UNIVERSITAT D'ALACANT
UNIVERSIDAD DE ALICANTE
Facultat de Ciències Econòmiques i Empresarials
Facultad de Ciencias Económicas y Empresariales

MASTER'S DEGREE IN ECONOMICS WITH DATA SCIENCE

ACADEMIC YEAR 2025 - 2026

TIME PRESSURE AND MULTIATTRIBUTE SEQUENTIAL SEARCH: EVIDENCE FROM AN EXPERIMENTAL SETTING

ALEJANDRO SÁNCHEZ NARANJO

Tutor
ADAM SANJURJO
DEPARTMENT OF FUNDAMENTALS OF ECONOMIC ANALYSIS

San Vicente del Raspeig, 7th September 2026

TIME PRESSURE AND MULTIATTRIBUTE SEQUENTIAL SEARCH: EVIDENCE FROM AN EXPERIMENTAL SETTING

Thesis by

ALEJANDRO SÁNCHEZ NARANJO

Supervisor

ADAM SANJURJO

In Partial Fulfillment of the Requirements
for the Master's Degree in
Economics with Data Science

Department of Fundamentals of Economic Analysis
Faculty of Economics & Business Sciences
University of Alicante
San Vicente del Raspeig, Spain

2026

Alejandro Sánchez Naranjo
University of Alicante, 2026

Abstract

This thesis examines how time pressure shapes multiattribute sequential search behavior, and whether the coping mechanisms subjects use to adapt to time pressure are associated with departures from optimal search. Using the Exogenous treatment of Gabaix et al.'s (2006) Mouselab experiment (390 subjects, 4,548 tasks), task-level measures of three coping mechanisms identified in the time-pressure literature (acceleration, filtration, and avoidance (strategy shift)) are constructed, and compliance with the four necessary conditions for optimal search proposed by Sanjurjo (2017) is evaluated. On their own, they respond unevenly to time pressure: pressure level is linked to fewer repeated look-ups but a higher probability of biased switching, with little or no relationship for the remaining conditions. Mixed-effects models show that the coping mechanisms are a stronger predictor of search quality than time pressure itself. Avoidance is linked to a higher probability of violation across nearly every condition. Acceleration shows no robust association with any of them. Filtration shows a mixed pattern: it worsens repeated look-ups but improves search sequencing, without affecting the accuracy of the final choice. These results suggest that coping with time pressure comes at a systematic cost to optimal search, and that this cost operates mainly through the specific mechanism a subject adopts, rather than through time pressure directly.

This thesis examines how time pressure shapes multiattribute sequential search behavior, and whether the coping mechanisms subjects use to adapt to time pressure are associated with departures from optimal search. Using the Exogenous treatment of Gabaix et al.'s (2006) Mouselab experiment (390 subjects, 4,548 tasks), task-level measures of three coping mechanisms identified in the time-pressure literature (acceleration, filtration, and avoidance (strategy shift)) are constructed, and compliance with the four necessary conditions for optimal search proposed by Sanjurjo (2017) is evaluated. On their own, condition violations respond unevenly to the time budget: more time is linked to more repeated look-ups but a lower probability of biased switching, with little or no relationship for the remaining conditions. Mixed-effects models show that the coping mechanisms are a stronger predictor of search quality than the time budget itself. Avoidance is linked to a higher probability of violation across nearly every condition. Acceleration shows no robust association with any of them. Filtration shows a mixed pattern: it worsens repeated look-ups but improves search sequencing, without affecting the accuracy of the final choice. These results suggest that coping with time pressure comes at a systematic cost to optimal search, and that this cost operates both through the specific mechanism a subject adopts and through a direct effect of the time budget itself.

Keywords: Multiattribute sequential search, Time pressure, Coping mechanisms, Necessary conditions for optimality.

Contents

List of Tables

List of Figures

1 Introduction

Many important decisions in everyday life require choosing among several alternatives that differ along multiple attributes, often within a limited amount of time. Choosing a job offer means weighing salary against location and growth prospects before an offer expires, selecting a health insurance plan means comparing premiums, coverage, and deductibles before an enrollment deadline closes, picking a supplier for a business means balancing price, quality, and delivery time under the pressure of an approaching decision. In each of these cases, the decision maker cannot examine every piece of information about every alternative, information is costly to acquire, and time to acquire it is often scarce.

Economic theory has traditionally modeled decision makers as fully rational agents who evaluate all available information before choosing the option that maximizes their utility. Yet this benchmark assumes that acquiring and processing information is costless and instantaneous, an assumption that becomes especially complicated once alternatives are defined by several attributes at once and time to search is limited. Understanding how real decision makers navigate this gap between the rational ideal and the practical constraints they face is central to behavioral and experimental economics, and is the starting point of this thesis.

Although the literature on multiattribute sequential search and the literature on time pressure in decision making are both well established, they have developed largely in parallel. On one hand, Gabaix et al. (2006) and Sanjurjo (2017) have shown that optimal search in multiattribute environments is analytically intractable, and that subjects systematically deviate from the necessary conditions that any optimal search policy must satisfy, over-searching within alternatives and switching between them in biased ways. On the other hand, a separate body of work, following Zur and Breznitz (1981) and Payne et al. (1988), has documented that decision makers cope with time pressure through mechanisms such as acceleration, filtration, and avoidance, typically in simpler decision environments.

What remains largely unexamined is how these two scenarios connect: whether the coping mechanisms identified in the time-pressure literature are themselves responsible for the departures from optimal search documented in the multiattribute search literature. Understanding this connection matters both theoretically, for characterizing how bounded rationality operates under the joint constraints of multiple attributes and limited time, and practically, for anticipating how real decision makers are likely to behave when facing complex choices under deadlines.

Then, this thesis is guided by one central research question: *how does time pressure shape multiattribute sequential search behavior?* Building on this, a secondary question follows naturally: *are the adaptive responses subjects use to cope with time pressure associated with violations of optimal search?*

1.1 Objectives

General Objective

To examine how time pressure shapes multiattribute sequential search behavior, and whether the adaptive responses subjects use to cope with time pressure are associated with violations of optimal search.

Specific Objectives

- To determine whether subjects exhibit acceleration, filtration, and strategy shift as adaptive responses when the time available for search decreases.
- To evaluate the extent to which subjects comply with the necessary conditions for optimal sequential search proposed by Sanjurjo (2017), and whether compliance is related to the time pressure.
- To assess whether these adaptive responses are associated with the probability of violating each necessary condition.

1.2 Hypotheses

- **H1:** As the time pressure increases, subjects will show evidence of acceleration, filtration and avoidance.
- **H2:** Violations of the necessary conditions for optimal search will increase as the time budget decreases.

- **H3:** The use of adaptive responses will help subjects cope with time pressure, reducing the probability of violating the necessary conditions relative to what the time budget alone would predict.

The remainder of this thesis is organized as follows. Chapter 2 develops the theoretical framework, situating multiattribute sequential search within the literature on bounded rationality, and reviewing the necessary conditions for optimal search and the mechanisms through which decision makers adapt to time pressure. It then reviews the empirical literature on multiattribute search and on adaptive decision making under time pressure, identifying the research gap this thesis addresses. Chapter 3 describes the experimental design underlying the data and the empirical strategy used throughout the analysis, including the construction of the search behavior, condition-violation, and coping-mechanism variables. Chapter 4 presents the results: how search behavior responds to the time budget, the extent of compliance with the necessary conditions for optimal search, and whether coping mechanisms are associated with condition violations and with economic performance. Chapter 5 concludes.

2 Literature Review

2.1 Theoretical Framework

This section establishes the conceptual foundations underlying the analysis of multiattribute sequential search under time pressure. We begin by situating the problem within the broader literature on bounded rationality and decision making, which provides the cognitive framework for understanding how agents navigate complex information environments. We then formalize the multi-attribute sequential search problem, outlining its structure as the evaluative standard for search behavior. Finally, we examine time pressure as a situational constraint that shapes how decision makers adapt their search strategies. Together, these three building blocks define the theoretical space within which the empirical analysis is conducted.

2.1.1 Bounded Rationality and Decision Making

Understanding how agents search among multiple alternatives first requires understanding how economics has traditionally modeled decision-making itself. Classical economics, however, was not built around the decisions of individual agents. Smith and Ricardo explained market outcomes, prices and the distribution of income among wages, profits, and rent, as the result of objective factors of production, the cost and the labor embodied in each good, known as the labor theory of value (Ricardo, 1817; Smith, 1776). Individual choice played no explicit role in this account, value was treated as a property of the good itself, not as the outcome of an agent's decision process. This worked reasonably well for goods whose price tracked their labor content, but it failed to explain cases where it did not, most notably the diamond-water paradox (Ekelund & Hébert, 1997).

The neoclassical school, emerging from the marginal revolution, addressed this gap by shifting the unit of analysis from aggregate classes to the individual agent. Price and value were now derived from the optimizing behavior of a decision maker choosing under a fixed set of preferences and constraints (Inoua & Smith, 2020). This shift is captured in Robbins's (1932) definition of economics itself as "the science which studies human behaviour as a relationship between ends and scarce means which have alternative uses". The agent implied by this definition, centered not on wealth or production but on the act of choosing, is endowed with complete information about the alternatives, their attributes, and the outcomes associated with each choice it faces in the market (Hicks & Allen, 1934), and is *rational* in the sense that it evaluates all available options and selects the one that maximizes expected utility (von Neumann & Morgenstern, 1944).

However, the model of *bounded rationality* challenges this assumption. Simon (1955) argued that agents do not maximize utility by comparing every alternative, but instead choose the first option that meets a minimum aspiration level. Search is no longer costless or instantaneous, it becomes a sequential process that stops as soon as a "good enough" option is found, not necessarily the best one. Section 2.1.2 develops this sequential search process in more detail, formalizing how agents decide when to stop searching among multiple alternatives.

Simon's critique opened the door to a broader research program: *Behavioral Economics*, which studies how people actually make decisions once the traditional assumptions are relaxed.Specifically, it challenges the notions that agents hold well-defined preferences and unbiased beliefs, that they make optimal choices based on them (implying unlimited cognitive ability and willpower) and that they possess complete information (Thaler, 2016). Rather than treating deviations from the rational choice model as errors to be corrected, behavioral economics treats them as systematic and predictable features of human decision-making, to be studied empirically.

Building on this behavioral turn, Payne et al. (1988) proposed the notion of *adaptive decision making*: rather than applying a single, fixed decision rule, agents select among a repertoire of heuristics depending on the demands of the task, the number of alternatives, the number of attributes, and the constraints under which the decision must be made. In this view, choosing how to decide is itself a decision: more exhaustive strategies tend to deliver more accurate choices, but demand more cognitive effort, so agents weigh this trade-off when selecting a strategy.

Taken together, this literature portrays the agent not as a perfect optimizer, but as a boundedly rational decision maker who does the best it can with the heuristics and cognitive tools available, adapting its search strategy to the demands and constraints of the task at hand.

2.1.2 Multi-attribute Sequential Search

Many real-world decisions require choosing among alternatives that differ along several dimensions at once, like a product that varies in price, quality, and brand. Keeney and Raiffa (1976) formalized this class of problems as *multi-attribute decision making*: the decision maker must aggregate information across multiple, often conflicting, attributes into a single overall evaluation of each alternative, rather than optimizing along a single dimension.

If information about every alternative's attributes were freely and instantly available, agents could evaluate the entire choice set at once, as the classical model assumes. In practice, however, acquiring information is costly: it takes time, attention, and effort to inspect each attribute of each alternative. Stigler (1961) was among the first to formalize this idea, showing that the acquisition of information about prices and alternatives has a cost, just like any other economic activity, so agents only keep searching as long as the expected benefit of finding a better alternative outweighs that cost (cost-benefit calculation). This is what makes search inherently *sequential*: rather than acquiring all information simultaneously, the decision maker samples pieces of information one at a time, deciding at each step whether to continue searching or to stop and choose.

When alternatives differ along a single dimension, the sequential search problem admits an elegant solution. Weitzman (1979) showed that an optimal stopping rule exists, based on a simple reservation price for each alternative. This tractability, however, does not extend to the multi-attribute case. Gabaix et al. (2006) show that once alternatives vary along several attributes at once, the high dimensionality of the search problem makes characterizing the optimal search strategy analytically and numerically intractable, even for a moderate number of alternatives and attributes. As a result, the literature has instead focused on characterizing *necessary conditions* that optimal search must satisfy, rather than deriving the optimal policy itself, an approach developed by Sanjurjo (2017), whose four conditions serve as the empirical benchmark for this thesis.

For a sequential search to be optimal, it must satisfy four basic requirements (Sanjurjo, 2017) . First, within any single alternative, information must be gathered starting from its most informative attribute, ***Condition V (Variance)*** rules out searching a less informative attribute before a more informative one in the same alternative. Second, when deciding whether to keep exploring the current alternative or move to another one, the searcher must always move toward the alternative that is more promising given what is currently known, ***Condition T (Trade-off)*** captures this choice between depth (staying within an alternative) and breadth (switching across alternatives). Third, no piece of information should ever be checked twice, ***Condition R (Repeat)*** simply requires that each attribute be searched at most once. Finally, once the search ends, the searcher must select the alternative that looks best given everything learned so far, ***Condition C (Choice)*** requires that the final decision be consistent with the information gathered during search. How violations of these conditions are measured and interpreted is discussed in detail in the Section 3.3.

Taken together, this section has shown that when a decision maker faces alternatives defined by multiple attributes, no fully optimal search policy is known to exist. Nevertheless, an optimal search must still satisfy certain necessary conditions, offering a partial, but rigorous, rational benchmark against which actual search behavior can be evaluated, and against which the heuristics decision makers use in its place can be understood.

2.1.3 Time Pressure in Decision Making

Time pressure refers to a situation in which a decision maker must reach a choice before a deadline, rather than being free to deliberate indefinitely. Research on time pressure has approached it in two ways: as a stressor that alters the decision maker's affective state, or as one more factor entering the cost-benefit calculation of which cognitive strategy to use (Maule & Hockey, 1993). This thesis follows the latter approach, treating time pressure as a constraint that shapes the choice of search strategy, consistent with the adaptive decision making framework introduced in Section 2.1.1. Under this view, the relevant question is not only whether time pressure worsens decision outcomes, but how it changes the process leading up to them.

Zur and Breznitz (1981), identify three ways in which decision makers cope with time pressure: ***acceleration***, processing the same information at a faster rate; ***filtration***, selectively processing only the most important subset of the available information; and a shift toward a less effortful decision strategy, which in its most extreme form collapses into near-random choice, a pattern they term ***avoidance***. This thesis follows Payne et al. (1988) in treating this third mechanism as a genuine *strategy shift* rather than mere avoidance, while retaining Zur and Breznitz's original terminology.

These three mechanisms need not operate in isolation, or with equal likelihood. Weenig and Maarleveld (2002) suggest that they may follow a rough hierarchy. Acceleration is the first response to time pressure, while increased selectivity (filtration) and switching to another strategy are reserved for later, more severe stages of time constraint. However, Wallsten (1993) points out that although this is not generally acknowledged, filtration and strategy shift are not mutually exclusive responses, a decision maker may adopt both simultaneously, rather than one replacing the other as time pressure intensifies.

Time pressure is especially consequential in multi-attribute search problems. As the number of alternatives and attributes grows, it becomes impossible to evaluate every option exhaustively without any time constraint (recall the computational intractability discussed in Section 2.1.2). Adding a time limit makes this worse. Payne (1982) shows that when a task is more complex, more options and more attributes, people rely more on simple and selective shortcuts instead of checking everything. So the more alternatives and attributes there are, the more likely time pressure is to push a decision maker toward acceleration, filtration, or a shift in strategy.

2.2 State of the Art

This section reviews the empirical evidence that informs the present analysis across two interconnected bodies of literature. We first examine experimental findings on multi-attribute search behavior, focusing on what is known about how agents actually navigate these environments and where they systematically deviate from optimal search. We then survey the evidence on adaptive decision making under time pressure, documenting what coping mechanisms have been observed and under what conditions they are activated. We bring both literatures together to establish the empirical foundation on which the present analysis is built.

2.2.1 Multi-attribute Search

Gabaix et al. (2006) provide the foundational experimental evidence on multi-attribute sequential search. They study how individuals acquire costly information when making decisions among multi-attribute scenarios. Their main objective is to examine whether information search is better explained by fully rational optimization or by a boundedly rational (myopic) process, which they formalize through the Directed Cognition (DC) model. The rational model considers the value of the entire future search process, including the possibility of continuing to search after each information acquisition. Directed Cognition is simpler, it evaluates only the next search step, as if it were the last opportunity to obtain information.

To test these predictions, the authors conduct two laboratory experiments. In the first experiment, participants sequentially acquire information about three uncertain investment projects, where each information acquisition carries an explicit monetary cost. This simple setting allows a direct comparison between both models. In the second experiment, participants complete a more complex multi-attribute search task using the Mouselab interface, where information is hidden behind boxes and search is costly because time is scarce. Subjects decide which attributes to reveal before selecting the alternative with the highest expected value, allowing the authors to observe the entire sequence of information acquisition. In this environment, a fully rational solution would require solving a high-dimensional dynamic programming problem (intractable as mentioned in Section 2.1.2), whereas the Directed Cognition model remains computationally tractable by evaluating one search operation at a time.

Across both experiments, the Directed Cognition model explains observed search behavior better than the fully rational benchmark. The authors conclude that individuals behave as boundedly rational decision makers who allocate attention according to approximate cost-benefit calculations rather than solving the computationally intractable optimal search problem.

Building on the experimental dataset of Gabaix et al. (2006), Sanjurjo (2017) addresses one of the main limitations acknowledged in the original study, the lack of a tractable rational benchmark for high-dimensional search problems. Rather than comparing observed behavior with the Directed Cognition model, the author derives a set of necessary conditions for optimal sequential search which make it possible to evaluate the rationality of search behavior in the same Mouselab environment.

Applying these conditions to the original dataset reveals systematic departures from optimal search. While violations of Condition V are relatively uncommon (5.6% of eligible look-ups), violations of Condition T occur in almost half of all eligible searches (48.6%). More specifically, subjects continue searching within the same alternative when they should switch in 47.2% of cases, and choose an incorrect alternative when switching in 54.9% of cases. Moreover, participants fail to leave an alternative when the optimal policy prescribes switching in 94% of the relevant decisions, indicating a strong tendency to search too deeply within alternatives. Violations of the remaining conditions are also substantial: 12% of look-ups repeat previously inspected attributes, violating Condition R, and subjects fail to choose the alternative with the highest revealed expected value in 30.8% of tasks, violating Condition C. Overall, 57% of eligible search actions violate at least one of Conditions T, V, or R, and 94% of search problems contain at least one violation of the necessary conditions.

The paper also shows that these deviations are economically meaningful. Subjects who violate Condition T more frequently obtain less valuable information during search and earn lower expected payoffs, whereas violations of Condition V have little measurable impact on performance. These findings suggest that the main inefficiency in multiattribute search is not the order in which attributes are inspected within an alternative, but the failure to optimally balance search depth and breadth across alternatives.

Relatedly, Bearden and Connolly (2007) examine whether decision makers can balance search depth within an option and search breadth across sequentially encountered options in a simpler two-attribute environment. Using a laboratory experiment with two-attribute alternatives, participants could pay to reveal information before deciding whether to accept an option or continue searching. The authors compare behavior under two information settings: one in which participants observed the exact value of an attribute and another in which they only learned whether it exceeded a chosen threshold. They find that participants searched too little when exact information was available (average stopping position 2.79 versus the optimal 3.21) but searched too much when information was presented as thresholds. Nevertheless, performance remained relatively close to the optimal benchmark because the additional search in the threshold condition led to better choices, offsetting the higher search costs. Overall, the results suggest that individuals systematically adapt their search behavior to the way information is presented, although they still deviate from the optimal search policy.

Complementing the experimental evidence on multi-attribute search, Caplin et al. (2011) investigate how individuals decide when to stop gathering information. In their laboratory experiment, participants sequentially acquired information before making a decision, allowing the authors to observe both the search process and the stopping decision. The results provide strong support for sequential search models: participants generally continued searching while the expected benefit of additional information remained high and stopped once the expected gains became sufficiently small. This finding suggests that stopping decisions are not random but respond systematically to the value of the remaining information, providing further evidence that individuals adapt their search behavior to the decision environment, even if they do not follow the fully optimal search policy.

Extending this literature beyond the laboratory, Gardete and Hunter (2024) analyze detailed online search data from consumers shopping for used cars. Their data show that consumers do not inspect products in an all-or-nothing manner; instead, they gather information gradually, looking at some attributes, returning to previously viewed cars (approximately 7% of consumers revisited previously inspected products), and sometimes choosing a vehicle without examining all the available information (only 32.5% inspected vehicle histories, and just 13.7% viewed inspection reports). These patterns are inconsistent with standard models in which consumers either inspect complete products sequentially or choose a fixed sample in advance. The authors therefore develop a model of gradual, attribute-level search and use it to study how the information shown by sellers affects search and market outcomes. Their results confirm that multi-attribute search is selective and dynamic, and that the way information is organized can influence both consumer behavior and seller performance.

Building directly on the deviations identified in his earlier work, Sanjurjo (2023) studies whether search patterns such as examining one alternative at a time and moving systematically across the screen may help reduce memory demands. He introduces the idea of space complexity, which measures how much information a decision maker must keep in working memory while comparing multi-attribute options. In a laboratory experiment, participants either followed imposed information-processing orders or chose their own. Subjects performed much better when they processed all the attributes of one alternative before moving to the next, and more than 90% of participants who were free to choose their order almost always adopted this low-complexity pattern on their own. These findings provide a possible explanation for the deep and adjacent search patterns documented by Sanjurjo (2017): although such behavior may violate the optimal search order, it can simplify the task and reduce costly choice errors.

2.2.2 Adaptive Decision Making Under Time Pressure

One of the earliest frameworks explaining how individuals adapt to time pressure was proposed by Zur and Breznitz (1981). They argue that decision makers cope with limited time through three main mechanisms: acceleration, filtration, and avoidance. To test these mechanisms, participants made repeated choices between monetary gambles under different levels of time pressure. The results showed strong evidence of acceleration, as participants spent less time processing each piece of information, and filtration, as they inspected fewer cues, the average number of information acquisitions decreased from 18.2 under low time pressure to 6.8 under high time pressure, and increasingly focused on information related to potential losses. However, little evidence was found for avoidance, since participants continued to make systematic rather than random decisions even under severe time constraints. These findings suggest that, rather than abandoning deliberation altogether, individuals primarily adapt by processing information more efficiently and selectively.

Building on that, Payne et al. (1988) examine whether time pressure changes not only how much information people process, but also the decision strategy they use. The paper combines computer simulations with two laboratory experiments using Mouselab. In the experiments, participants chose among multi-attribute alternatives either without a time limit or under limits of 15 and 25 seconds. Under time pressure, they examined less information and processed each item more quickly; in the 15-second condition, the average number of information acquisitions fell from 35.0 to 16.7. Under moderate pressure, adaptation occurred mainly through acceleration and some filtration, whereas severe pressure also produced a shift toward simpler, attribute-based heuristics. The results therefore show that strategy change is an additional mechanism of adaptation when acceleration and filtration are no longer sufficient.

Extending the adaptive decision-maker framework to strategic settings, Spiliopoulos et al. (2018) examine whether players also respond to time pressure through acceleration, filtration, and avoidance (strategy shift). Participants played 29 two-player games using Mouselab, under three conditions: no time limit, a 20-second time limit, and a condition requiring them to wait at least 45 seconds before deciding. Comparing the no-limit and 20-second conditions, average response time fell from 44.3 to 13 seconds under time pressure, and participants examined less information overall, particularly information about the other player: the share of the opponent's payoffs they checked dropped from 91.9% to 54.7%, and in 27% of pressured decisions, subjects did not look at any of the opponent's payoffs at all. This narrower search was accompanied by a shift toward simpler decision rules: subjects increasingly relied on *Level-1* reasoning (choosing a best response to the assumption that the opponent moves randomly, a simple heuristic that ignores strategic depth), which rose from 26.5% to 46.8% of decisions, while *Dominance-1* reasoning (eliminating one's own dominated strategies before choosing, a more demanding rule that requires comparing payoffs across the full game) fell sharply from 31.6% to 2.5%. Despite these substantial changes in both search and decision rule, expected payoffs did not decline significantly under time pressure, suggesting that this adaptation can reduce cognitive effort without necessarily reducing performance.

Reutskaja et al. (2011) study how people search and choose under extreme time pressure. Using eye-tracking, they showed 41 subjects sets of 4, 9, or 16 snack foods and gave them only 3 seconds to choose. They found that as the choice set grew, subjects saw a shrinking share of the options, from 83% of items seen when there were 4 to just 34.5% when there were 16, and their choice quality fell accordingly, from 70.8% to 50.1% efficiency. Subjects still chose the best item they had seen most of the time (69–76% depending on set size), showing that time pressure mainly cuts how much gets searched, not how well people decide given what they saw.

A related but distinct question is whether this cost of time pressure fades with experience, which is what Klimm et al. (2023) study in a different kind of sequential search. Here, 192 participants repeatedly decided whether to accept a price or pay a cost to see another offer, with a monetary penalty for slow decisions in the high-pressure condition. Averaged across ten rounds of practice, time pressure barely changed how much people searched. But in the very first round, before subjects had learned the task, time pressure made them stop searching too early and cut their profits by more than 40% (13.33 vs. 23.14 experimental currency units).

Together, these five studies show that acceleration, filtration, and strategy shift are consistently observed across very different settings (gambles, multiattribute Mouselab choices, strategic games, and consumer search)establishing them as the dominant, well-replicated responses to time pressure. They also show that these mechanisms are not simply harmless shortcuts: while they sometimes preserve overall performance, they can distort how broadly and deeply people search, and in some cases directly reduce profits. What remains largely unexamined, however, is how the three mechanisms operate jointly within a single multiattribute search environment, where switching between alternatives and searching within them are governed by distinct necessary conditions for optimality. Payne et al. (1988) and Spiliopoulos et al. (2018) document the mechanisms themselves but do not test them against a formal rationality benchmark. Sanjurjo (2017) provides that benchmark but does not examine the coping mechanisms that might explain the violations. This thesis brings both together, using the coping-mechanism framework reviewed here to explain deviations from Sanjurjo's (2017) necessary conditions in the same Mouselab dataset, and thereby closes the gap between these two literatures directly.

3 Methodology

3.1 Experimental Design

The data used in this thesis come from the multi-attribute search experiment conducted by Gabaix et al. (2006). In each task, subjects chose among eight alternatives, with the value of each alternative equal to the sum of ten individual attribute values. Only the first attribute of each alternative was observable at no cost, the remaining nine were concealed and could only be revealed through costly search, one at a time. Attributes were normally distributed with mean zero and variance decaying linearly from the first attribute to the tenth, so that earlier attributes carried more information than later ones. Subjects were paid the sum of all ten attribute values of the alternative they ultimately chose, regardless of how many attributes they had actually searched.

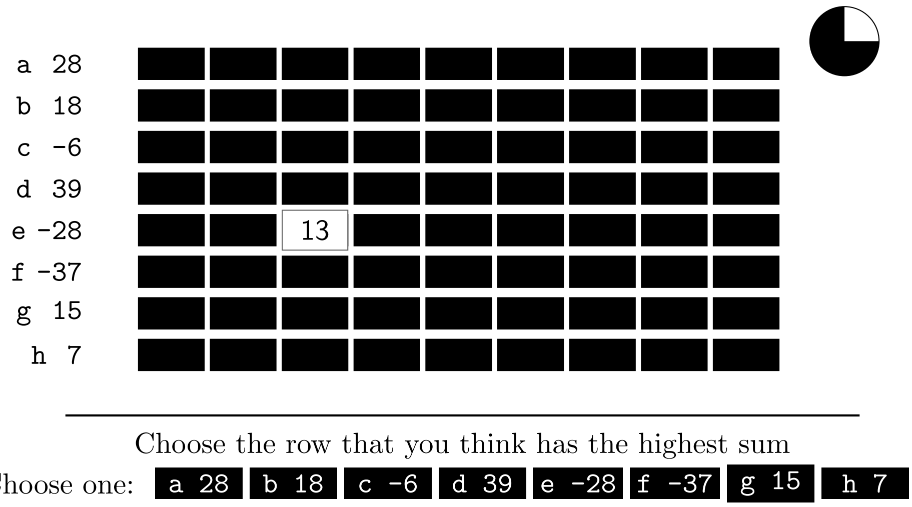

Figure 1: Sample Game with Values Concealed

Note: Illustrative representation of the search interface, designed by the author based on the experimental design of Gabaix et al. (2006). Alternatives are arranged in rows (1 to 8) and attributes in columns (1 to 10). The value of the first attribute is visible for the entire duration of the task, while the remaining nine attributes per alternative are concealed and can only be revealed one at a time by clicking on the corresponding box. A choice box for each alternative appears at the bottom of the display, and the remaining time is shown continuously through a countdown timer in a corner of the screen.

Gabaix et al.'s experiment included two treatments that varied how time constraints were imposed. In the *Endogenous* treatment, subjects were given a fixed budget of 25 minutes to complete as many problems as they chose, allocating their own time across tasks. In the *Exogenous* treatment, subjects were instead allocated a fixed amount of time, drawn from a uniform distribution between 10 and 49 seconds, for each of 12 tasks.

This thesis focuses exclusively on the Exogenous treatment. Because time is assigned externally and independently of the subject's own choices, it provides a cleaner source of variation in time pressure: differences in search behavior across tasks can be attributed to the amount of time available, rather than to the subject's own decision of how much time to allocate.

Gabaix et al. used the MouseLab experimental interface, building on earlier work by Payne et al. (1988) and Johnson et al. (1989), to track subjects information acquisition in complete detail. MouseLab is essentially a mechanical analog to eye-tracking, attribute values are hidden inside boxes on the screen, and each box must be opened with a left-click of the mouse to reveal its value, then closed with a right-click before another box can be opened. This interface records the complete sequence and duration of every attribute look-up performed by a subject, rather than only the final choice or an aggregate summary of what was searched. As a result, the dataset preserves not just what information subjects acquired, but the exact order and timing with which they acquired it, the level of detail required to construct the search-based variables and necessary-condition violations used in this thesis.

The experiment's 390 subjects were Harvard undergraduate students. The design was within-subject: each subject completed problems under both the Exogenous and Endogenous treatments, with roughly half of the subjects beginning in the Exogenous treatment and the other half in the Endogenous, in order to control for order effects. Subjects were given complete information about the distribution of all attributes and, in each problem, the amount of time remaining.

3.2 Data

The dataset underlying this thesis is structured across three levels of observation: events, tasks, and subjects. An *event* is a single attribute look-up, a subject opening one box in one alternative, at one point in time. A *task* is one of the 12 search problems a subject completed under the Exogenous treatment, and consists of the full sequence of events performed with the subject's final choice. Each subject completed multiple tasks, and there are 390 *subjects* in the filtered sample used.

Each event record includes the alternative and attribute searched and the time at which it occurred. The sequence and timing of these events determines the order and duration of search within each task. A separate variable records the subject's final choice, from which the task's *score* is calculated as the sum of the ten attribute values belonging to the chosen alternative, regardless of how many were actually searched. In a small number of task, however, this choice variable is missing, these task are excluded from the analysis, together with the time-related inconsistencies described below.

The raw dataset does not record the time actually allotted to each task. As a proxy, this thesis uses *endtime*: the moment at which the allotted time expired, for task in which the subject searched until the time limit, or the time at which the subject made their choice, for task in which the subject decided before time ran out. Tasks with an *endtime* outside the experiment's possible time range (10 to 49 seconds) are also excluded. This proxy was validated against the theoretical uniform distribution used to assign time budgets (see Figure A.1 in the Appendix).

After applying these filters, 132 tasks were removed due to inconsistencies, leaving 4,548 tasks across 390 subjects and 160 distinct task types. Table 1 shows the resulting distribution of the number of task completed per subject. The large majority of subjects (301, or 77.18%) completed all 12 task, and a further 67 (17.18%) completed 11. The remaining subjects completed between 6 and 10 games, together accounting for less than 6% of the sample. Given this concentration near the maximum, no additional subjects were excluded.

Games Played	Freq.	Percent
6	1	0.26
7	1	0.26
8	5	1.28
9	4	1.03
10	11	2.82
11	67	17.18
12	301	77.18
Total	390	100.00

Table 1: Distribution of Task Completed per Subject

3.3 Empirical Strategy

3.3.1 Defining Time Pressure

To classify tasks by the severity of time pressure, this thesis uses the Endogenous treatment of Gabaix et al. (2006)'s experiment as an external reference point. In this treatment, subjects are given a fixed budget of 25 minutes to allocate freely across as many problems as they choose, rather than being assigned a fixed time for each individual task. The median time subjects take per problem under this self-paced condition, approximately 38.06 seconds (mean: 43.68 seconds), is taken as a proxy for the amount of time subjects choose to use when they are not constrained by an externally imposed deadline. A threshold of 40 seconds is adopted, slightly above the median to account for the higher mean and take a somewhat conservative view of what counts as time pressure: tasks in the Exogenous treatment with a time budget below 40 seconds are classified as being under time pressure, while tasks with a time budget at or above it are classified as *No pressure* and serve as the baseline against which the coping mechanisms described below are evaluated.

Within the range of tasks classified as being under time pressure (10 to 40 seconds), this thesis further distinguishes three levels of severity using fixed, evenly spaced 10-second intervals: *Mild* pressure (30 to 40 seconds), *Moderate* pressure (20 to 30 seconds), and *Strong* pressure (10 to 20 seconds). This approach follows the standard practice in the experimental time-pressure literature of defining severity using a small number of fixed, researcher-imposed time levels (Payne et al., 1988; Spiliopoulos et al., 2018; Zur & Breznitz, 1981).

3.3.2 Search Behavior and Coping Mechanisms

This section describes how search behavior is measured at the task level, and how these measures are used to test for the presence of the three coping mechanisms introduced in Section 2.1.3: acceleration, filtration, and avoidance (strategy shift).

The analysis proceeds in two steps. First, to test whether each mechanism is present in aggregate, mean levels of the relevant variable are compared across pressure levels using parametric statistics: pairwise t-tests check whether the raw differences between levels are statistically significant, and the same variables are then regressed on *endtime* (Section 3.2), using the following specification:

$$y_{it} = \alpha_i + \beta \cdot \text{endtime}_{it} + \varepsilon_{it}$$

where y_{it} is the search behavior variable for subject i in task t, and α_i is a subject-specific fixed effect. Because each subject completes multiple tasks, including α_i removes any systematic difference across subjects that is constant across their tasks, such as one subject being generally faster or more thorough than another, so that β captures how behavior changes *within* the same subject as the time budget varies, rather than differences between subjects who happened to receive different time budgets. Standard errors are further clustered at the subject level in all regressions, allowing the error term to be correlated across the multiple tasks completed by the same subject rather than assuming full independence across all observations (Cameron & Miller, 2015).

Second, beyond checking whether each mechanism holds on average, this thesis builds a subject-task-level indicator for each mechanism. Each indicator flags tasks where a subject's behavior departs sharply from the No pressure baseline, the condition where time is not binding. For each relevant search behavior variable y_{it}, a z-score is computed relative to the sample mean and standard deviation under No pressure:

$$z_{it} = \frac{y_{it} - \bar{y}_{\text{No pressure}}}{\sigma_{\text{No pressure}}}$$

so that z_{it} measures how many standard deviations a subject's behavior in a given pressured task lies from the average unpressured behavior observed across all subjects, rather than from their own individual baseline. A mechanism is coded as active in a task when this deviation is large enough to be judged an adaptive response, rather than ordinary task-to-task variation, using a $z < -1.645$ threshold (the one-tailed 5% critical value). Tasks under No pressure are coded as zero for every mechanism by construction, since the baseline cannot deviate from itself.

Acceleration. To evaluate whether subjects accelerate their search under time pressure, this thesis examines how quickly subjects gather information. *Average time per look-up* is the mean duration of all look-ups within a task, calculated as the time elapsed between each look-up and the one immediately preceding it. Following Zur and Breznitz (1981), acceleration implies that average time per look-up should decrease as time pressure increases.

$$\text{acceleration\_dummy}_{it} = \mathbf{1}[z_{it}^{avgtime} < -1.645] \quad \text{for pressure\_level}_{it} > 0$$

Filtration. A second, related question is whether subjects narrow the overall scope of their search, rather than simply searching it faster. This thesis measures filtration as the share of all searchable information a subject actually examines within a task. Specifically, *selectivity index* is the number of distinct alternative-attribute pairs inspected, divided by the 72 pairs available in total (eight alternatives, nine searchable attributes each), and therefore ranges from 0 to 1. If subjects filter information under time pressure, selectivity should decrease as the time budget decreases: subjects examine a smaller share of the total information available that they should consider more important, rather than simply distributing a similar amount of search differently across alternatives and attributes.

$$\text{filtration\_dummy}_{it} = \mathbf{1}[z_{it}^{selectivity\_index} < -1.645] \quad \text{for pressure\_level}_{it} > 0$$

Avoidance *(Strategy Shift).* Finally, time pressure may lead subjects to change not just how much they search, but how they organize it. To capture this, this thesis constructs a *search index* (SI), an adaptation of Payne's (1976) measure of interdimensional versus intradimensional search orientation, to characterize the pattern, rather than the amount, of information search. For each transition between consecutive look-ups within a task, a *between-alternative* transition is coded when the alternative (row) changes relative to the previous look-up, and a *within-alternative* transition is coded otherwise. The task-level search index is then calculated as $\text{search\_index}_t = 1 - 2 \cdot \bar{p}_{\text{between},t}$, where $\bar{p}_{\text{between},t}$ is the share of between-alternative transitions within task t. The index ranges from $+1$, when search proceeds entirely within alternatives before moving to the next, to -1, when search proceeds entirely by comparing alternatives on the same attribute before moving to the next one. A shift toward simpler heuristics under time pressure would be reflected in the search index moving toward -1 as the time budget decreases.

$$\text{avoidance\_dummy}_{it} = \mathbf{1}[z_{it}^{\text{search\_index}} < -1.645] \quad \text{for pressure\_level}_{it} > 0$$

3.3.3 Necessary Conditions for Optimal Search

This section describes how each of the four necessary conditions for optimal sequential search introduced in Section 2.1.2 — V, T, R, and C — is operationalized as a violation indicator in the data.

Condition V (Variance). A look-up is eligible for this condition only if the alternative being searched still has at least two unsearched attributes left, if only one remains, there is nothing left to order. Among eligible look-ups, a violation occurs if the subject opens an attribute other than the most informative one still available in that alternative, in other words, skipping ahead to a less informative attribute when a more informative one was still there to be searched.

Condition T (Trade-off). Every look-up except the first in a task is eligible for this condition, since it checks whether a search decision makes sense given what has already been found and the first look-up has nothing to compare against yet. A violation means the subject's next move was the wrong one: either they kept searching an alternative they should have left, or they switched to an alternative that was not actually the best one to switch to. These two cases are split based on what the subject did just before:

- *Within-alternative violations (over-search)*: when the subject was already searching the same alternative, a violation means they should have switched but did not.
- *Between-alternative violations (biased switch)*: when the subject had just switched alternatives, a violation means they switched to the wrong one.

Condition R (Repeat). A look-up violates Condition R if the same attribute box (the same alternative-attribute pair) had already been opened at an earlier point within the same task. Every look-up in a task is eligible for this condition, since any previously searched attribute can, in principle, be reopened.

Condition C (Choice). Unlike the previous three conditions, Condition C is evaluated once per task rather than once per look-up. Every completed task with a recorded final choice is eligible. A violation occurs if the chosen alternative is not the one with the highest expected value revealed by search at the moment the task ends.

The analysis of the necessary conditions proceeds in two steps. First, violation rates for Conditions V, T (and its two subtypes), R, and C are computed for the Exogenous-treatment sample, with standard errors clustered at the subject level, in order to establish the overall rate of compliance and allow comparison with Sanjurjo (2017). Second, each violation indicator is regressed on time budget, with subject fixed effects and standard errors clustered at the subject level, following the same specification described before, in order to test whether the probability of violating each condition is systematically related to time pressure.

Although each violation indicator is binary, these regressions are estimated as linear probability model (LPM) rather than as fixed-effects logit or probit models. Nonlinear panel models with fixed effects suffer from the incidental parameters problem: when the number of tasks per subject is small, as is the case here, maximum likelihood estimates of fixed-effects logit and probit models are biased and inconsistent (Greene, 2004). The linear probability model avoids this issue and its coefficients have a direct interpretation as changes in the probability of violation.

3.3.4 Coping Mechanisms and Search Optimality

The regressions in Sections 3.3.2 and 3.3.3 address whether search behavior and condition violations respond to the time budget, using subject fixed effects to absorb differences across individuals. The question addressed here is different: whether the coping mechanisms themselves are associated with the probability of violating each necessary condition.

Because multiple tasks are observed for each participant, the data have a hierarchical structure in which tasks are nested within subjects. A fixed-effects specification, as used in the previous sections, is not well suited to this question because it gives each subject their own separate intercept, it removes all subject-level variation from the model, and they may differ in their overall tendency to violate the conditions. A mixed-effects (multilevel) model instead treats each subject's baseline tendency to violate a condition as a random effect, drawn from a common distribution rather than estimated separately for each person (Raudenbush & Bryk, 2002). This makes it possible to both estimate the effect of the mechanisms and the time budget, and capture how much subjects differ from one another.

For the conditions evaluated at the look-up level (V, T, R), the data have a further layer of hierarchy: individual look-ups are themselves nested within tasks, and the coping mechanism indicators, which are constructed at the task level, are constant across all look-ups belonging to the same task. Treating each look-up as an independent observation, without accounting for this intermediate level, would understate the correlation among look-ups within the same task and could bias the estimated standard errors. To address this, a second random intercept is included at the task level, nested within subjects, so that the model accounts for both subject-level and task-level heterogeneity. For Condition C, which is evaluated once per task, only the subject-level random intercept is needed, since there is no lower level of nesting to account for.

Because the outcome is binary, the model is estimated as a mixed-effects logistic regression rather than as a linear probability model. This differs from the choice made in Section 3.3.3, where the linear probability model was preferred specifically to avoid the incidental parameters problem that affects *fixed*-effects logit with a small number of tasks per subject. That problem does not arise here: a random-intercept logit model estimates only the parameters of the distribution from which subject-level intercepts are drawn, rather than a separate parameter for each subject, so it remains consistent even when the number of tasks per subject is small.

The model is:

$$\begin{gathered}\operatorname{logit}\left[P(\text{violation}_{ijt} = 1)\right] = \beta_0 + \beta_1 \cdot \text{endtime}_{jt} + \beta_2 \cdot \text{acceleration\_dummy}_{jt} + \\ \beta_3 \cdot \text{filtration\_dummy}_{jt} + \beta_4 \cdot \text{avoidance\_dummy}_{jt} + u_i + v_{j(i)}\end{gathered}$$

where violation_{ijt} is a binary indicator equal to one if look-up t in task j (completed by subject i) violates the condition of interest, and zero otherwise. The term u_i is a subject-specific random intercept and $v_{j(i)}$ is a task-specific random intercept nested within subject i, both assumed to follow independent normal distributions with mean zero. For Condition C, evaluated once per task, the model simplifies to the two-level specification with only u_i, as in the original formulation. A separate model is estimated for each condition.

Beyond their association with condition violations, the coping mechanisms may also be related to subjects actual economic performance. The final step of this analysis examines whether the mechanisms are associated with *score*: the sum of the ten attribute values belonging to the alternative chosen in a task, regardless of how many attributes were actually searched. Unlike the condition violations analyzed, score is a continuous outcome measured once per task, so the model is estimated as a linear mixed-effects regression, with a random intercept at the subject level only, as with Condition C, there is no lower level of nesting to account for, since payoff has no observations below the task level.

The baseline model is:

$$\text{Score}_{it} = \beta_0 + \beta_1 \cdot \text{endtime}_{it} + \beta_2 \cdot \text{acceleration\_dummy}_{it} + \beta_3 \cdot \text{filtration\_dummy}_{it} + \beta_4 \cdot \text{avoidance\_dummy}_{it} + u_i + \varepsilon_{it}$$

which estimates the total association between each mechanism and score, net of the time budget.

To examine whether these associations operate through condition violations, four additional models are estimated, each extending the baseline specification with one task-level violation measure as an additional covariate: the Condition C violation indicator (violation_C_{jt}), and the task-level violation rates for Conditions R, V, and T, denoted $\overline{\text{violation\_R}}_{jt}$, $\overline{\text{violation\_V}}_{jt}$, and $\overline{\text{violation\_T}}_{jt}$, respectively, calculated as the share of eligible look-ups in task j that violate the corresponding condition. Condition T is further examined by separately including the violation rates for its two subtypes, over-search and biased switching.

If a mechanism's coefficient shrinks substantially once a given violation measure is included, this indicates that the mechanism's association with score operates largely through that specific form of search inefficiency, if the coefficient remains stable, the mechanism's relationship with score operates through some other channel.

Declaration of AI Use

During the preparation of this thesis, artificial intelligence tools (Claude, Anthropic) were used in two capacities. First, as a writing aid, to improve the clarity, grammar, and flow of the English prose throughout the document, and to help restructure and simplify explanations of theoretical and methodological content. Second, as a search and verification tool, to locate relevant academic literature and to identify and confirm appropriate methodological approaches (e.g., econometric specifications, standard error corrections, and index construction methods) alongside their corresponding citations.

All research design, data analysis, interpretation of results, and the substantive arguments developed in this thesis are the author's own. AI-assisted searches for literature and methods were used to identify candidate sources and approaches, all of which were independently reviewed and verified by the author before being incorporated into the thesis. The author takes full responsibility for the content, accuracy, and conclusions of this work.

4 Results

4.1 Search behavior under time pressure

This section examines how subjects' search behavior changes as the time available for a task varies, as a first step toward identifying which coping mechanisms are at play. Figure 2 compares mean look-up duration, selectivity index, and search index across the four pressure levels defined in Section 3.3.2, and Table 2 reports pairwise t-tests for each mechanism.

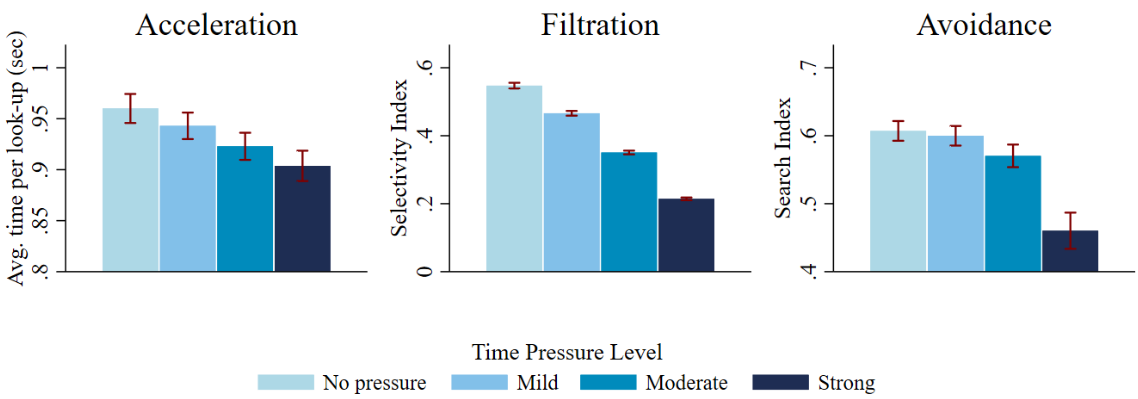

Figure 2: Search Behavior by Time Pressure Level

The first mechanism considered is acceleration: on average, a look-up lasts 0.93 seconds

(SD = 0.24, range: 0.48–3.50), and if subjects speed up under time pressure, this duration should fall as the time available for a task shrinks. The graph shows a decline that is roughly monotonic as pressure increases.

Level (vs. No pressure)	Acceleration	Filtration	Avoidance
Mild	-0.017^{*}	-0.081^{***}	-0.007
Moderate	-0.037^{***}	-0.197^{***}	-0.037^{***}
Strong	-0.056^{***}	-0.333^{***}	-0.147^{***}

Table 2: Pairwise t-test Differences, by Mechanism

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Pairwise t-tests support this pattern, though the effect is not uniform across all comparisons. Differences between non-adjacent levels (No pressure vs. Moderate, No pressure vs. Strong, Mild vs. Strong) are all significant at $p < 0.001$, while two adjacent-level comparisons (No pressure vs. Mild ($p = 0.082$) and Moderate vs. Strong ($p = 0.059$)) are significant only at the 10% level.

Turning to filtration, the proportion of unique boxes visited per task (*selectivity index*) is 0.39 on average (SD = 0.16, range: 0.04 to 1), indicating that, even before accounting for time pressure, subjects leave more than half of the available information unexamined. Figure shows a decline that is both monotonic and substantially larger in magnitude than the one observed for acceleration: selectivity under the strongest time constraint falls to roughly a third of its level under no constraint at all. Pairwise t-tests confirm that this decline is robust across every comparison, with all six pairwise differences significant at $p < 0.001$.

Finally, for avoidance (strategy shift), the search index is 0.56 on average (SD = 0.33, range: -1 to 1), indicating that search proceeds predominantly within alternatives before switching, rather than by comparing alternatives on the same attribute. Unlike acceleration and filtration, the pattern in the figure is not a gradual decline across all four levels, but closer to a threshold effect, the search index is essentially unchanged between No pressure and Mild, and declines only under Moderate pressure, before dropping sharply under Strong pressure.

Pairwise t-tests confirm this threshold pattern directly, the No pressure vs. Mild comparison is far from significant ($p = 0.491$), while every comparison involving Strong pressure is significant at $p < 0.001$. Moderate pressure occupies an intermediate position, differing significantly from both No pressure ($p = 0.001$) and Strong ($p < 0.001$).

Taken together, the three mechanisms display markedly different profiles: acceleration is only partially robust across comparisons, filtration is large and consistent across every level, and avoidance (strategy shift) behaves as a threshold response concentrated in the most severe pressure levels. These unconditional comparisons, however, do not account for heterogeneity across subjects, an issue addressed next through subject fixed-effects regressions.

Table 3 compares these unconditional differences to the corresponding coefficients from a fixed-effects regression of each measure on time budget, with standard errors clustered by subject.

Dependent Variable:	**(1) Avg. Time per Look-up**	**(2) Selectivity Index**	**(3) Search Index**
Time Budget (`endtime_`)	0.00197***	0.01124***	0.00477***
	(0.00022)	(0.00015)	(0.00049)
Constant (`_cons`)	0.87326***	0.05751***	0.41708***
	(0.00647)	(0.00451)	(0.01447)
Observations	4,548	4,548	4,548
Number of Players (Clusters)	390	390	390
Within R^2	0.0301	0.7965	0.0568
ρ (share due to u_i)	0.698	0.593	0.533

Table 3: Fixed-Effects Regressions for Acceleration, Filtration, and Avoidance

Notes: Robust standard errors clustered at the player level are in parentheses.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

The average time spent per look-up increases significantly with the time budget ($p < 0.001$), though the effect is modest in size: each additional second of time budget is associated with only 0.002 additional seconds per look-up. This gap reflects substantial between-subject heterogeneity in search speed, 69.8% of the total variance in look-up duration is attributable to differences across subjects ($\rho = 0.698$), which adds noise to the simple between-group comparison but is netted out once each subject serves as their own control. Time pressure itself explains a comparatively small share of the remaining within-subject variation (within $R^2 = 0.030$).

Unlike acceleration, a fixed-effects regression of selectivity on *endtime* yields a coefficient of 0.0112 ($p < 0.001$), implying that each additional second of time available is associated with a 1.1 percentage point increase in the share of unique boxes visited. This relationship is considerably stronger than the one found for acceleration: 59.3% of the total variance in selectivity is attributable to differences across subjects ($\rho = 0.593$, versus 0.698 for acceleration), and the time budget explains a substantially larger share of within-subject variation (within $R^2 = 0.797$, versus 0.030 for acceleration).

For strategy shift, a fixed-effects regression of the search index on time budget yields a coefficient of 0.0048 ($p < 0.001$), implying that each additional second of time available is associated with 0.0048 increase in the search index. This effect is closer in size to acceleration than to filtration: 53.3% of the total variance in the search index is attributable to differences across subjects ($\rho = 0.533$), and the time budget explains a modest share of within-subject variation (within $R^2 = 0.057$, versus 0.797 for filtration and 0.030 for acceleration).

Taken together, the fixed-effects results confirm the pattern already visible in Figure 2 and Table 2: filtration responds to time pressure far more strongly, and far more consistently across subjects, than either acceleration or strategy shift. The within R^2 values tell largely the same story as the unconditional comparisons, time pressure accounts for the large majority of within-subject variation in filtration, but only a small share for the other two mechanisms, indicating that narrowing the scope of search, rather than processing information faster or reorganizing how it is searched, is the primary channel through which subjects adjust to time pressure.

Beyond these aggregate patterns, the subject-task-level indicators constructed in Section 3.3.2 make it possible to examine how often each mechanism is present at the same time, and whether subjects tend to rely on more than one at once. Table 4 shows how many of the three mechanisms co-occur within a task. The majority of tasks (59.3%) show no mechanism crossing the activation threshold at all, 32.4% show exactly one, and only 8.2% show two mechanisms activated simultaneously, joint activation of all three is essentially absent, occurring in a single task out of 4,548. Rather than a uniform activation of all coping mechanisms at once, the typical pattern is a partial response involving only one or two mechanisms. This activation threshold is set at the

5% tail of each mechanism's distribution under No pressure, Appendix A.1.2 shows that the relative ranking of the three mechanisms is robust to this choice.

No. Mechan.	Freq.	Percent
0	2,698	59.32
1	1,475	32.43
2	374	8.22
3	1	0.02
Total	4,548	100.00

Table 4: Number of Mechanisms Implemented

Table 5 reports how these activations pair up, and the resulting frequencies broadly line up with the magnitudes reported earlier for each mechanism. Filtration, which showed the largest and most consistent effect of pressure on the aggregate measures, is also the most frequently activated mechanism at the task level (37.9% of tasks), followed by strategy shift (9.5%). Acceleration, whose effect on look-up duration was statistically robust but small in magnitude, crosses the stricter activation threshold in only 1.6% of tasks. This suggests that the modest average shift documented above may reflect a small subset of subjects adjusting their pace substantially, rather than a more uniform acceleration across the sample.

Where mechanisms do co-occur, the pairing is similarly uneven. Filtration and strategy shift jointly activate in 7.7% of tasks, considerably more often than either mechanism co-occurs with acceleration (0.6% with filtration, and close to zero with strategy shift). This asymmetry echoes the earlier pattern: filtration and strategy shift, the two mechanisms with the clearest aggregate effects, also appear to be deployed together more often, while acceleration tends to operate largely on its own when it does occur

Figure 3 shows, for tasks in which each mechanism is activated, the density of activation across the time budget. All three mechanisms peak within a narrow window at the most severe end of the range (roughly 13 to 15 seconds), rather than at successively later points as time pressure eases, as the hierarchy suggested by Weenig and Maarleveld (2002) would predict. Filtration's peak is both the earliest and by far the most pronounced, while acceleration and strategy shift track each other closely and peak only marginally later. Beyond the peak, however, the three mechanisms diverge: filtration declines the most steeply, while strategy shift maintains a comparatively higher

	Acceleration = 1	Filtration = 1	Total Mechanism
Avoidance = 1	0.02%	7.72%	**9.48%**
	(1)	(351)	(431)
Filtration = 1	0.55%		**37.89%**
	(25)		(1,723)
Total Mechanism	**1.58%**	**37.89%**	
	(72)	**(1,723)**	

Table 5: Co-occurrence of Coping Mechanisms

relative density through the Moderate range (20 to 30 seconds) before tapering off, and acceleration follows a similar path to strategy shift without a distinct peak of its own. Overall, this pattern offers little support for a strict sequential activation of coping mechanisms, all three become active under severe pressure at roughly the same time.

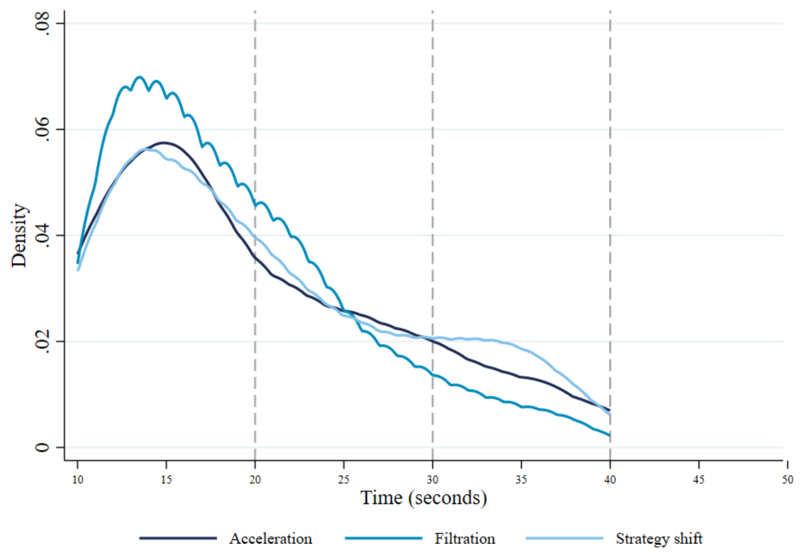

Figure 3: Density of Mechanism Activation by Time Budget

4.2 Compliance with optimal search conditions

Table 6 reports violation rates for the four necessary conditions in the Exogenous-treatment sample, alongside the corresponding figures reported by Sanjurjo (2017) for comparison.

Condition	Violation Rate (%)	95% CI	Sanjurjo (2017) (%)
V (Variance)	5.8	[4.9, 6.9]	5.6
T (Trade-off, aggregate)	49.4	[48.4, 50.4]	48.6
T$\leftrightarrow$ (over-search)	47.9	[46.9, 48.9]	47.2
T$\updownarrow$ (biased switch)	54.3	[53.2, 55.3]	54.9
R (Repeat)	9.7	[9.1, 10.4]	12.0
C (Choice)	36.8	[35.2, 38.5]	30.8

Table 6: Violation Rates of the Necessary Conditions for Optimal Search

Notes: Standard errors are clustered at the subject level. Confidence intervals use a logit transformation to keep them within the valid 0–100% range.

Overall, these results closely replicate the pattern of systematic and substantial deviations from optimal sequential search documented by Sanjurjo (2017): violations are common across all four conditions, with over-search and biased switching remaining the most frequent forms of departure from optimality. The main difference from the original study is a somewhat lower rate of repeated look-ups (9.7% versus 12.0%), which is consistent with the exclusion of the Endogenous treatment from this analysis: when subjects manage their own time budget across problems, they may have more freedom, and more opportunity, to revisit previously searched attributes.

To test whether the probability of violating each condition responds to the time available, each violation indicator was regressed on *endtime*, with subject fixed effects and standard errors clustered at the subject level. Table 7 reports the results.

The clearest relationship is found for Condition R: each additional second of time budget is associated with a 0.44 percentage point increase in the probability of repeating a look-up ($p < 0.001$). This is not necessarily evidence that more time makes subjects less careful, rather, since a longer time budget mechanically allows for more look-ups overall (Section 4.1), it also creates more opportunities for a previously searched attribute to be revisited.

Condition V shows no relationship with the time budget ($p = 0.676$), the already low rate of variance violations does not change systematically as time becomes more or less scarce. Condition C shows a modest, only marginally significant relationship

	(1)	(2)	(3)	(4)	(5)	(6)
	C	R	V	T	T$\leftrightarrow$	T$\updownarrow$
Time Budget	-0.0011^{*}	0.0044^{***}	0.0001	-0.0001	0.0004	-0.0010^{***}
	(0.0006)	(0.0002)	(0.0001)	(0.0003)	(0.0003)	(0.0003)
Constant	0.3997^{***}	-0.0516^{***}	0.0562^{***}	0.4971^{***}	0.4676^{***}	0.5764^{***}
	(0.0185)	(0.0053)	(0.0048)	(0.0087)	(0.0099)	(0.0104)
Observations	4,548	140,306	130,315	140,306	108,031	27,727
Within R^2	0.0007	0.0253	0.0000	0.0000	0.0001	0.0005
ρ (share due to u_i)	0.108	0.043	0.170	0.039	0.043	0.053

Table 7: Effect of Time Budget on Condition Violations

Notes: Robust standard errors clustered at the subject level in parentheses.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
Each column's dependent variable is the corresponding violation indicator (1 = violation). Column (4) reports the aggregate Condition T violation rate; columns (5) and (6) split this into within-alternative (over-search) and between-alternative (biased switch) transitions, respectively.

($p = 0.090$), tighter time budgets are weakly associated with a higher probability of failing to choose the best revealed alternative.

The results for Condition T reveal a pattern that would be missed by looking at the aggregate rate alone, which shows no significant relationship with the time budget ($p = 0.715$). Splitting the condition by transition type shows that this null result masks two offsetting relationships: over-search violations (T$\leftrightarrow$) show no significant association with the time budget ($p = 0.224$), while biased-switch violations (T$\updownarrow$) show a clear, significant one ($p = 0.001$), each additional second of time budget is associated with a 0.10 percentage point reduction in the probability of switching to the wrong alternative. In other words, more time does not make subjects less likely to over-search within an alternative, but it does help them switch to a better one when they do decide to move.

4.3 Mechanisms and Search Quality

The regressions in Section 3.3.3 showed that condition violations respond, at best, weakly to the time budget once considered on its own. This section asks a related but different question: whether the coping mechanisms constructed in Section 3.3.2, beyond the time budget itself, are associated with the probability of violating each necessary condition, using the mixed-effects logistic model described in Section 3.3.4. Table 8 reports the results for all four conditions, with Condition T further split into its two subtypes.

	(1) C	(2) R	(3) V	(4) T	(5) T$\leftrightarrow$	(6) T$\updownarrow$
Time Budget	0.993*	1.110***	1.009*	0.995***	0.995**	0.992***
	(0.004)	(0.004)	(0.005)	(0.002)	(0.002)	(0.002)
Acceleration = 1	0.922	0.996	0.766	0.849	0.845	0.897
	(0.245)	(0.223)	(0.297)	(0.093)	(0.108)	(0.110)
Filtration = 1	0.875	1.994***	1.173	0.780***	0.758***	0.824***
	(0.086)	(0.156)	(0.152)	(0.031)	(0.036)	(0.036)
Avoidance = 1	1.351***	2.704***	1.217	1.481***	1.334***	1.214***
	(0.156)	(0.262)	(0.190)	(0.077)	(0.093)	(0.054)
Constant	0.716**	0.0011***	0.0051***	1.168**	1.117	1.525***
	(0.112)	(0.0001)	(0.0011)	(0.075)	(0.084)	(0.105)
Subject random intercept	Yes	Yes	Yes	Yes	Yes	Yes
Task random intercept	–	Yes	Yes	Yes	Yes	Yes
Observations	4,548	140,306	130,315	140,306	108,031	27,727

Table 8: Mixed-Effects Logistic Regressions: Mechanisms and Condition Violations

Notes: Odds ratios reported, with standard errors in parentheses. All models include a random intercept by subject. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Column (3) is restricted to eligible look-ups (Section 3.3.3); columns (5) and (6) split Condition T into within-alternative (over-search) and between-alternative (biased switch) transitions.

Avoidance (strategy shift) is the only mechanism associated with a higher probability of violation across five of the six conditions, and by a wide margin: odds ratios range from 1.21 for biased switching to 2.70 for repeated look-ups. The exception is Condition V, where the association with avoidance is not statistically significant. Tasks in which subjects shift toward comparing alternatives on the same attribute, rather than exploring one alternative at a time, are consistently more likely to violate most of the necessary conditions.

Acceleration shows no significant association with any of the six conditions. This is consistent with the low frequency of acceleration activation at the task level (Section 4.1): faster processing, on its own, is not systematically linked to search quality in this setting.

Filtration's association with violations points in two different directions depending on the condition. It is strongly associated with a higher probability of violating Condition R (OR = 1.99): visiting fewer alternatives and attributes appears to make subjects more likely to return to information they have already seen. For Condition T, the relationship runs the other way: filtration is associated with significantly lower odds of violation, both for the aggregate condition and for its two subtypes (OR = 0.78 for T, 0.76 for over-search, 0.82 for biased switching), suggesting that a narrower search also comes with fewer sequencing problems. Filtration shows no significant relationship with Condition V or Condition C.

Once the mechanisms are accounted for, the time budget itself retains a significant direct association with four of the six conditions: repeated look-ups increase with more time (OR = 1.11), while the aggregate Condition T and both of its subtypes become significantly less likely as more time becomes available (OR between 0.99 and 0.995). Conditions C and V show only marginally significant direct effects ($p = 0.099$ and $p = 0.089$, respectively), with more time associated with fewer violations for Condition C (OR = 0.993) but more violations for Condition V (OR = 1.009). This indicates that, once mechanisms are accounted for, time pressure continues to operate on search quality partly through a direct channel, beyond the coping mechanisms themselves.

Having shown that the coping mechanisms are systematically associated with violations of the necessary conditions, a natural next question is whether relying on these mechanisms also translates into worse economic outcomes for the subjects who use them. Table 9 reports the results.

In the baseline model (column 1), a longer time budget is associated with a higher score, as expected (0.39 units per additional second, $p < 0.001$). Of the three coping mechanisms, only avoidance is associated with a lower score net of the time budget (-11.81, $p < 0.001$); neither acceleration nor filtration shows a significant association in this baseline specification.

Columns (2) through (7) examine whether these associations operate through the condition violations documented earlier. Adding the Condition C violation indicator (column 2) leaves the pattern largely unchanged: avoidance remains significant, though

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Base	+C	+R	+V	+T	+T$\leftrightarrow$	+T$\updownarrow$
Time Budget	0.388***	0.300***	0.546***	0.388***	0.258**	0.268**	0.360***
	(0.112)	(0.100)	(0.124)	(0.112)	(0.108)	(0.108)	(0.113)
Acceleration = 1	0.486	−0.196	1.193	−0.092	−0.833	−1.211	0.505
	(7.148)	(6.321)	(7.163)	(7.117)	(6.835)	(6.801)	(7.087)
Filtration = 1	0.797	−1.054	1.933	0.679	−5.273**	−4.710*	−0.433
	(2.663)	(2.367)	(2.692)	(2.656)	(2.577)	(2.557)	(2.664)
Avoidance = 1	−11.810***	−7.785***	−10.432***	−10.968***	−3.128	−5.315*	−10.562***
	(3.150)	(2.787)	(3.188)	(3.145)	(3.046)	(3.101)	(3.140)
Violation_C		−54.582***					
		(1.595)					
$\overline{\text{Violation\_R}}$			−28.154***				
			(9.591)				
$\overline{\text{Violation\_V}}$				−23.611***			
				(6.565)			
$\overline{\text{Violation\_T}}$					−78.778***		
					(4.168)		
$\overline{\text{Over-Search}}$						−71.904***	
						(3.552)	
$\overline{\text{Biased-Switch}}$							−13.663***
							(3.255)
Constant	37.515***	60.557***	34.462***	38.878***	81.178***	76.830***	45.714***
	(4.236)	(3.824)	(4.363)	(4.241)	(4.677)	(4.479)	(4.669)
Observations	4,548	4,548	4,548	4,548	4,548	4,514	4,514

Table 9: Mixed-Effects Regressions: Coping Mechanisms and Score

Notes: Coefficients reported, with standard errors in parentheses. All models include a random intercept by subject. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. The dependent variable is Score. Violation_C is coded so that 1 indicates a violation (the chosen alternative was not the best revealed). Columns (2)–(7) each add one task-level violation measure to the baseline specification; columns (6) and (7) split Condition T into its two subtypes and have a slightly smaller sample because a small number of tasks contain no transitions of one type.

its coefficient shrinks by roughly a third, while acceleration and filtration remain non-significant. Conditions R and V (columns 3 and 4) likewise leave the three mechanisms' coefficients close to their baseline values, avoidance stays significant throughout, and acceleration and filtration remain non-significant.

Filtration only emerges as a significant predictor of score once Condition T is considered. Adding the aggregate Condition T violation rate (column 5) makes filtration significantly negative (-5.27, $p < 0.05$), while avoidance's coefficient falls sharply and loses significance. Splitting Condition T into its two subtypes clarifies this pattern: controlling for over-search (column 6), filtration remains marginally significant, and avoidance becomes marginally significant as well, whereas controlling for biased switching (column 7) restores avoidance to a large, highly significant coefficient and returns filtration to non-significance. Acceleration remains non-significant across every specification in the table.

Taken together, these results indicate that avoidance is the mechanism most consistently associated with lower economic performance, an association that persists across nearly every specification and is only substantially attenuated once the aggregate Condition T violation rate is included. Filtration's relationship with score is comparatively fragile, emerging only once search-sequencing violations are taken into account, and acceleration shows no robust association with score in this specification. This pattern echoes the results found for condition violations themselves: avoidance is the mechanism most consistently linked to worse outcomes, both in terms of rule violations and in terms of economic performance.

5 Conclusions

5.1 Summary of Findings

This thesis examined how time pressure shapes multiattribute sequential search, using the Exogenous treatment of Gabaix et al. (2006)'s experiment and the necessary conditions for optimal search introduced by Sanjurjo (2017). The findings partially support, and partially challenge, the three hypotheses guiding this thesis.

Consistent with **H1**, subjects show clear evidence of adapting their search to the time available, though not uniformly across the three mechanisms considered. Filtration responds most strongly and consistently to pressure, declining sharply and monotonically across all four levels. Acceleration also responds significantly, though the adjustment is modest in size. Strategy shift follows a different pattern altogether: rather than a gradual response, it remains flat between No pressure and Mild, and only shifts sharply once pressure becomes Moderate or Strong. Filtration, rather than acceleration, is the dominant response to time pressure in this setting.

H2 is only partly supported. Violation rates are widespread and closely replicate those reported by Sanjurjo (2017), despite this thesis relying only on the Exogenous treatment and a smaller sample. However, once the time budget is considered on its own, its relationship with violations is uneven rather than uniform: more time is associated with a higher probability of repeating a look-up, but with a lower probability of biased switching; Condition V shows no relationship with the time budget, Condition C only a marginally significant one, and the aggregate Condition T no relationship once its two subtypes are combined. The time budget, on its own, therefore explains some, but not all, of the observed violations.

The evidence runs contrary to **H3**. Rather than helping subjects cope with time pressure and reducing violations relative to what the time budget alone would predict, the coping mechanisms are associated with a higher probability of violating most conditions studied. Avoidance (strategy shift) is linked to more violations across five of the six conditions. Filtration shows a mixed profile: it is linked to more repeated look-ups, but to fewer violations of Condition T, without affecting the accuracy of the final choice. Acceleration, by contrast, shows no significant association with any condition, suggesting that faster processing on its own is not a reliable predictor of search quality in either direction. Far from being a purely adaptive response that preserves search quality, coping with time pressure in this setting is, for two of the three mechanisms, associated with moving further away from optimal search.

5.2 Theoretical Implications

The findings of this thesis speak to several strands of the literature reviewed in Chapter 2. First, they reinforce the view that decision makers under time pressure are boundedly, rather than fully, rational: departures from the necessary conditions for optimal search are widespread and systematic, consistent with Simon (1955)'s original account of satisficing and with Sanjurjo (2017)'s characterization of search behavior in this same experimental environment.

Second, this thesis contributes to closing the gap identified in Section 2.2 between the literature on multiattribute search optimality and the literature on coping mechanisms under time pressure. Rather than treating time pressure as an undifferentiated cost that simply worsens search, the results show that its effect on the necessary conditions operates partly through specific, identifiable mechanisms, avoidance and filtration in particular, and partly through the time budget itself, which retains a significant direct association with several conditions even once these mechanisms are accounted for. This suggests that future work seeking to explain departures from optimal search should consider both channels: the particular way in which decision makers choose to adapt to time pressure, and the direct effect of the time budget that operates independently of any single coping mechanism.

Finally, the results complicate the traditionally adaptive framing of coping mechanisms found in (Payne et al., 1988) and related work, which emphasizes that strategy shift and related adjustments can preserve accuracy while reducing effort. In this richer, ten-attribute search environment, avoidance is instead associated with a *higher* probability of violating most of the necessary conditions studied. This does not necessarily contradict the adaptive account outright, subjects may still be behaving reasonably given their cognitive constraints, but it indicates that adaptiveness in this setting does not translate into a uniform or straightforward improvement in search optimality.

5.3 Limitations and Future Research

Several limitations of this thesis point to natural directions for future work. First, the exact time budget for each task was not recorded in the data, so it had to be estimated using *endtime* instead. This proxy was checked against the expected distribution of assigned times and matched it reasonably well (Section 3.2), but it cannot perfectly tell apart the time a subject was given from the time a subject actually used, in cases where they stopped searching early. Future studies could use data where the time budget is recorded directly, to avoid this issue entirely.

Second, the coping mechanism indices are built as an average over each task, so they may hide what happens within a task. For example, a subject who starts searching broadly and only speeds up near the end would look no different, on average, from one who searches at the same pace the whole time. Future work could use the detailed, look-up-by-look-up timing already available in this dataset to study whether these mechanisms kick in gradually as time runs out, rather than being applied the same way from the start.

Third, this thesis only uses the Exogenous treatment of Gabaix et al. (2006)'s experiment, leaving out the Endogenous treatment on purpose, in order to have a cleaner measure of time pressure (Section 3.1). Future research could look at the Endogenous treatment too, where subjects decide for themselves how to split their time across tasks, to see whether the same mechanisms and results still hold when time pressure is self-imposed instead of assigned by the experimenter.

Finally, this thesis only shows that coping mechanisms and condition violations are related, not that one causes the other. It is possible that some other trait of the subject, not measured here, is what actually drives both. Future research could address this by directly manipulating time pressure in a way that isolates one specific mechanism, to test more directly whether it causes worse search outcomes.

References

Bearden, J. N., & Connolly, T. (2007). Multi-attribute sequential search. *Organizational Behavior and Human Decision Processes*, *103*(1), 147–158.

Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, *50*(2), 317–372. https://doi.org/10.3368/jhr.50.2.317

Caplin, A., Dean, M., & Martin, D. (2011). Search and satisficing. *American Economic Review*, *101*(7), 2899–2922.

Ekelund, J., Robert B., & Hébert, R. F. (1997). *A history of economic theory and method* (4th). McGraw-Hill.

Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, *96*(4), 1043–1068.

Gardete, P. M., & Hunter, M. (2024). Multiattribute search: Empirical evidence and information design [Published online]. *Marketing Science*.

Greene, W. (2004). The behavior of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects. *The Econometrics Journal*, *7*(1), 98–119. https://doi.org/10.1111/j.1368-423X.2004.00123.x

Hicks, J. R., & Allen, R. G. D. (1934). A reconsideration of the theory of value. *Economica*, *1*(1), 52–76.

Inoua, S. M., & Smith, V. L. (2020). Classical economics. *The Independent Review*, *25*(1), 79–90. https://www.jstor.org/stable/48583695

Johnson, E. J., Payne, J. W., Bettman, J. R., & Schkade, D. A. (1989). *Monitoring information processing and decisions: The Mouselab system* (Technical Report No. TR-89-4-ONR). Center for Decision Studies, Duke University. Durham, NC.

Keeney, R. L., & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs*. John Wiley & Sons.

Klimm, F., Kocher, M. G., Opitz, T., & Schudy, S. (2023). Time pressure and regret in sequential search. *Journal of Economic Behavior & Organization*, *206*, 406–424. https://doi.org/10.1016/j.jebo.2022.12.019

Maule, A. J., & Hockey, G. R. J. (1993). State, stress, and time pressure. In O. Svenson & A. J. Maule (Eds.), *Time pressure and stress in human judgment and decision making* (pp. 83–101). Springer.

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance*, *16*(2), 366–387. https://doi.org/10.1016/0030-5073(76)90022-2

Payne, J. W. (1982). Contingent decision behavior. *Psychological Bulletin*, *92*(2), 382–402. https://doi.org/10.1037/0033-2909.92.2.382

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(3), 534–552. https://doi.org/10.1037/0278-7393.14.3.534

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (2nd). Sage Publications.

Reutskaja, E., Nagel, R., Camerer, C. F., & Rangel, A. (2011). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, *101*(2), 900–926. https://doi.org/10.1257/aer.101.2.900

Ricardo, D. (1817). *On the principles of political economy and taxation*. John Murray.

Robbins, L. (1932). *An essay on the nature and significance of economic science*. Macmillan.

Sanjurjo, A. (2017). Search with multiple attributes: Theory and empirics. *Games and Economic Behavior*, *104*, 535–562.

Sanjurjo, A. (2023). *Complexity in choice* (Working Paper No. 4580158). Social Science Research Network (SSRN). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4580158

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118. https://www.jstor.org/stable/1884852

Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations*. W. Strahan; T. Cadell.

Spiliopoulos, L., Ortmann, A., & Zhang, L. (2018). Complexity, attention, and choice in games under time constraints: A process analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(10), 1609–1640. https://doi.org/10.1037/xlm0000535

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, *69*(3), 213–225. https://www.jstor.org/stable/1831147

Thaler, R. H. (2016). Behavioral economics: Past, present, and future. *American Economic Review*, *106*(7), 1577–1600. https://doi.org/10.1257/aer.106.7.1577

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.

Wallsten, T. S. (1993). Time pressure and payoff effects on multidimensional probabilistic inference. In O. Svenson & A. J. Maule (Eds.), *Time pressure and stress in human judgment and decision making* (pp. 83–98). Springer.

Weenig, M. W. H., & Maarleveld, M. (2002). The impact of time constraint on information search strategies in complex choice tasks. *Journal of Economic Psychology*, *23*(6), 689–702. https://doi.org/10.1016/S0167-4870(02)00135-4

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, *47*(3), 641–654.

Zur, H. B., & Breznitz, S. J. (1981). The effect of time pressure on risky choice behavior. *Acta Psychologica*, *47*(2), 89–104. https://doi.org/10.1016/0001-6918(81)90001-9

Appendix

A Additional Tables and Figures

A.1 Robustness Checks

A.1.1 Validity of *endtime* as a Proxy for the Time Budget

Figure A.1 compares the empirical distribution of *endtime* against the uniform distribution over 10 to 49 seconds originally used to assign time budgets in the Exogenous treatment. The empirical distribution closely tracks the theoretical uniform density across most of the range, with a small number of spikes exceeding the uniform benchmark. This pattern supports the use of *endtime* as a reasonable proxy for the time actually allotted to each task.

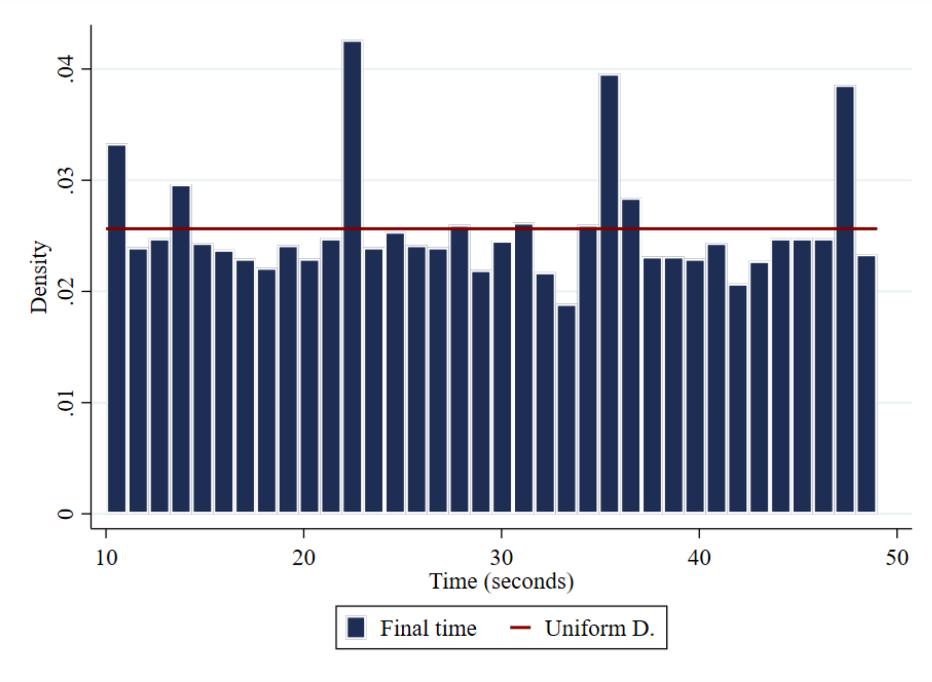

Figure A.1: ENDTIME VS. UNIFORM DISTRIBUTION (10–49S)

A.1.2 Sensitivity to the Activation Threshold

As a robustness check, the same z-score procedure was repeated using a range of alternative activation thresholds, from a stricter 1% cutoff to a more lenient 20% cutoff, rather than the 5% threshold used throughout the main analysis. Table A.1 reports the resulting activation rate for each mechanism.

Threshold (z-cutoff)	Acceleration	Filtration	Strategy Shift	≥2 mechanisms
1% ($z < -2.326$)	0.0%	18.7%	6.8%	3.8%
5% ($z < -1.645$)	1.6%	37.9%	9.5%	8.2%
10% ($z < -1.282$)	7.3%	46.4%	11.6%	12.7%
15% ($z < -1.036$)	14.7%	54.3%	13.7%	18.8%
20% ($z < -0.842$)	20.9%	58.8%	15.5%	24.6%

Table A.1: Sensitivity of Mechanism Activation Rates to the Threshold

Notes: Each cell reports the share of tasks in which the corresponding mechanism indicator equals 1, using the z-score threshold in the row. The main analysis uses the 5% threshold ($z < -1.645$). The last column reports the share of tasks with two or more mechanisms simultaneously active.

The relative ranking of the three mechanisms is stable across every threshold considered: filtration is the most frequently activated mechanism at every cutoff, followed by strategy shift, with acceleration consistently the least frequently activated. This ordering matches the pattern found in the aggregate regressions (Section 4.1), where filtration showed the largest and most consistent effect of pressure, and acceleration the smallest.

The three mechanisms differ, however, in how sensitive their activation rate is to the threshold itself. Acceleration is the most sensitive by far: at the strictest 1% cutoff, no task in the sample crosses the threshold, while at the most lenient 20% cutoff, over a fifth of tasks do. Filtration is comparatively stable in relative terms but becomes a majority classification well before the most lenient threshold is reached, crossing 50% of tasks between the 10% and 15% cutoffs. Strategy shift is the most stable in absolute terms, rising gradually from 6.8% to 15.5% across the full range of thresholds considered.

This pattern supports the 5% threshold used in the main analysis: it is lenient enough to activate acceleration in a non-trivial share of tasks, while remaining strict enough that filtration does not yet represent the majority of the sample, a boundary crossed only at more lenient cutoffs. The substantive conclusion that filtration is the most common and consistent coping mechanism, followed by strategy shift, and with acceleration the least frequently activated, holds regardless of which threshold in this range is used.