of Economic Studies* 93.1, pp. 664–696.

Reutskaja, Elena, Rosemarie Nagel, Colin F Camerer, and Antonio Rangel (2011). "Search dynamics in consumer choice under time pressure: An eye-tracking study". In: *American Economic Review* 101.2, pp. 900–926.

Schmieder, Johannes F., Till von Wachter, and Stefan Bender (2016). "The Effect of Unemployment Benefits and Nonemployment Durations on Wages". In: *American Economic Review* 106.3, pp. 739–777.

Sergeyev, Dmitriy, Chen Lian, and Yuriy Gorodnichenko (2025). "The economics of financial stress". In: *Review of Economic Studies* 92.5, pp. 3403–3437.

Sewell, Patrick (2026). *Attention Allocation under Scarcity: Theory and Evidence*. AEA RCT Registry. June 29. DOI: 10.1257/rct.18956-1.0.

Shah, Anuj K, Sendhil Mullainathan, and Eldar Shafir (2012). "Some consequences of having too little". In: *Science* 338.6107, pp. 682–685.

— (2019). "An exercise in self-replication: Replicating Shah, Mullainathan, and Shafir (2012)". In: *Journal of Economic Psychology* 75, p. 102127.

Shah, Anuj K, Eldar Shafir, and Sendhil Mullainathan (2015). "Scarcity frames value". In: *Psychological Science* 26.4, pp. 402–412.

Simon, Herbert A (1955). "A behavioral model of rational choice". In: *The Quarterly Journal of Economics* 69.1, pp. 99–118.

— (1967). "Motivational and emotional controls of cognition." In: *Psychological Review* 74.1, pp. 29–39.

Sims, Christopher A (2003). "Implications of rational inattention". In: *Journal of Monetary Economics* 50.3, pp. 665–690.

Smith, Adam (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: W. Strahan and T. Cadell.

Spiller, Stephen A (2011). "Opportunity cost consideration". In: *Journal of Consumer Research* 38.4, pp. 595–610.

Stigler, George J (1961). "The Economics of Information". In: *Journal of Political Economy* 69.3, pp. 213–225.

Stone, Richard (1954). "Linear expenditure systems and demand analysis: An application to the pattern of British demand". In: *The Economic Journal* 64.255, pp. 511–527.

Thaler, Richard H (1999). "Mental accounting matters". In: *Journal of Behavioral Decision Making* 12.3, pp. 183–206.

tho Pesch, Fiona, Antonia F Langenhoff, and Mahesh Srinivasan (2026). "Revisiting evidence for an attentional tunneling effect caused by scarcity". In: *Nature Communications* 17.1, p. 10233.

Weitzman, Martin L. (1979). "Optimal Search for the Best Alternative". In: *Econometrica* 47.3, pp. 641–654.

Wu, Liangjie (2024). "Partially directed search in the labor market". In: *Journal of Economic Theory* 220, p. 105858.

Zhao, Jiaying and Brandon M Tomm (2017). "Attentional trade-offs under resource scarcity". In: *International Conference on Augmented Cognition*. Springer, pp. 78–97.

Zubrickas, Robertas (2026). "Reference dependence and lottery participation". In: *Economic Inquiry* 64.3, pp. 886–894.

Appendix A Proofs

Properties of the consumption utility function $u(\cdot)$: (i) continuous, (ii) monotonic, (iii) strictly quasi-concave.

Proof: Recall that $u(x;\eta) = v(c(x);\eta)$, where $c(x) = \left[\sum_i \alpha_i x_i^r\right]^{\frac{1}{r}}$ with $\alpha_i > 0$ and $r < 1$.

(i) The aggregator $c(\cdot)$ is continuous, and $v(\cdot;\eta)$ is continuous since $\Delta(\eta)$ connects its two pieces at η. Hence, $u(x;\eta)$ is continuous.

(ii) Since $\alpha_i > 0$, $c(\cdot)$ is increasing in each good, and since $\frac{dv_1}{dc} > 0$ and $\frac{dv_2}{dc} > 0$, $v(\cdot;\eta)$ is strictly increasing in the level of consumption. Hence, $u(x;\eta)$ is increasing in each good.

(iii) Since $v(\cdot;\eta)$ is strictly increasing, for any utility level k the upper contour set $\{x : u(x;\eta) \geq k\}$ coincides with $\{x : c(x) \geq \bar{c}\}$ for some level of consumption $\bar{c}$. For $r < 1$, the CES aggregator is concave, so these sets are convex. Hence, $u(x;\eta)$ is quasi-concave, and strictly so since $c(\cdot)$ is strictly quasi-concave for $r < 1$. ■

Lemma 3.1 For every a, there exists a threshold $|T_0|$, which may depend on a, such that $\frac{\partial \mu(T|a)}{\partial a} \leq 0 \; \forall |T| \leq |T_0|$, and $\frac{\partial \mu(T|a)}{\partial a} > 0 \; \forall |T| > |T_0|$

Proof: Since μ is a probability distribution, we know $\sum_{T \subseteq S} \mu(T|a) = 1$, so **(i)** $\sum_{T \subseteq S} \frac{\partial \mu(T|a)}{\partial a} = 0$. By supermodularity (Assumption 3), for any $a' > a$ and $|T'| > |T|$, we have $\mu(T'|a') - \mu(T'|a) \geq \mu(T|a') - \mu(T|a)$. Dividing both sides by $a' - a > 0$ and letting $a' \to a$ yields **(ii)** $\frac{\partial \mu(T'|a)}{\partial a} \geq \frac{\partial \mu(T|a)}{\partial a}$ whenever $|T'| > |T|$. That is, the effect of attention on the probability of a consideration set is weakly increasing in its size.

By Assumption 1, $\frac{\partial \mu(T|a)}{\partial a} < 0$ for $|T| = 1$, so by **(i)** there must be some size for which $\frac{\partial \mu(T|a)}{\partial a} > 0$. Let $|T_0|$ be the largest size for which $\frac{\partial \mu(T|a)}{\partial a} \leq 0$, which exists since singletons satisfy this condition. By **(ii)**, every size $|T| \leq |T_0|$ satisfies $\frac{\partial \mu(T|a)}{\partial a} \leq \frac{\partial \mu(T_0|a)}{\partial a} \leq 0$, and by the definition of $|T_0|$, every size $|T| > |T_0|$ satisfies $\frac{\partial \mu(T|a)}{\partial a} > 0$. ■

Lemma 3.2: The higher the allocated attention to consumption, the greater the expected consumption utility.

Proof: The proof proceeds in 2 steps. First I show that expected utility weakly increases when the agent considers larger choice sets, and second I prove that $\frac{\partial \mathbb{E}[u(x(a))]}{\partial a} > 0$.

1. A larger consideration set weakly increases utility.

Fix the budget set $B(p,b)$ and η. Take two consideration sets $T_L \subset T_H$, so $|T_L| < |T_H|$. Denote by x^L and x^H the optimal bundles given T_L and T_H, respectively. Since $B(p,b)$ is fixed and $T_L \subset T_H$, any bundle feasible under T_L is also feasible under T_H. Therefore, $u(x^H;\eta) \geq u(x^L;\eta)$, since the agent can always choose the same bundle when considering more goods.

To extend this result to consideration sets that are not nested, denote by V_m the expected utility conditional on considering m goods:

$$V_m = \frac{1}{\binom{|S|}{m}} \sum_{\{T \subseteq S : |T| = m\}} u(x_T;\eta).$$

Since $\mu(T|a) = \mu(|T|,a)$, all consideration sets of the same size are equally likely. Consider a randomly drawn set of size m and add one randomly chosen good from the remaining goods. The resulting set of size $m+1$ is also uniformly distributed and contains the original set. Since a larger nested consideration set weakly increases utility, we obtain $V_{m+1} \geq V_m$.

2. Prove that $\frac{\partial \mathbb{E}[u(x(a))]}{\partial a} > 0$.

Since sets of the same size are equally likely at every level of attention, the utility of each set can be replaced by the average utility $V_{|T|}$ of sets of its size:

$$\frac{\partial \mathbb{E}[u(x(a))]}{\partial a} = \sum_{T \subseteq S} \frac{\partial \mu(T|a)}{\partial a} u(x_T;\eta) = \sum_{T \subseteq S} \frac{\partial \mu(T|a)}{\partial a} V_{|T|}.$$

By **(i)** in the proof of Lemma 3.1, $\sum_{T \subseteq S} \frac{\partial \mu(T|a)}{\partial a} = 0$, so subtracting $V_{|T_0|}$ from every term does not change the sum:

$$\frac{\partial \mathbb{E}[u(x(a))]}{\partial a} = \sum_{T \subseteq S} \frac{\partial \mu(T|a)}{\partial a} \left(V_{|T|} - V_{|T_0|}\right).$$

By Lemma 3.1, $\frac{\partial \mu(T|a)}{\partial a} > 0$ for $|T| > |T_0|$, where $V_{|T|} - V_{|T_0|} \geq 0$ since V_m is weakly increasing in m by the first step of the proof, and $\frac{\partial \mu(T|a)}{\partial a} \leq 0$ for $|T| \leq |T_0|$, where $V_{|T|} - V_{|T_0|} \leq 0$. Therefore, no term is negative: attention shifts probability from small consideration sets to large ones, which yield higher utility on average.

Finally, the derivative is strictly positive. Intuitively, attention raises the probability of considering all goods, and an agent who considers fewer goods may miss a good of the optimal bundle and choose a strictly worse one. Formally, as shown in the proof of Lemma 3.1, some set sizes become more likely with attention, so $|T_0| < |S|$ and $\frac{\partial \mu(S|a)}{\partial a} > 0$. With all goods considered, the agent chooses the unique optimal bundle x^*, so $V_{|S|} = u(x^*;\eta)$.

Since $|T_0| < |S|$, some sets of size $|T_0|$ miss a good of x^*, and their best bundle is therefore strictly worse than x^*. Hence $V_{|T_0|} < V_{|S|}$, so the term for $T = S$ is strictly positive and $\frac{\partial \mathbb{E}[u(x(a))]}{\partial a} > 0$. ■

Lemma 3.3: The agent will perceive scarcity if $b < b_\eta$. Moreover, when he considers all goods, he perceives scarcity if and only if $b < b_\eta$.

Proof: ($\Leftarrow$) Suppose not. We have $b < b_\eta$ and $c(x) \geq \eta$, where $x = (x_1, x_2, \ldots, x_n) \in B(p, b)$

$\implies$ $b \geq p \cdot x \geq b_\eta$, where the last inequality follows from the definition of b_η. This contradicts with $b < b_\eta$.

Therefore $c(x) < \eta$.

($\Rightarrow$) Suppose that $b \geq b_\eta$, and let $\hat{x}$ be a bundle that attains b_η, which exists since $c(\cdot)$ is continuous and prices are positive. Then $c(\hat{x}) \geq \eta$ and $p \cdot \hat{x} = b_\eta \leq b$, so $\hat{x} \in B(p, b)$. Since utility is increasing in the level of consumption, the bundle x^* chosen when all goods are considered satisfies $c(x^*) \geq c(\hat{x}) \geq \eta$. Hence, the agent does not perceive scarcity. ■

Lemma 3.4: Holding objective resources fixed, the marginal value of attention to consumption is greater for an individual who perceives scarcity than for an otherwise identical individual who does not, at every level of attention. Formally, if $b < b_{\eta^H}$ and $b > b_{\eta^L}$, then $\frac{\partial \mathbb{E}[u(x(a);\eta^H)]}{\partial a} > \frac{\partial \mathbb{E}[u(x(a);\eta^L)]}{\partial a}$ for every $a > 0$.

Proof: Fix the budget $b > 0$, prices p, and the attention rule μ, and compare two need levels: $\eta^H > \eta^L$, such that one perceives scarcity ($b < b_{\eta^H}$) and the other does not ($b > b_{\eta^L}$). Resources are identical in both cases, and only the need differs.

Step 1: The need does not affect which bundle is chosen. For a consideration set T, let c_T denote the highest level of consumption that the agent can afford using only the goods in T. Since $v(\cdot; \eta)$ is strictly increasing, the best bundle in T is the one that maximizes $c(x)$, regardless of η, so that $u(x_T; \eta) = v(c_T; \eta)$. Moreover, c_T is weakly increasing as goods are added to T.

Step 2: A higher need makes additional consumption more valuable. Define the scarcity premium of a level of consumption c as $g(c) := v(c; \eta^H) - v(c; \eta^L)$, and let $h := v_1 - v_2$, which is strictly increasing since $\frac{dv_1}{dc} > \frac{dv_2}{dc}$. By the definition of the kinked utility,

$$g(c) = \begin{cases} 0 & \text{if } c < \eta^L, \\ h(c) - h(\eta^L) & \text{if } \eta^L \leq c < \eta^H, \\ h(\eta^H) - h(\eta^L) & \text{if } c \geq \eta^H. \end{cases}$$

Thus, g is weakly increasing, and strictly increasing between η^L and η^H: raising the need increases the value of consumption precisely in the range where consumption falls short of the higher need.

Step 3: Attention is more valuable under perceived scarcity. As in the proof of Lemma 3.2, let $V_m(\eta)$ denote the average of $v(c_T;\eta)$ over consideration sets of size m, so that

$$\frac{\partial \mathbb{E}[u(x(a);\eta)]}{\partial a} = \sum_{T\subseteq S} \frac{\partial \mu(T|a)}{\partial a} V_{|T|}(\eta).$$

Since μ does not depend on η, the difference in the marginal value of attention between the two need levels is

$$\left.\frac{\partial \mathbb{E}[u(x(a);\eta)]}{\partial a}\right|_{\eta^H} - \left.\frac{\partial \mathbb{E}[u(x(a);\eta)]}{\partial a}\right|_{\eta^L} = \sum_{T\subseteq S} \frac{\partial \mu(T|a)}{\partial a} D_{|T|}, \qquad D_m := V_m(\eta^H) - V_m(\eta^L),$$

where D_m is the average scarcity premium $g(c_T)$ over consideration sets of size m. By Steps 1 and 2, $g(c_T)$ is weakly increasing as goods are added to T, so the random-addition argument in the proof of Lemma 3.2 implies that D_m is weakly increasing in m. The rest of the proof mirrors that of Lemma 3.2, with D_m in place of V_m. Since $\sum_{T\subseteq S} \frac{\partial \mu(T|a)}{\partial a} = 0$,

$$\sum_{T\subseteq S} \frac{\partial \mu(T|a)}{\partial a} D_{|T|} = \sum_{T\subseteq S} \frac{\partial \mu(T|a)}{\partial a} \left(D_{|T|} - D_{|T_0|}\right) \geq 0,$$

since, by Lemma 3.1, attention raises the probability of every set larger than $|T_0|$, for which $D_{|T|} \geq D_{|T_0|}$, and weakly lowers that of every other set, for which $D_{|T|} \leq D_{|T_0|}$.

Finally, the inequality is strict. Intuitively, even when he considers all goods, the scarce agent falls short of his need, so missing a good of his optimal bundle costs him strictly more than it costs an agent who does not perceive scarcity. Formally, with all goods considered, consumption meets the low need but not the high one, $\eta^L < c_S < \eta^H$, so any shortfall from c_S strictly lowers the scarcity premium (Step 2). Since $|T_0| < |S|$, some sets of size $|T_0|$ miss a good of the unique optimal bundle x^*, and therefore yield less consumption than c_S. Hence $D_{|T_0|} < g(c_S) = D_{|S|}$, and since $\frac{\partial \mu(S|a)}{\partial a} > 0$, the term for $T = S$ is strictly positive. ■

Proposition 1: Consider two otherwise identical individuals with the same budget b, but different subjective needs η^H and η^L, such that $b < b_{\eta^H}$ and $b > b_{\eta^L}$. Then, for any optimal allocation of attention:

$$a_x^*(b,\eta^H) > a_x^*(b,\eta^L) \text{ and } a_y^*(b,\eta^H) < a_y^*(b,\eta^L).$$

Therefore, the individual perceiving scarcity allocates more attention to consumption at the expense of reduced attention in the productive task.

Proof: By Lemma 3.3, an agent with $b < b_\eta$ perceives scarcity, whereas an agent with $b > b_\eta$ does not when he considers all goods. I therefore compare two agents with the same budget b, prices, attention technology and production function, but different needs: η^H, with $b < b_{\eta^H}$, and η^L, with $b > b_{\eta^L}$.

Since both $\mathbb{E}[u(x(a);\eta)]$ (Lemma 3.2) and $\omega(y(a_y))$ are increasing in attention, the attention constraint binds, so that $a_y = (A - k_x a_x)/k_y$. The problem therefore reduces to choosing a_x to maximize expected utility, which, with some abuse of notation, I denote by $U(a_x;\eta)$:

$$U(a_x;\eta) = \mathbb{E}[u(x(a_x);\eta)] + \omega\left(y\left(\tfrac{A-k_x a_x}{k_y}\right)\right), \qquad \frac{\partial U(a_x;\eta)}{\partial a_x} = \frac{\partial \mathbb{E}[u(x(a_x);\eta)]}{\partial a_x} - \frac{k_x}{k_y}\frac{\partial \omega(y(a_y))}{\partial a_y}.$$

The second term of the derivative does not depend on η. Hence, by Lemma 3.4, $\frac{\partial U(a_x;\eta^H)}{\partial a_x} > \frac{\partial U(a_x;\eta^L)}{\partial a_x}$ at every a_x: perceived scarcity raises the marginal value of attention to consumption at every level of attention. Equivalently, the difference $U(a_x;\eta^H) - U(a_x;\eta^L)$ is strictly increasing in a_x.

Let $a_H = a_x^*(b,\eta^H)$ and $a_L = a_x^*(b,\eta^L)$ be any optimal attention to consumption under η^H and η^L, respectively. Intuitively, since attention to consumption is more valuable at every level under η^H, the agent perceiving scarcity gains more from any increase in it, and so never chooses less of it than the agent who does not. Formally, suppose, by contradiction, that $a_H < a_L$. Since a_L is optimal under η^L, raising attention from a_H to a_L weakly increases utility under η^L. Because $U(a_x;\eta^H) - U(a_x;\eta^L)$ is strictly increasing in a_x, the same increase raises utility by strictly more under η^H:

$$U(a_L;\eta^H) - U(a_H;\eta^H) > U(a_L;\eta^L) - U(a_H;\eta^L) \geq 0.$$

Thus, a_L yields strictly higher utility than a_H under η^H, contradicting the optimality of a_H. Hence, $a_H \geq a_L$.

Moreover, $a_H \neq a_L$. By the Inada conditions, a_L is interior, so $\frac{\partial U(a_L;\eta^L)}{\partial a_x} = 0$ and hence $\frac{\partial U(a_L;\eta^H)}{\partial a_x} > 0$: at the allocation that is optimal without scarcity, an agent perceiving scarcity would gain by shifting attention from the productive task to consumption, so a_L is not optimal for him. Therefore, $a_x^*(b,\eta^H) = a_H > a_L = a_x^*(b,\eta^L)$ and, since the attention constraint binds, $a_y^*(b,\eta^H) < a_y^*(b,\eta^L)$. ■

Remark 1: Without a subjective need and with substitutable goods, a lower

budget increases attention to consumption if $\varepsilon(c) > 1$ for all c, decreases it if $\varepsilon(c) < 1$ for all c, and leaves it unchanged under logarithmic utility.

Proof: Without a subjective need, so that $u(x) = v(c(x))$ with v smooth, strictly increasing, and concave, a lower budget has two opposing effects on the value of attention to consumption. On the one hand, lower consumption increases marginal utility, making improvements in the chosen bundle more valuable. On the other hand, a smaller budget reduces the consumption gains attainable from making better choices, making the consumption task less attractive. Which effect dominates depends on the elasticity of marginal utility,

$$\varepsilon(c) := -\frac{d \log v'(c)}{d \log c} = -\frac{c\, v''(c)}{v'(c)},$$

which measures how rapidly marginal utility declines with consumption.

Formally, suppose that goods are substitutes ($r \in (0,1)$). Since $c(\cdot)$ is homogeneous of degree one, the highest level of consumption attainable with the goods in T is $c_T = \kappa_T b$, where κ_T denotes consumption per unit of budget and is weakly increasing as goods are added to T. Expected utility from the consumption task can therefore be written as

$$V(a,b) = \sum_{T \subseteq S} \mu(T \mid a)\, v(\kappa_T b).$$

Differentiating with respect to attention and the budget gives

$$\frac{\partial^2 V(a,b)}{\partial a\, \partial b} = \sum_{T \subseteq S} \frac{\partial \mu(T \mid a)}{\partial a}\, \kappa_T v'(\kappa_T b).$$

The sign of this cross-partial depends on how $\kappa_T v'(\kappa_T b)$, which captures the utility gain from an additional unit of b given consideration set T, varies with κ_T. In particular,

$$\frac{d}{d\kappa}\left[\kappa v'(\kappa b)\right] = v'(\kappa b) + \kappa b\, v''(\kappa b) = v'(\kappa b)\left[1 - \varepsilon(\kappa b)\right].$$

Thus, if $\varepsilon(c) < 1$ for all c, the expression is strictly increasing in κ. By the argument in the proof of Lemma 3.2, the cross-partial is positive, so a higher budget increases the marginal value of attention to consumption. Conversely, if $\varepsilon(c) > 1$ for all c, the cross-partial is negative, and a lower budget increases the marginal value of attention. With logarithmic utility, $\varepsilon(c) = 1$ and $\kappa_T v'(\kappa_T b) = 1/b$ for every T, so the cross-partial is zero.

Since the marginal cost of moving attention away from production does not depend on b, the comparative statics follow as in the proof of Proposition 1, with b in place of

η. In particular, under CRRA utility, $\varepsilon(c) = \sigma$: a lower budget draws attention toward consumption if $\sigma > 1$, away from consumption if $\sigma < 1$, and leaves attention unchanged if $\sigma = 1$. ■

Corollary 1: Under perceived scarcity, (i) decision quality in the consumption task improves if actual performance in that task is increasing in attention, and (ii) productivity in the productive task falls if actual output in that task is increasing in attention.

Proof: Let $Q_x(a_x)$ and $Q_y(a_y)$ denote the actual performance in the consumption and productive tasks as functions of the attention allocated to each, which need not coincide with the perceived mappings that the agent uses to allocate attention. By Proposition 1, $a_x^*(b, \eta^H) > a_x^*(b, \eta^L)$ and $a_y^*(b, \eta^H) < a_y^*(b, \eta^L)$. (i) If Q_x is increasing on $[a_x^*(b, \eta^L), a_x^*(b, \eta^H)]$, then $Q_x(a_x^*(b, \eta^H)) > Q_x(a_x^*(b, \eta^L))$, so decision quality improves. (ii) If Q_y is increasing on $[a_y^*(b, \eta^H), a_y^*(b, \eta^L)]$, then $Q_y(a_y^*(b, \eta^H)) < Q_y(a_y^*(b, \eta^L))$, so productivity falls. Each part depends only on the condition for its own task. In the special case where the actual mappings coincide with the perceived ones, both conditions hold: expected consumption utility is increasing in attention by Lemma 3.2, and $y(\cdot)$ is strictly increasing by assumption. ■

Appendix B Job Search with Limited Attention

Job search is a complex task that requires looking for jobs, evaluating vacancies to assess how well they fit, contacting potential employers, and writing applications. In this extension, I apply the model to job search and focus on two uses of attention: sending applications, which raises the chance of finding a job, and evaluating job opportunities, which helps him find jobs that better match his skills and therefore pay more. The extension formalizes the interpretation of the Hungarian evidence in Section 6.

Related literature. Job search models usually study one of these margins at a time. Models of search intensity study how much effort workers put into search or how many applications they send; when they add a second margin, it is usually how selective workers are, through a reservation wage or the wage of the jobs they target (DellaVigna and Paserman 2005; Kircher 2009; DellaVigna et al. 2017; Nekoei and Weber 2017). Models of partially directed search study how precisely workers can target the jobs that suit them, but keep the intensity of search fixed (Menzio 2007; Cheremukhin et al. 2020; Wu 2024; Lentz et al. 2026). The formulation below builds on models of simultaneous search, in which workers choose a portfolio of applications (Chade and L. Smith 2006; Kircher 2009), and on evidence that what job seekers know about available jobs shapes where they apply (Belot et al. 2019; Altmann et al. 2026). In these models, what workers know about available jobs is given. Here, it depends on how much attention workers devote to evaluating job opportunities, and this attention competes with attention to applications.

Setup. I keep the setup as close as possible to a standard job search model with hand-to-mouth workers, as in DellaVigna et al. (2017), and add the attention constraint and the subjective need of the main model. Time is discrete. Consider a worker who is unemployed at the start of the period right after his benefits drop, which corresponds to day 270 in the Hungarian setting. From then on, he receives a benefit b while unemployed, and b is the same for workers who entered unemployment before and after the reform. He consumes his income in every period, so his consumption is b while unemployed and w while employed at wage w. His flow utility is the kink utility $v(c;\eta)$ of the main model, and he discounts the future at rate $\beta \in (0,1)$.

Need. The need η is formed by past consumption, as in Section 3.5: $\eta_t = \rho\,\eta_{t-1} + (1-\rho)\,c_{t-1}$, with $\rho \in (0,1)$. The following lemma shows that the reform leaves the pre-reform group with a higher η at the moment benefits drop.

Lemma B.1. *Suppose both groups start with the same need η_0. The pre-reform group (B) receives $\bar{b}$ in each of the first T periods. The post-reform group (A) receives $b_H > \bar{b}$ in the first $T_1 < T$ periods and $b_L < \bar{b}$ in the remaining $T - T_1$ periods, with the same total: $T_1(b_H - \bar{b}) = (T - T_1)(\bar{b} - b_L)$. Then for any $\rho \in (0, 1)$, $\eta_T^B > \eta_T^A$.*

Proof. Iterating the law of motion gives $\eta_T = \rho^T \eta_0 + (1-\rho) \sum_{s=0}^{T-1} \rho^{T-1-s} c_s$, so $\eta_T^B - \eta_T^A = (1-\rho) \sum_{s=0}^{T-1} \rho^{T-1-s} d_s$, where $d_s = c_s^B - c_s^A$ is negative for $s < T_1$, positive for $s \geq T_1$, and $\sum_s d_s = 0$. Let $\bar{\omega} = \rho^{T-1-T_1}$. Since $\sum_s d_s = 0$, $\sum_s \rho^{T-1-s} d_s = \sum_s (\rho^{T-1-s} - \bar{\omega})\, d_s$. For $s < T_1$, both factors are negative; for $s \geq T_1$, both are non-negative. Hence the sum is strictly positive. □

In words, recent benefits weigh more in η than older ones, and the pre-reform group received the higher benefit in the most recent periods. In Hungary, the total was slightly higher in the new system. The result does not depend on exact equality: with the benefit amounts in Lindner and Reizer (2020), η^B is still higher than η^A at day 270 as long as η adjusts to new consumption with a half-life shorter than about 20 months. I assume that the drop at day 270 brings the pre-reform group below its need, $b < \eta^B$, which is natural since the unemployment assistance paid after day 270 is much lower than the benefit before it. I also assume that every job pays enough to meet the need, $\eta^B \leq \underline{w}$, where $\underline{w}$ is the lowest possible wage. This is natural in this setting: η is formed by recent consumption, and an unemployed worker consumes his benefit, so over a long unemployment spell η moves toward the benefit level. Since benefits replace only part of previous earnings, a job can pay considerably less than the previous one and still pay at least η.[28] I make no assumption on whether the post-reform group perceives scarcity.

Vacancies and the consideration set. The market offers a finite set of vacancies J, and vacancy j pays wage w_j. Wages differ across vacancies partly because jobs differ in how well they match the worker's skills: a vacancy that makes good use of his experience pays more than a generic one. Since his previous wage reflects these skills, the empirical analysis measures reemployment wages relative to previous earnings. The worker does not know every vacancy. As in the consumption problem of the main model, he only considers a subset $T \subseteq J$, drawn according to the attention rule $\mu(T \mid a_m)$, where a_m is the attention he devotes to evaluating job opportunities: looking through vacancies and learning which ones suit him. A vacancy in the consideration set is one whose wage he knows, as with an inspected box in Weitzman (1979). The attention rule satisfies the same assumptions as in

[28] The results extend to the case in which some job types pay less than η. For these types, a higher η also raises the value of finding a better type, which pushes toward evaluation. This effect is small when few jobs pay less than η.

the main model, and in particular it depends only on the size of the set, so that all sets of a given size are equally likely. The only difference is that even without evaluating anything, the worker knows some vacancies, for instance those listed at the employment office or passed on by acquaintances:

Assumption 4 (Baseline set)**.** $\mu(T \mid a_m) = 0$ *for every* T *with* $|T| < \bar{n}$*, where* $\bar{n}$ *is the largest number of applications the worker can send in a period (see below), and Assumption 1 and the Inada condition of the main model apply to sets of size* $\bar{n}$ *instead of single goods.*

Since all sets of a given size are equally likely, the vacancies a worker knows without evaluation are a random draw from the market. He still applies to the best of them, but the more attention he devotes to evaluation, the larger the set he chooses from, and the more directed his search becomes.

Applications. Attention to applications, a_i, covers writing applications and contacting employers. Let $N(a_i)$ denote the number of applications the worker sends given this attention. Preparing an application takes attention, and how much varies: some applications take a few minutes, while others require a cover letter or contacting the employer. I assume that each application requires an amount of attention that is exponentially distributed with mean $1/\kappa$, independently across applications.[29] The worker works down his list, starting with the best-paid vacancy he knows, and completes as many applications as his attention allows, up to $\bar{n}$ per period. The cap $\bar{n}$ can be thought of as the largest number of applications he could complete in a period given his attention capacity A, and the baseline set ensures that he never runs out of vacancies. Then $N(a_i)$ follows a Poisson distribution with mean κa_i. Each application leads to an offer with probability $p \in (0, 1)$, independently of the other applications and of the vacancy,[30] and he accepts the best offer he receives. He therefore finds a job with probability

$$\lambda(a_i) = 1 - \mathbb{E}\Big[(1-p)^{N(a_i)}\Big]$$

which is strictly increasing in a_i. Without the cap, offers would arrive according to a Poisson process with rate $p\kappa a_i$, as in standard models of search intensity (Mortensen 1986). Thus,

[29] The exponential distribution is assumed for tractability. The results do not depend on it and hold for a broad class of distributions of preparation times.

[30] The same offer probability for all vacancies is a simplifying benchmark. Higher-paying vacancies need not be harder to get: wage differences may reflect how well a job matches the worker's skills rather than differences in requirements or competition, and employers are, if anything, more willing to hire applicants who fit well. If they were easier to get, evaluation would also raise the probability of finding a job, and the results below continue to hold as long as applications matter more for finding a job than evaluation does. If better-paying vacancies were harder to get, as in directed search models where higher wages attract more applicants (Nekoei and Weber 2017), a worker who perceives scarcity would also choose to apply to lower-paying vacancies. This would reinforce the wage loss, but through a channel other than attention.

when $\bar{n}$ is large enough that the cap rarely binds, $\lambda(a_i)$ is close to $1 - e^{-p\kappa a_i}$.

If $w_{(1)} \geq w_{(2)} \geq \ldots$ denote the wages of the vacancies in his set, ranked from highest to lowest, the worker ends up in the j-th best vacancy if he sends at least j applications, receives no offers from the $j-1$ better vacancies, and receives an offer from the j-th:

$$q_j(a_i) = \Pr\Big(N(a_i) \geq j\Big)\,(1-p)^{j-1}\,p, \qquad j \leq \bar{n}.$$

Summing these probabilities over all vacancies gives $\lambda(a_i)$, the probability of finding a job. As in the main model, the two uses of attention share the constraint $k_i a_i + k_m a_m \leq A$.

Values. Once employed, the worker keeps his job, so the value of a job paying w is $W(w;\eta) = v(w;\eta)/(1-\beta)$. As in models where workers choose which jobs to apply to (Chade and L. Smith 2006; Kircher 2009; Nekoei and Weber 2017), I assume that workers apply only to jobs they would be willing to accept and choose the best offer they receive. This abstracts from a separate reservation-wage decision, as in the baseline model of DellaVigna et al. (2017).

In the spirit of models of simultaneous search (Stigler 1961; Chade and L. Smith 2006; Kircher 2009), I focus on a single search round immediately after the benefit drop. A worker who does not find a job remains unemployed, with value $U(\eta) = v(b;\eta)/(1-\beta)$. This simplifies the analysis, but the results also extend to a dynamic search setting.[31]

With $b < \eta \leq \underline{w}$, the gain from finding a job paying w is

$$W(w;\eta) - U(\eta) = \frac{C(\eta) + v_2(w)}{1-\beta}, \qquad C(\eta) = v_1(\eta) - v_1(b) - v_2(\eta). \tag{8}$$

The gain has two parts. The first, $C(\eta)$, captures the value of moving from unemployment benefits up to η, where marginal utility is high. It is the same for every job and is strictly increasing in η. The second, $v_2(w)$, captures the value of a better-paying job and does not depend on the need. When $b \geq \eta$, the need no longer affects the gain from employment.

The worker solves

$$\begin{aligned} &\max_{a_i, a_m}\; U(\eta) + \frac{1}{1-\beta}\Big[\lambda(a_i)\,C(\eta) + R(a_i, a_m)\Big] \\ &\text{s.t.}\quad k_i a_i + k_m a_m \leq A, \end{aligned}$$

[31] In a dynamic model, η still affects the value of finding a job, but not the relative value of better-paying jobs. If η also adjusts to the lower benefit during unemployment, the two groups' needs and attention allocations converge, consistent with the wage gap closing by day 315 in Section 6.

where $R(a_i, a_m) = \sum_{T \subseteq J} \mu(T \mid a_m) \sum_j q_j(a_i)\, v_2\Big(w_{(j)}(T)\Big)$ is the expected value of the job he ends up in. This problem mirrors the main model: the worker chooses the best options in an attention-dependent consideration set, and that attention competes with attention to another task. Evaluation raises R, by the same argument as in Lemma 3.2: adding a vacancy to a set weakly raises the wage of its j-th best vacancy, for every j, and more attention makes larger sets more likely. More attention to applications raises the probability of finding a job, and the value of this, $\lambda'(a_i)\, C(\eta)/(1-\beta)$ plus a term that does not depend on η, rises with η. The value of evaluation does not depend on η. Perceived scarcity therefore makes applications more valuable relative to evaluation.

Proposition 2 (Job search under perceived scarcity)**.** *Let $\eta^B > \eta^A$ and $b < \eta^B \leq \underline{w}$. In the period after benefits drop, the pre-reform worker allocates more attention to applications and less to evaluating job opportunities than the post-reform worker: $a_i^B > a_i^A$ and $a_m^B < a_m^A$.*

Proof. Since λ and R are increasing in both uses of attention, the constraint binds, so $a_m = (A - k_i a_i)/k_m$, and the worker maximizes $\Phi(a_i; \eta) = \lambda(a_i)\, C(\eta) + R\Big(a_i, (A - k_i a_i)/k_m\Big)$. Since η enters only through $\lambda(a_i)C(\eta)$,

$$\frac{\partial \Phi(a_i; \eta^B)}{\partial a_i} - \frac{\partial \Phi(a_i; \eta^A)}{\partial a_i} = \lambda'(a_i)\Big[C(\eta^B) - C(\eta^A)\Big] > 0$$

at every a_i, since λ is strictly increasing and $C(\eta^B) > C(\eta^A)$, as $\eta^B > \max\{b, \eta^A\}$. The optimum is interior: with no attention to applications, the worker applies nowhere, so evaluation is worthless, while the first unit of attention to applications has a strictly positive return; and the Inada condition on the attention rule rules out zero evaluation. The argument in the proof of Proposition 1, with a_i in place of a, then gives $a_i^B > a_i^A$, and the binding constraint gives $a_m^B < a_m^A$. □

Corollary 2 (Predictions)**.** *In the period after benefits drop, compared with post-reform workers, pre-reform workers*

(i) find jobs faster: $\lambda(a_i^B) > \lambda(a_i^A)$;

(ii) accept lower-paying jobs: conditional on finding a job, their wages are lower in the sense of first-order stochastic dominance.

Proof. Part (i) follows from Proposition 2 and the fact that $\lambda(\cdot)$ is an increasing function. For (ii), two channels lower the wage. First, less evaluation a_m makes larger consideration sets less likely (Lemma 3.1), and by the random-addition argument in the proof of Lemma 3.2,

each ranked wage improves with the size of the set, so with less attention to evaluation the wage at every rank is lower in the sense of first-order stochastic dominance. This channel lowers wages even with the same attention to applications. Second, more applications move the worker further down his ranking: conditional on finding a job, he ends up in the j-th best vacancy of his set with probability $q_j(a_i)/\lambda(a_i)$, which does not depend on the set, since $N(a_i) \leq \bar{n} \leq |T|$ by Assumption 4. More attention to applications raises the mean κa_i, and Poisson distributions with a higher mean are larger in the likelihood-ratio order, as are their versions capped at $\bar{n}$, so $\Pr(N(a_i^B) \geq j)/\Pr(N(a_i^A) \geq j)$ is increasing in j: the pre-reform worker ends up further down his ranking, in the sense of first-order stochastic dominance. Both channels go in the same direction: the first, which is specific to this model, is reinforced by the second. □

Where in the wage distribution. The model predicts lower wages but does not restrict where in the wage distribution the differences appear, which depends on how wages vary across vacancies. Importantly, the wage loss need not be concentrated among the lowest-paying jobs. With less evaluation, workers discover fewer opportunities, reducing the quality of even their best-ranked vacancies. More applications also lead them further down their ranking, but these vacancies may pay more than the best opportunities discovered by other workers. As a result, the mix of jobs found can change throughout the wage distribution. This differs from a lower reservation wage, which only adds jobs below the previous acceptance threshold, leaving the relative shares of higher-paying jobs unchanged (see below). I return to this distinction when comparing the two mechanisms with the Hungarian evidence in Section 6.

The evidence in Section 6 is consistent with these predictions. In the six weeks after day 270, pre-reform workers have a 56% higher job-finding hazard (part (i)) and find jobs paying about 9% less relative to their previous earnings (part (ii)). Moreover, wages are similar across the two groups in the bottom 40% of the distribution, while the gap emerges in the middle. This pattern is consistent with changes in the opportunities workers discover and apply to, but is harder to reconcile with a lower reservation wage alone, as discussed next.

Comparison with reference-dependent search with a reservation wage. The main alternative explanation is reference-dependent job search (DellaVigna et al. 2017). In its baseline version, workers accept every offer and the wage does not depend on their choices, so reference dependence affects how fast workers find a job but not which job they find. It can explain the higher hazard, but not the wage gap. The natural way to generate a wage gap in that framework is to add a reservation wage. Consider a reference-dependent worker

who chooses his search effort and a reservation wage w^*, and who receives offers from a distribution $F(w)$ that does not depend on his choices. A drop below the reference point makes unemployment more painful, so he searches harder and lowers w^*, as in models with both search effort and a reservation wage (DellaVigna and Paserman 2005). Accepted wages then follow F truncated at w^*, with quantile function

$$Q^{RD}(\tau; w^*) = F^{-1}\Big(F(w^*) + \tau\,[1 - F(w^*)]\Big).$$

Lowering w^* adds jobs at the bottom of the accepted distribution. Measured by its rank in the offer distribution, each quantile falls by $(1-\tau)\Big[F(w_A^*) - F(w_B^*)\Big]$, which is largest at the bottom and zero at the top; with uniform offers, the wage gap falls linearly from the bottom to the top of the distribution. Accepted wages also become more dispersed, since for any log-concave offer distribution the variance of the truncated distribution rises as the truncation point falls (Heckman and Honoré 1990). A sharper prediction holds for any offer distribution. Every job that the post-reform group accepts pays at least its reservation wage, and therefore also at least the lower reservation wage of the pre-reform group. A lower reservation wage only adds jobs below the old one, so above any wage that post-reform workers accept, the share of pre-reform jobs paying at least that wage must be the same fraction of the post-reform share. In other words, among jobs that both groups would accept, the mix of good and very good jobs is identical. Less evaluation and more applications instead change the jobs found throughout the distribution, so this proportionality does not hold in general. These predictions apply to workers who are similar, such as the high earners in the sample; with very different workers, a lower reservation wage among those with high reservation wages could add jobs in the middle of the pooled distribution. Table 2 in Appendix D compares the two models with the data.

The two mechanisms are not mutually exclusive. With the kink utility, a worker who perceives scarcity would also lower his reservation wage if it were binding. The evidence in Section 6, however, is hard to reconcile with a lower reservation wage alone: among jobs that both groups would accept, pre-reform workers find fewer of the better-paid ones (Figure 11). Reference-dependent workers who also change where they apply could produce a similar pattern, but this requires a margin of search beyond effort and the reservation wage, which is precisely the margin this extension formalizes.

Appendix C Theoretical Extensions

C.1 Extension: Task Dependence

An intuitive extension of the model is to allow for the budget in the consumption task to be equal to the agent's output in the productive task $b = y(a_y)$, thus relaxing the assumption that tasks are independent. Specifically, the budget set becomes $B(p, y(a_y)) = \{x \in S \mid p \cdot x \leq y(a_y)\}$. To save on notation, we will assume in this section that prices and the production function $y(\cdot)$ are are fixed, so we can write $B(a_y) := B(p, y(a_y))$. Here, the agent no longer derives direct utility from production; instead, production affects utility only through its effect on the budget for consumption. The agent's new maximization problem is:

$$\begin{aligned} &\max_{a_x, a_y} \mathbb{E}\Big[u(x\Big(a_x, a_y\Big); \eta)\Big] \\ &\text{s.t. } k_x a_x + k_y a_y \leq A \end{aligned}$$

where

$$x(a_x, a_y) = \underset{\mathbf{x} \in \mathbf{T}(\mathbf{a_x}) \cap \mathbf{B}(\mathbf{a_y})}{\arg\max} u(x; \eta)$$

and $T(a_x)$ denotes the random attention-dependent consideration set.

In this extension, the agent faces a new tradeoff: allocating attention a_x improves the quality of the consumption decision by expanding the consideration set $T(a_x)$, while allocating attention a_y expands the budget set $B(a_y)$ by improving production output. The optimal allocation of attention is characterized by:

$$\frac{\frac{\partial \mathbb{E}[u(x(a_x, a_y); \eta)]}{\partial a_x}}{\frac{\partial \mathbb{E}[u(x(a_x, a_y); \eta)]}{\partial a_y}} = \frac{k_x}{k_y}$$

meaning that, as before, the agent equates the marginal rate of return in expected utility from attention to consumption with that of attention to production, adjusted by the relative complexities of the two tasks.

Using the chain rule, the marginal effect of a_y on expected utility can be decomposed as:

$$\frac{\partial \mathbb{E}[u(x(a_x, a_y); \eta)]}{\partial a_y} = \frac{\partial \mathbb{E}[u(x)]}{\partial b} \cdot \frac{dy(a_y)}{da_y}$$

where:

$\frac{\partial \mathbb{E}[u(x)]}{\partial b}$ captures the marginal utility of expanding the budget,

$\frac{dy(a_y)}{da_y}$ captures how production grows with attention.

Thus, attention now trades off between expanding the set of considered options (via a_x) and expanding the feasible budget set (via a_y). The latter will depend on the production function, i.e. how effective attention is in producing money, and on how increasing the budget set improves consumption utility. Therefore, if the agent's productive task directly affects the budget of consumption **and** the agent is not myopic, it is unclear whether we should observe the effects in proposition 1. This could also explain why some papers do not consistently observe the scarcity effects on the outcomes of interest (Carvalho et al. 2016; Fehr et al. 2022). In an experiment, we can control for that by making the tasks independent, as in the baseline model.

C.2 Additional Sources of Subjective Need

This appendix discusses two additional sources through which subjective needs may be formed. These sources are not used directly in the experimental design, but they illustrate that the need threshold η can arise from several economically relevant environments.

Social comparison. Following the literature of social comparison and "Keeping up with the Jonses" (Gali 1994; Clark and Oswald 1996), an individual l observes the consumption level of those around him, and through social comparison, forms a subjective threshold of need η. For instance, consider someone living in a low- or middle-income neighborhood who then moves to a high-income area where the norm is to have a flush green lawn, luxury cars, expensive clothing, and private education for kids. In this new environment, the individual's perception of what is a necessary standard of living shifts upward, as goods and services once seen as luxuries now feel essential to maintain a good social standing.

Formally, let $\eta^l = f\Big((c(x^j))_{j\neq l}\Big)$, where $f : \mathbb{R}^{N-1} \to \mathbb{R}$ is a social comparison function. A simple example is to assume that individual l's threshold of need depends on the average level of consumption of others $\eta^l = \alpha \frac{1}{N-1} \sum_{j\neq l} c(x^j)$ where $\alpha > 0$ captures the individual's sensitivity to social comparison and N is the total number of people around him.

Uncertain environments and precautionary needs. Subjective needs may also depend on the uncertainty individuals face about future requirements. A household may not know future expenses, a worker may not know how long job search will last, or a driver may

not know how much fuel will be needed under adverse conditions. In such environments, the need threshold can reflect a precautionary buffer: the amount of resources the individual feels is necessary to be comfortable in difficult states of the world.

Formally, let $\sigma \geq 0$ denote the uncertainty associated with future requirements. The subjective need can be written as

$$\eta = h(\sigma), \qquad h'(\sigma) > 0.$$

Thus, greater uncertainty raises the amount of the resource that feels necessary. A simple functional form is

$$\eta = \bar{\eta} + \gamma\sigma, \qquad \gamma > 0,$$

where $\bar{\eta}$ is the need threshold in a stable environment and γ captures how strongly uncertainty increases precautionary needs.

C.3 Identification of Scarce Domain

One of the main goals of this paper is to provide a useful framework for researchers and practitioners to better understand the impacts of perceptions of scarcity on economic behavior. Since the model can predict both positive and negative effects on performance, a valid concern is that it could be used to explain any type of behavioral effect[32]. To address this point, it is important to recall the two main ingredients of the model that are essential for capturing these effects. The first ingredient is limited attention that needs to be allocated to two or more tasks[33]. The second ingredient is the globally concave reference-dependent utility function, with the reference point determined by a *subjective* threshold of need[34]. Therefore, if an analyst would like to verify the predictions of this model in any context, they should clearly identify that the individual faces more than one task, and that at least one of the tasks involves a reference-dependent utility function, where the marginal utility would be greater below the reference point.

The cleanest approach would be to have a proxy for attention, such as time allocated to each task, and a measurement of the reference point or a credible manipulation of the

[32] Thanks to Charlie Sprenger for raising this issue.

[33] A task does not necessarily need to be observable, for example one of the main points in the literature is that scarcity triggers mind wandering and ruminating, which take up attentional resources

[34] An S-shaped reference-dependent utility function (à la Prospect Theory) could generate the same effects but not always. The curvature of this function becomes flatter when moving away from the reference point in the loss domain, which would violate the necessary condition that the marginal utility for a good is greater when the individual perceives scarcity compared to not.

reference point. Although some datasets do contain a measurement of time allocation, this is often not attainable and we have to rely directly on measures of decision quality or performance. A common approach in the literature to study the effects of scarcity is to median split individuals based on income. This is indeed the right approach to study the effects of objective scarcity; however, for *perceived* scarcity, a closer behavioral marker that is available in many surveys is their answer to questions that ask about their financial concerns and how they feel relative to their situation with those around them.

C.4 Extension: Efficiency Bonus in the Consumption Task

In the experiment, participants receive a bonus $\beta \geq 0$ if they select the optimal bundle. This extension shows that the bonus does not affect the main results. Let x^* denote the bundle chosen when all goods are considered. Since the agent's best bundle maximizes $c(x)$ over $B(p, b)$, x^* does not depend on the need η. The consumption payoff becomes $\tilde{u}(x; \eta) = u(x; \eta) + \beta \, \mathbb{I}\{x = x^*\}$, so that

$$\mathbb{E}[\tilde{u}(x(a); \eta)] = \mathbb{E}[u(x(a); \eta)] + \beta \Pr(x(a) = x^*).$$

The agent chooses x^* if and only if his consideration set contains all the goods in x^*, since x^* is then affordable and is the unique best bundle. If x^* contains k goods, a consideration set of size m contains all of them with probability $q_m = \binom{|S|-k}{m-k} \Big/ \binom{|S|}{m}$, which is weakly increasing in m. Hence, $\Pr(x(a) = x^*) = \sum_{T \subseteq S} \mu(T|a) \, q_{|T|}$ is increasing in attention, by the same argument as in the proof of Lemma 3.2.

Two implications follow. First, the bonus raises the marginal value of attention to the consumption task. Second, since $\Pr(x(a) = x^*)$ does not depend on η, the bonus term is identical with and without perceived scarcity. It therefore cancels out in the comparison of Lemma 3.4, and the proof of Proposition 1 carries over unchanged. The bonus thus shifts attention toward the consumption task for all agents, but does not alter the effect of perceived scarcity.

Appendix D Additional Field Evidence

This appendix reports additional results and robustness checks for the Hungarian analysis in Section 6.

The main results are robust to additional controls. Pre-reform claimants were slightly more likely to have been dismissed by their employer and less likely to hold a college degree. Controlling for the reason for job loss and education leaves the results almost unchanged: the job-finding rate remains 56% higher ($p < 0.001$) and reemployment wages remain about 9% lower ($p = 0.01$). The wage gap is also similar when using average earnings in the year before unemployment, the two years before unemployment, or earnings in 2002 as the reference wage, as in Lindner and Reizer (2020) (8–11%, all $p < 0.05$).[35]

The lower wages are not simply explained by workers changing occupations. The data record both the previous and new occupations, allowing me to compare workers who change occupation with those who return to the same one. The wage gap is present in both groups and is even larger among workers returning to the same occupation.[36] This suggests that the wage difference is not mainly driven by workers switching to lower-paying occupations.

I also find no significant differences in how long the new jobs last. Among workers reemployed in the first 45 days after day 270, both groups are similarly likely to be out of work again within six months. However, job losses within this period are rare (around 9%), so the estimates are relatively imprecise. I also find no clear differences in whether workers move to less skilled occupations.

Finally, I examine the distribution of reemployment wages to better distinguish the attention mechanism from a reduction in reservation wages. If pre-reform workers simply lowered their reservation wage, they would accept additional low-paying jobs. However, among jobs that both groups would accept, the relative frequency of different wages should remain unchanged, assuming the same wage-offer distribution. The ratio of the shares of jobs paying above a given wage should therefore be constant above the post-reform reservation wage.

Figure 11 shows a different pattern. Among reemployed workers, the pre-reform group has a 28% lower share of jobs paying above the post-reform 60th percentile, and a 42% lower share above the 90th percentile. This suggests that pre-reform workers are less likely to find

[35] The gap is also similar when using log reemployment wages in levels, adjusted for consumer prices (9%, $p < 0.01$). No such gap is observed before day 270.

[36] Among workers with both occupations recorded, the wage gap is 8% ($p = 0.02$). It is 7% among those changing occupation ($p = 0.08$) and 23 log points among the 123 workers returning to the same occupation ($p = 0.03$). Since changing occupation may itself be affected by the reform, these comparisons are descriptive.

better-paying jobs, consistent with the attention mechanism. However, the decline in the ratio is imprecisely estimated and does not rule out the reservation-wage explanation.[37]

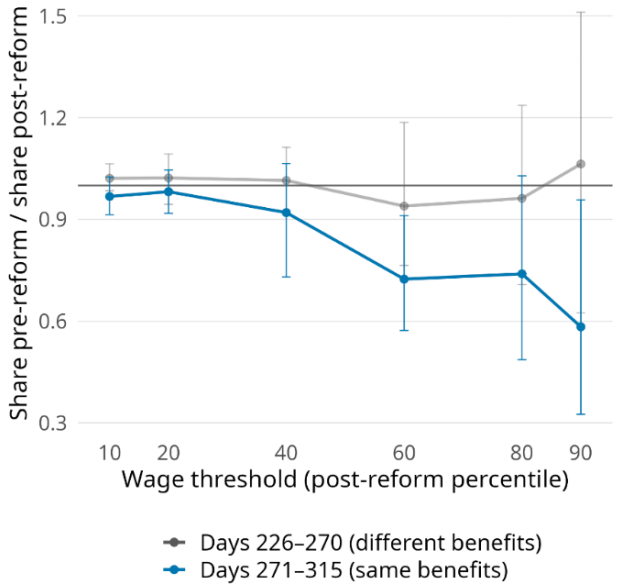


Figure 11: Share of pre-reform jobs paying at least a given amount, relative to the same share for post-reform jobs, before day 270 (grey) and in the six weeks after day 270 (blue). Wage levels are percentiles of the post-reform distribution. If pre-reform workers had only lowered their reservation wage, the line would be flat. Bars are 95% confidence intervals, obtained by resampling local unemployment insurance offices.

Table 2 summarizes how the attention mechanism and reference-dependent search with a reservation wage compare with the data.

[37]The estimated decline in the ratio between the 80th and 20th percentiles is -0.24 ($p = 0.12$), after adjusting for controls.

	Reference dependence with reservation wage	Attention allocation	Data (days 271–315)
Job-finding hazard	Higher	Higher	56% higher
Average reemployment wage	Lower	Lower	9% lower
Where the wage gap is	Bottom[a]	Not restricted	Bottom 40% similar; 9–12% lower in the middle
Wage dispersion	Higher[b]	Not restricted	Lower ($p = 0.04$)
Ratio of shares of jobs that both groups would accept	Flat	Falls with the wage	Falls with the wage: 0.72 at p60, 0.58 at p90

Table 2: Predictions for pre-reform relative to post-reform workers after day 270. The reference-dependent model has search effort and a reservation wage; the attention model is the one in Appendix B. Both are stated for similar workers. [a]In terms of rank in the offer distribution; in wage levels this depends on the offer distribution. [b]For log-concave offer distributions. The last row is the share of pre-reform jobs paying at least a given wage relative to the post-reform share, above the post-reform reservation wage; data at the post-reform 60th and 90th percentiles (Figure 11).

Appendix E Additional Figures

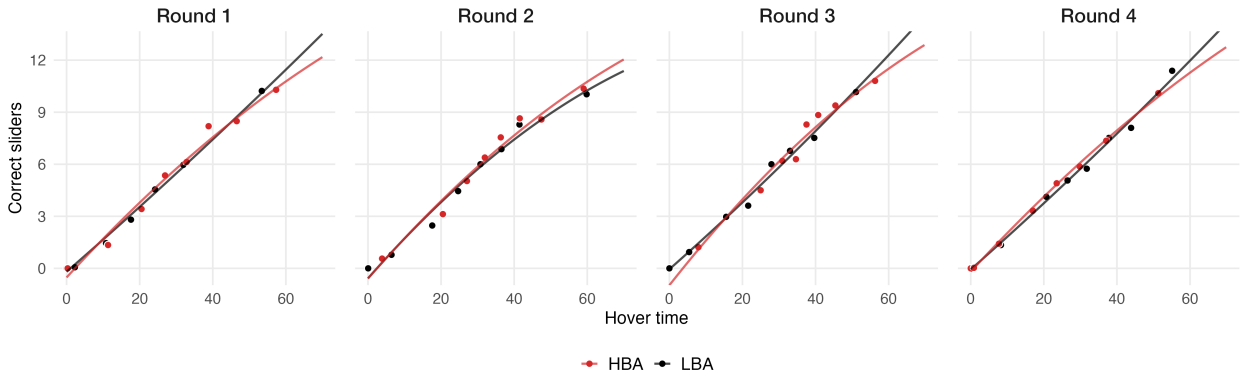


Figure 12: Correct sliders against hover time in the productivity task, by round. Within each group and round, participants are split into 8 equally sized bins of hover time; each point shows a bin's mean, and lines are quadratic fits for each group.

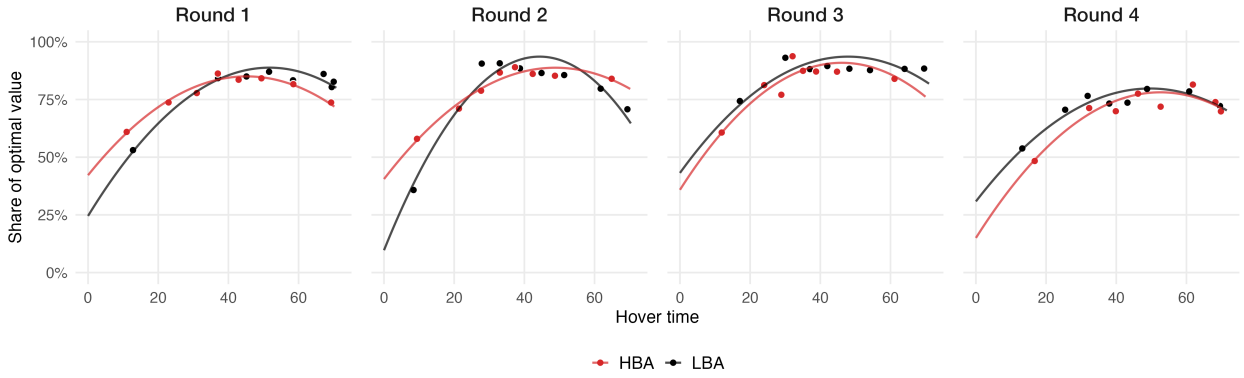


Figure 13: Decision quality (the value of the chosen bundle as a share of the value of the optimal bundle) against hover time in the consumption task, by round. Within each group and round, participants are split into 8 equally sized bins of hover time; each point shows a bin's mean, and lines are quadratic fits for each group.

Table 3: Performance and hover time

	Hover time	Hover time^2 (×1000)	Slope per 10 seconds		N
			below 35s	above 35s	
Panel A. Consumption task: decision quality					
All rounds	0.024 (0.002)	−0.257 (0.021)	0.179 (0.016)	−0.017 (0.005)	2,004
Round 1	0.022 (0.003)	−0.219 (0.033)	0.154 (0.027)	−0.017 (0.008)	501
Round 2	0.030 (0.002)	−0.338 (0.027)	0.206 (0.017)	−0.038 (0.010)	501
Round 3	0.022 (0.004)	−0.231 (0.039)	0.167 (0.026)	−0.003 (0.007)	501
Round 4	0.022 (0.003)	−0.213 (0.028)	0.167 (0.020)	0.003 (0.009)	501
Excl. time limit	0.026 (0.002)	−0.282 (0.025)	0.179 (0.016)	−0.012 (0.006)	1,788
Panel B. Productivity task: correct sliders					
All rounds	0.221 (0.013)	−0.478 (0.283)	1.960 (0.043)	1.386 (0.224)	2,004
Round 1	0.202 (0.014)	−0.199 (0.303)	1.886 (0.067)	1.294 (0.287)	501
Round 2	0.242 (0.018)	−0.945 (0.376)	2.008 (0.068)	1.140 (0.278)	501
Round 3	0.218 (0.022)	−0.242 (0.470)	2.013 (0.078)	1.662 (0.346)	501
Round 4	0.207 (0.017)	−0.184 (0.392)	1.938 (0.053)	1.739 (0.336)	501

Notes: Regressions of performance on hover time (in seconds) and the squared term. The coefficient on hover time is the slope at 0s. Slope per 10 seconds is ten times the coefficient on hover time in separate linear regressions for participant-rounds with a hover time below and above 35 seconds, the midpoint of the 70-second round. In Panel A, decision quality is the value of the chosen bundle as a share of the value of the optimal bundle; "excl. time limit" drops participant-rounds with a hover time of 69 seconds or more. Standard errors in parentheses, clustered by participant in the regressions that pool all rounds and heteroskedasticity-robust in the single-round regressions.