

Coarse and Precise Information in Food Labeling

JOB MARKET PAPER

[Click here for the latest version](#)

Silvio Ravaioli[†]

November 20, 2021

Abstract

Public authorities and companies often adopt simple categorical labels to convey information and promote the purchase of healthy, ethical, or environmentally-friendly products. Why are these “coarse” labels favored over more detailed ones which should allow the consumer to make better decisions? This paper investigates whether precise labels can be more effective and informative than coarse ones. In a preregistered online study conducted on a representative US sample, I manipulate front-of-package labels about foods’ calorie content. I find that coarse-categorical labels generate a larger reduction in calories per serving compared to detailed-numerical labels despite providing less information (-3% and -1% calories, respectively). Results also show that participants prefer coarse labels. Choices violate the predictions of Bayesian decision theory, suggesting that consumers are less responsive to detailed information. A bounded rationality model with precision overload can capture the main experimental results: detailed labels are more complex and harder to understand, and consumers face a tradeoff between simplicity and precision. Some information helps, but too much detail can be confusing, and lead to less healthy food choices.

Keywords: Consumer Behavior, Experiment, Attention, Nutrition, Food Labels.

JEL codes: C91, D83, D91, Q28.

[†] Department of Economics, Columbia University, 420 West 118th Street, 10027 New York. E-mail address sr3300@columbia.edu. Many thanks to Mark Dean and Michael Woodford for the invaluable advice and support on this project. I also thank Alessandra Casella, Eric Johnson, Jacopo Perego, Bernard Salanié, the participants at Columbia University’s Economics Department Colloquia, and the member of the Cognition and Decision Laboratory for the valuable feedback. This project has received funding from Columbia University’s Program for Economic Research and Columbia Experimental Laboratory for Social Sciences. All data were collected with the approval of the Columbia University Institutional Review Board (protocol AAAT5268). The experimental design, hypotheses, and sample size were preregistered on AER Social Science Registry under the name “Coarse and Granular Nutritional Labels” (trial AEARCTR-0007856).

1 Introduction

Coarse labels — i.e., labels that indicate that a product’s attribute falls within a range of values — are widespread. You might have started your day by grabbing a low-fat yogurt (how many grams of fat?) from an energy-efficient refrigerator (how many kilowatt-hours?) while you drink a fair-trade coffee (how much was the farmer paid?). Public authorities and private companies often adopt these simple categorical labels to convey information and promote the purchase of healthy, environmentally-friendly, or ethically produced goods.¹ Food labels across the world provide examples of coarse messages, some of which are in Figure 1: warnings in Chile (binary messages), traffic lights in Ecuador (three colors), and Nutri-Score in France (five grades), among others. Food labels can have a significant impact on consumer behavior. For example, Barahona et al. (2020) estimate that the introduction of warning labels in Chile in 2016 led to a 6.5% reduction in consumed calories.



Figure 1: Examples of coarse nutrient labels used in different countries. From the left: Warning label (Chile), Traffic light (Ecuador), Nutri-Score (France), Guiding Stars (US), Health Star Rating (Australia). The first two labels refer to the amount of specific nutrients (e.g., calories and sugar), the other three provide a summary based on multiple nutrients.

In the examples above, the label designers selected a small number of possible categories and the criteria for assigning a category based on the product characteristics. Why did they not use more categories? In principle, coarse labels could provide the same information *and more* by refining each category further. For example, a nutrient label with five grades, from A to E,² could become more informative³ for the consumers by employing a richer scheme that partitions the A grade into three sub-grades (A+, A, and A-). In standard information models, higher precision allows consumers to be more informed and make better choices based

¹The effect of labels on consumer behavior has been studied in various contexts, including food choice (Ikonen et al., 2020; Crosetto, Muller and Ruffieux, 2016), energy efficiency (Davis and Metcalf, 2016; Andor, Gerster and Götte, 2019), and fair trade (Hainmueller, Hiscox and Sequeira, 2015; Rousseau, 2015). Studies in these areas typically compare the effectiveness of labels that differ in multiple dimensions, including coarseness.

²The Nutri-Score is an example of a coarse label with five grades from A to E used to rate the healthiness of food products. The grade is obtained through an algorithm that considers multiple nutrients and generates a raw score between -15 (healthiest) and +40. This raw score is coarsened into five grades, from A (-15 to -1 points) to E (19 to 40 points). The Nutri-Score has been officially recommended by health authorities in France, Germany, Switzerland, and other European countries.

³A finer partition of the signal space is more informative in the Blackwell sense (Blackwell, 1953).

on their preferences. The underlying assumption is that consumers can coarsen information without loss. Under this assumption, the use of coarse labels might be justified by regulators’ persuasion motives or manufacturers’ strategic response. Regulators could attempt to manipulate the consumers’ choices by pooling together desirable and undesirable products.⁴ Coarse labels could also affect manufacturers’ response through price changes and product reformulation.⁵ However, there is an additional reason why coarse labels may be adopted, absent strategic considerations: consumers may simply find detailed labels confusing and difficult to interpret, thereby making them less effective than coarse labels.

In this paper, I study the direct effect of coarse-categorical and precise-numerical calorie labels on consumer behavior and the related underlying mechanisms. I conduct an experiment with incentivized food choices and vary the characteristics of the label to answer three questions.

First, how do coarse and precise labels affect calorie content of food choices? I find that increasing the number of categories in the label has a U-shaped effect on calories. Coarse-categorical labels reduce the number of calories by up to 3% relative to having no labels, but detailed-numerical labels have a weaker effect, only 1% less. I also find that, when details labels are used, choices are less responsive to the calorie content of the product.

Second, are the results consistent with Bayesian decision theory, the standard approach in decision making? In my analysis, I estimate the effect of labels on choices in a simple demand model, and reject Bayesianism. Results suggest that, despite more information contained in the detailed labels, consumers may extract less information from them.

Third, how can the standard model be modified to capture the results? I show that the experimental results cannot be characterized using standard approaches of limited rationality, such as salience and rational inattention. To address these shortcomings, I introduce a limited attention model with *precision overload* and show that it is able to capture the main effects. Consumers with limited cognitive resources face an explicit tradeoff between simplicity and precision: detailed messages are more complex than coarse ones and are associated with larger noise in the evaluation process. The effect can be modeled by assuming that agents face two separate types of attention costs for recognizing and retrieving information accurately.

⁴In the Bayesian persuasion literature (e.g., Kamenica and Gentzkow (2011); Bloedel and Segal (2018)) principal and agent have misaligned incentives, and the principal can design a signal structure that induces the agent to make choices more aligned with the principal’s goal. In food choices, the regulator could assign a warning label to high- and medium-calorie products to promote the choice of low-calorie products only.

⁵Barahona et al. (2020) show that, after the Chilean label reform, producers of breakfast cereals with warning labels changed the recipe and reduced the amount of calorie and sugar content to be just below the warning threshold.

In the experiment, participants are randomly assigned to treatments that differ in level of calorie labels’ granularity.⁶ I refine coarse labels into more detailed ones by dividing each category into a more granular partition. In the main part of the study, participants choose between either two or four products which appear side-by-side in an interface similar to the one typically used for online groceries. Using a preregistered design, I collect data on food choice, preference for information, and consumer characteristics. The study is conducted online on a representative US sample (n=856), using products broadly available in US stores.

Conducting an experiment to study the effect of labels has three significant advantages. First, the experimental manipulation provides a clean randomization of the treatment. Empirical studies that consider policy implementations are limited by the successfully implemented labels and can compare only across time (before and after the reform) or countries (with and without a label). Second, the experimental setup provides additional control with respect to field data. Participants in the study need to read the instructions and answer the comprehension questions correctly before proceeding with the main task. I provide contextual information about the typical range of calories and how to interpret the calorie labels. This makes it possible to minimize concerns about the nature of the information conveyed through the labels. Third, by collecting choices and process data I am able to study how participants acquire information, particularly if they consult the nutrition panels by flipping the packages.

The main experimental result is that coarse-categorical labels are more effective than detailed-numerical ones in reducing the calorie content of food choices. The coarse label with four categories (ranging from *very low* to *very high* calories) leads more often to the choice of products with lower calories compared to treatments with more or fewer categories. I decompose the label effect and observe that detailed labels are less effective for large calorie differences across products. When faced with detailed labels, participants are less sensitive to the information about the calorie content. The result is not driven by search (consulting the back-of-package) nor by taste (disliking low-calorie products). The magnitude of the effect is heterogeneous across demographic groups, with a larger impact across low-income and low-education groups. Participants also indicate that they prefer coarse labels to detailed ones. When asked to rank various types of labels, the coarse label is the most popular first choice and the Condorcet winner in the pairwise comparisons.

As a first step in understanding this effect, I estimate a simple structural demand model and I test whether subjects behave in line with Bayesian decision theory, the standard approach in choice under uncertainty. The premise of the structural model is that consumers

⁶I focus on one nutrient, calories, that is already available on the back-of-package as an exact amount, and I only manipulate the front-of-package label.

care about nutritional content and other characteristics of food products, but calories are not directly observed. Under Bayesianism, labels help consumers by providing a signal about the calorie content. More granular labels reduce the uncertainty and generate more accurate posterior beliefs about calories.⁷ In the context of the structural model, I estimate weights representing the informativeness of the labels. The central prediction of the Bayesian model is that coarse labels are associated with less extreme, smaller weights compared with detailed labels.⁸ However, I find that the estimates from the experimental dataset show the opposite pattern, with coarse labels having systematically more extreme weights, at odds with the model’s predictions.

To capture the main results of the experiment, I extend the choice model by introducing limited attention. Classic behavioral models like salience (Bordalo, Gennaioli and Shleifer 2012) and rational inattention (Sims 2003) are not able to capture the main experimental results. The precision overload model is motivated by existing evidence on information overload (Roetzel, 2019), and cognitive limits in perception and memory (Pollack, 1952; Miller, 1956; Cowan, 2001).⁹ Following closely the definition introduced by Roetzel (2019) for information overload, I define *precision overload* as a state in which the precision of the information inhibits the decision maker’s ability to determine the best possible decision. In lay terms, decision makers could find the detailed information so confusing that they make worse use of it compared to the case in which they were receiving a coarser message. Differently from the choice overload effect (Iyengar and Lepper, 2000), precision overload refers to the degree of granularity of information, not to the number of options available in a choice set.¹⁰

I model precision overload by assuming that the agent has a noisy understanding of the message conveyed by the label, with the level of noise depending on the granularity of the message. The variable noise emerges in a setup in which the agent has limited cognitive resources available to process label information. These resources are consumed by two types of attention costs — a recognition cost and an retrieval cost. The novel feature of the model is the introduction of the recognition cost, which is incurred to recognize the labels

⁷As it is the case in the experiment, I assume that detailed labels are a partition of coarse ones and thus more informative in the Blackwell sense.

⁸More precisely, coarse labels give weights that can be expressed as convex combinations of those associated with detailed labels.

⁹In a series of studies on the ability to classify auditory stimuli, Pollack (1952) reports that subjects can discriminate up to five tones with perfect accuracy, but accuracy decreases when the number of tones is larger. The interpretation the author offers is that decision makers have limited processing capacity, for example, facing a constraint in Shannon’s entropy reduction of the signal (Shannon, 1948; Norwich, 1993).

¹⁰Choice overload can emerge when consumers are Bayesian agent and face either search costs, as in Kuksov and Villas-Boas (2010), or are uncertainty about own preferences and use the choice set to make inference about them, as in Kamenica (2008).

and depends only on the complexity of messages. After some resources are exhausted to recognize the labels, the remaining resources are allocated toward the retrieval. The retrieval cost depends on how accurately the label information is retrieved from memory. As in the rational inattention literature, more accurate beliefs are costlier. Intuitively, while detailed-numerical labels contain more information than coarse-categorical labels, it may be more difficult to extract. Hence, simplifying labels may lead to less noisy and better choices.

This paper contributes to the vast and interdisciplinary literature studying the effect of food labels on consumer behavior using laboratory experiments (Bix et al., 2015; Ono and Ono, 2015; Crosetto et al., 2020), field experiments (Dubois et al., 2020; Shangguan et al., 2019), and reforms (Barahona et al., 2020; Sandoval, Carpio and Sanchez-Plata, 2019). This is the first paper that compares coarse and detailed food labels designed as partitions, tests whether more detailed messages lead to low-calorie choices, and shows that excessive precision can backfire after controlling for other factors (contextual information, process data). The results are connected to the empirical literature on limited consumer attention, which provides evidence that consumers do not use all the information available and make costly mistakes (Chetty, Looney and Kroft, 2009; Brown and Jeon, 2019). My study highlights that mistakes can be associated with an excessive granularity of information, similarly to what has been discussed for choice overload (Iyengar and Lepper, 2000) and information overload (Roetzel, 2019). The adoption of coarse messages in communication problems relates to a broad theoretical literature. Coarsening can emerge in principal-agent settings in which incentives are not aligned (Crawford and Sobel, 1982; Kamenica and Gentzkow, 2011). I show that another possible explanation is that agents face larger noise when they process detailed information and cannot fully absorb the content of the message. This framework is analogous to the rational inattention approach (Sims, 2003; Caplin and Dean, 2015) with the addition of an explicit recognition cost that depends on the complexity of the message and not on how it will be used in the decision process.

To summarize, the experimental results highlight that coarsening can be desirable both for consumers and regulators, and coarse labels can be more effective than detailed ones in generating publicly desirable behavior, such as promoting a healthier diet with fewer calories. The effect of food labels is relevant both for policymakers and companies who want to inform consumers and promote the purchase of healthier products. Mandatory and standardized food labels have been proposed as tools to reduce obesity and diet-related chronic diseases,¹¹ and as less controversial alternatives to product bans or sugar taxes.¹² This study points out

¹¹72% of the US population (Center for Disease Control and Prevention, 2018) and 40% of the world's adult population (World Health Organization, 2018) is overweight or obese. The estimated annual medical cost of obesity in the US in 2008 was \$147 billion, about \$500 per person.

¹²Allcott, Lockwood and Taubinsky (2019) provide an extensive discussion on sugar taxes and why they

that the success of a food labeling reform might also depend on the coarseness used to provide nutrient information. Precision overload provides a compelling mechanism to motivate the large use of coarse information: even a benevolent information designer, whose interests are aligned with the consumer, might want to simplify the labels' information by reducing their granularity to make them less confusing. The experimental results refer to food choices and calorie information, but precision overload might be relevant in many other decisions that involve imperfect attention to information. Another relevant aspect is the heterogeneity in the response to the label information. I observe that the effect of coarse labels is relatively stronger among consumers from low socioeconomic status, suggesting that regulators should consider how different consumers respond to intervention. Some information helps, but too much detail can be confusing and lead to less healthy food choices.

The remainder of this paper is organized as follows. Section 2 describes the experimental design, hypotheses, and data collected in the online study. I present the main results of the experiment in Section 3, including search data and preference over types of labels. In Section 4, I present and estimate a choice model in which consumers infer the product nutrient from the label. Section 5 generalizes the model to include precision overload. I discuss alternative interpretations of the findings in Section 6. Section 7 contains further discussion of the relevant literature and Section 8 concludes.

2 Experiment

I use a field-framed online study with incentivized choices between food items. In the first part of the session, subjects choose a product (cereals or snacks) from each choice set, with an interface that mimics an online grocery store. Products display front-of-package labels that indicate the calorie content of the product. Subjects are randomly assigned to one out of four treatments, that correspond to different labeling regimes. The labeling regimes are nested, with more granular labels defined through a partition of the coarser ones. In the second part of the session, subjects answer additional questions about label preferences, calorie preferences, and food habits.

may be regressive.

2.1 Experimental Design

2.1.1 Choice Task

In the product choice task, subjects choose one food item among those presented on the screen. Figure 2 show an example of choice between two cereals. The task is incentive compatible, as subjects have a chance of receiving one of the selected products directly at home (details in the Recruitment section). Subjects complete 80 trials of the task at their own pace. The interface contains the same elements that customers would see in an online platform: an image of the package, a brief description, and the price of the product. By clicking on the buttons “front” and “back,” subjects can flip the package and read the nutrition facts. Subjects are not required to consult the nutrition facts in the task to make a choice.¹³

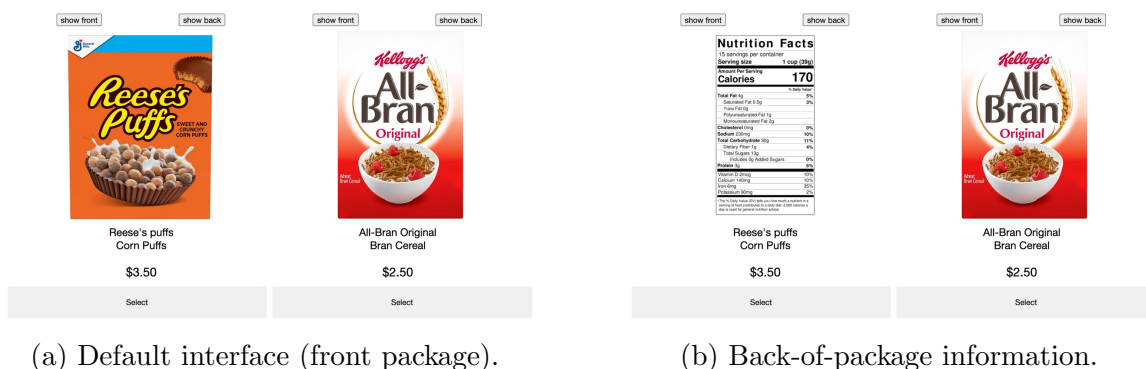


Figure 2: Experimental interface for the choice task (control treatment, Absent label). Examples of information displayed on the front and back of the package. Subjects can freely consult the nutrition facts by clicking the buttons above each product.

2.1.2 Treatments

Each subject is assigned to one out of four experimental treatments, that differ in the type of front-of-package labels used to provide information on the calories per serving. The products used in the study have a number of calories per serving that range between 60 and 240 Calories. The four treatments have different number of labels and therefore different partitions of the calorie range. The four treatments are displayed in Figure 3 and are:

- Absent (A): no label.

¹³All the subjects are made aware of the “flip the package” feature in the instructions. In a practice trial, with a slightly different interface, they are not allowed to proceed until they have consulted the information on the back of the package.

- Binary (B): 2 categorical labels (Low calorie, High calorie).
- Coarse (C): 4 categorical labels (Very low calorie, Low c., High c., Very high c.).
- Detailed (D): 10 numerical labels (3% recommended c. Daily Value, 4% DV, ...).

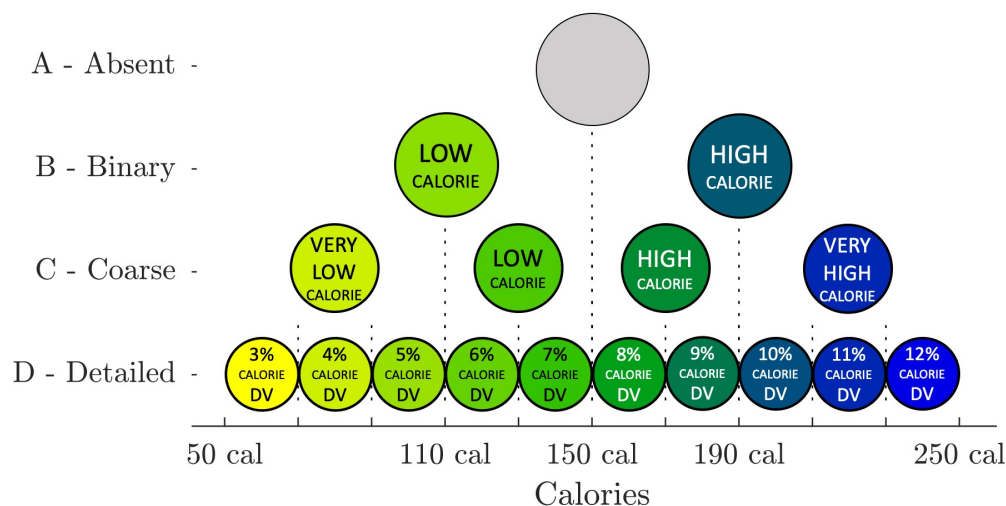


Figure 3: Treatments. Labels used in the four treatments and corresponding range of calories values. Labels are expressed as a combination of color (yellow to blue) and message (categorical for treatments B and C, numerical for treatment D).

The message on the label is presented on a colored background, ranging from yellow (for low calories) to blue (for high calories). I avoid the classic *traffic light* color palette,¹⁴ and red in particular, to minimize the framing of high-calorie products as dangerous or undesirable. We know from other studies Schuldt (2013); Shen, Shi and Gao (2018) that labels' colors affect their effectiveness, with red and black being typically associated with warning messages.

Instructions and general explanation of the meaning of labels are unchanged across treatments, control treatment included. The instructions include the information that the number of labels is randomly assigned in the study, and explanations on the recommended daily values, the relation between DV percentages and number of calories, and other resources available to understand nutrition facts.

Packaged food often display information about nutrient fact on the front of the package, in the form of claims (e.g., low-fat) and exact value (e.g., facts up front). In order to minimize concerns about interaction between labels and existing information, I manually remove all of

¹⁴Traffic lights colors (green/red) are commonly used to express recommendations in the form of go/no go binary messages. Yellow and intermediate shades of green and red are also used to indicate intermediate levels. Nutri-score labels, for example, have five labels ranging from A/green to E/red.

them from the front of the package. Figure 4 shows an example of how the original package is edited by removing the existing nutrient claims and adding a calorie label according to the treatment and the true calorie content of the product.




Figure 4: Product front package manipulation. The original pictures of the packages (left) are edited to remove existing nutrient claims and obtain a clean version of the package (center). The calorie label, if any, is added to the package according to the treatment and the true calorie content of the product (right).

2.1.3 Instructions and Meaning of the Labels

The instructions of the experiment contain information about the context and the meaning of the labels. The general instructions show the distribution of calories across products according to the USDA National Nutrient Database for Standard References. For cereals and snacks over 95% of the products fall between 60 and 240 calories per serving, and mostly between 100 and 200, as displayed in Figure 5.a.¹⁵ The general instructions show the probability distribution of calories, explain how to interpret the calories amount in terms of recommended daily values, and introduce the concept of calorie labels to provide information about the range of values for the product. The comprehension questions test whether the subjects remember and test this information. In order to proceed with the study, they need to answer the questions correctly and indicate the range of calories for the products in the study, both as numerical value (60 to 240 calories) and percentage of recommended daily value (3 to 12%). After the random assignment to one of the treatments, the same distribution is shown again with the label mapping that is implemented in the rest of the

¹⁵The products used in the study provide an approximation of the empirical distribution, with more products in the medium-calorie range.

study. Figure 5.b shows the information for the Detailed treatment. Subjects can consult again this information at any time during the experiment, as it is available in the choice task page.



(a) Calorie distribution (general instructions).

(b) Calorie distribution (Treatment D)

Figure 5: Explanation of the label meaning in the general instructions (left, example with three categories) and Detailed treatments (right).

2.1.4 Features of the Experimental Design

The experimental design has two features that are worth discussing before proceeding further: subjects have the possibility to consult back-of-package information, and detailed labels contain numerical messages. I want to clarify why I chose these features and what are the advantages compared to alternative versions of the same task.

Subjects can consult the back-of-package, but an alternative approach would have been forcing them to choose based on front-of-package information only. This feature is motivated by a stronger test of the precision overload effect. Suppose the effect exists but consumers can overcome the confusion that arises from the precise labels by consulting the back-of-package: then consumers might not be worse off under precise labels, as they are just substituting front-package with back-package information. As a consequence, I would be able to find precision overload as a result of the experiment but this would not be a real concern in the real world. Instead, the fact that I observe the effect in an experiment in which subjects have the possibility to search further is a stronger evidence that consumers might be more confused, but not take further actions to address this issue. The free search feature has two additional advantages. First, it allows to collect process data (what do consumers consult before making the choice) and test secondary hypotheses (how do labels affect the search behavior). Second, this provides a more natural choice environment; online and in-store grocery experience allows to collect information about the product before making choices.

Detailed labels in the experiment contain numerical messages (e.g., 3% calorie Daily Value) whereas coarser labels display categorical messages without numbers (e.g., low calorie). A possible concern is that any effect of precision from this comparison is due to the type of message and not to the granularity by itself. A limitation of the current design is that it does not aim to separate whether the effect of precision is uniquely due to one of these two aspects. The experiment allows to test the joint hypothesis that detailed-numerical labels lead to worse (calorie-wise) choices compared to coarse-categorical labels: in everyday experience precise information is in fact typically associated with numerical values. An alternative experimental design could use only labels with the same type of messages: for example, coarse-numerical label that indicates the calorie range (3-7% calorie DV) or detailed-textual labels with higher precision (very-very-very-low calorie). Both these approaches rise a different type of concerns because they are dissimilar to labels we can encounter in the real world, and can lead subjects to behave differently because of lack of familiarity. The experimental design addresses the more general concern that subjects would not be able to interpret or compare the numerical information. Numbers are integers in the 3-12 range, instructions contain information about the calorie range (also tested in the comprehension questions), and the color of the label indicates its meaning in all the treatments.

2.1.5 Implementation

Instructions and Comprehension Questions. After signing the consent form, subjects are walked through the general instructions for the study (types of tasks, length of the study, incentives) and detailed instructions for the choice task. The detailed instructions include practice rounds to familiarize with the interface, details on the type of products, explanations of the labels that might be presented in the study,¹⁶ and of the meaning of the recommended daily value of calories according to FDA guidelines.¹⁷

At the end of the instructions, subjects are presented with three comprehension questions on the nature of the task (choose between food items), the incentives (receive one of the chosen products), and the range of calories for the products used in the study (between 3% DV [60 Calories] and 12% DV [240 Calories]). Subjects cannot start the task until they answer

¹⁶Subjects in all the treatments are presented a trinary label scheme as example (low, medium, and high calorie). They would not be presented that specific label scheme, but whatever they will see has similar properties (e.g., yellow-blue palette).

¹⁷The instructions include information on the 2,000 daily calorie guideline (and how the exact recommendation varies based on personal characteristics and behavior), the use of daily value percentages (e.g., 5% to indicate 100 calories), and an explanation that all the nutrition information in the study are expressed in calories per serving. Subjects can also use the links provided to visit the FDA website and learn more.

correctly all the questions.¹⁸ The latter question tests whether subjects remember what is the lowest and highest amount of calories they can expect to find among the products in the study. A summary of the instructions is also available during the main task by clicking on the instructions button.

Products and Calories Distributions. The study contains choice sets comprised of two or four items. The products used are cereals and salty snacks that can be typically purchased in grocery stores. Within each trials, all the product belong to the same category. These product categories have the advantages of being largely popular food (Hamrick, 2016; Rhodes et al., 2017) and displaying a large heterogeneity in the nutrient characteristics. Using publicly available data of the US Department of Agriculture, I extract the nutritional information on calories per serving and I include that in the study as part of the instructions (to explain calorie heterogeneity across products) and task (to explain how each label is associated with a particular range of calorie values). I truncate the distribution at 60 and 240 Calories per serving (a range that includes more than 95% of the observations for each category).¹⁹

Randomization. Randomization is operated at the subject level and at the trial level. Subjects are randomly assigned a treatment and an order of the blocks of choice trials. Trials are grouped into blocks to avoid that the same item appears several times at short distance, and to maintain the same choice set size for several rounds in a row.²⁰ For each trial, I randomize the position of the products on the screen to remove position concerns. All subjects face the same choice sets, and the choice sets are designed with desirable properties for the subsequent analysis. First, the marginal calorie distribution approximates the empirical one shown in Figure 5, with mid-calorie levels being more frequent than extreme ones. Second, choice sets are heterogeneous in the combinations of calories of the products: calorie content is similar in some cases and very different in others.²¹ Each choice set is presented twice to each subject with different prices, once with the higher price for the high-calorie products, and once with the higher price for the low-calorie products.

¹⁸The main results of the study are qualitatively and quantitatively robust if I restrict the analysis to the subjects that pass the comprehension questions at the first attempt.

¹⁹The USDA National Nutrient Database for Standard References can be found on the [USDA website](#).

²⁰There is no significant interaction between order of appearance and healthy choices. Subjects are equally likely to choose a healthy product in early and late trials of the study.

²¹For choice sets with two products I group the products in four quartile levels according to the coarse labels, and I generate ten combinations with replacement of two levels out of four. For choice sets with four products I generate choice sets that are balanced or skewed towards low or high calories.

Programming. The study is conducted using the online platform Qualtrics and custom JavaScript code. Qualtrics does not have specific features for mousetracking, and other features, for example the ranking questions, have limited flexibility. I developed various JavaScript tools that facilitate product visualization and data collection and can be used for a variety of online studies.

2.1.6 Additional Data

After the choice task, the subjects complete other two short tasks and a questionnaire.

Preferences over Labels. Subjects are asked to indicate which type of mandatory front-package label they would prefer. They rank the four types of labels used in the treatments:²² no labels, two types, four types, or daily value percentage (with an image displaying ten types). After the first choice (most preferred option), the selected type disappears and the subject is asked to repeat the choice among the remaining options, until there is only one option left.²³

Preferences over Calories. The study contains two categories of products (cereals and salty snacks). For each category, I use an incentivized task to elicit subjects' preferences over calories amounts — whether they prefer low-calorie products, high-calorie, or anything in between.

First, I remind the subjects that similar products can have very different calories amount, and that most of the products are typically between 50 and 250 Calories per serving. Then, I ask them to rank the products based on calories (in ten bins of 20 Calories each), from the most to the least preferred.

Food Literacy Questionnaire. I use eight items from the food literacy questionnaire validated in Poelman et al. (2018). These questions focus on the subject's habitual behavior when it comes to consult and compare products' information (e.g., Do you check the nutritional labels of products for calories, fat, sugar, or salt content?) and reflect about own food choice (e.g., If you have something to eat, do you reflect on what you have eaten earlier that day?). I collect answers using a Likert 7 scale, with values ranging from 1 (no, never) to 7 (yes, always). To facilitate comparison of answers, in the analysis I refer to the FLI (food literacy index) that is a score that ranges between 0 (if all the answers are no) to 100

²²Note that at the moment of the question subjects had experienced only one of the four labels in the choice task. The results show a small and significant familiarity effect (the familiar label is chosen on average more than in the other treatments), but the overall ranking of the four options is the same across all the treatments.

²³The label ranking question is not incentivized and does not affect any subsequent question in the study.

Suppose the FDA (Food and Drug Administration) decides to introduce mandatory FRONT-PACKAGE labels for calories.

Which of these "label systems" would you prefer?

Select one among the food label systems below.

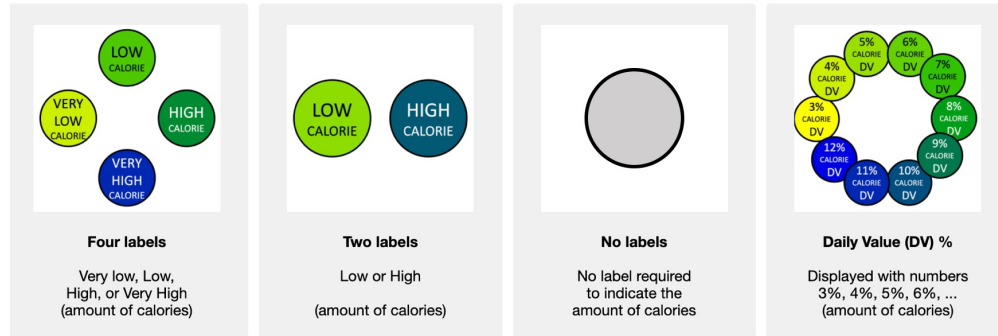


Figure 6: Preference over types of labels.

Snacks

Do you prefer high-calorie or low-calorie snacks?

Click the green buttons to rank the options based on your preferences
from the most (on the top) to the least preferred option.

▶ Instructions Details

% DV	Description	Ranking
<input checked="" type="checkbox"/> 3%	3% recommended Daily Value: 50-69 calories	7% recommended Daily Value
<input type="checkbox"/> 4%	4% recommended Daily Value: 70-89 calories	6% recommended Daily Value
<input type="checkbox"/> 5%	5% recommended Daily Value: 90-109 calories	5% recommended Daily Value
<input type="checkbox"/> 6%	6% recommended Daily Value: 110-129 calories	4% recommended Daily Value
<input type="checkbox"/> 7%	7% recommended Daily Value: 130-149 calories	
<input type="checkbox"/> 8%	8% recommended Daily Value: 150-169 calories	
<input type="checkbox"/> 9%	9% recommended Daily Value: 170-189 calories	
<input type="checkbox"/> 10%	10% recommended Daily Value: 190-209 calories	
<input type="checkbox"/> 11%	11% recommended Daily Value: 210-229 calories	
<input type="checkbox"/> 12%	12% recommended Daily Value: 230-249 calories	
		<input type="button" value="clear"/>

Figure 7: Preference over calorie amounts.

(if all the answers are yes).

Final Questionnaire. Finally I collect additional information about subjects' food habits (allergies, food intolerance, and dietary restrictions), body weight (whether they are underweight, healthy weight, overweight, or obese), and drivers of purchase (importance of price, brand, and healthiness of food). Sociodemographic data (e.g., gender, age, household income) are collected by the panel at the beginning of the study and used to screen the subjects according to the target quotas (details in the Recruitment section).

2.2 Preregistered Hypotheses

The experimental design, hypotheses, and sample size were preregistered on the American Economic Review Social Science Registry under the name “Coarse and Granular Nutritional Labels,” trial AEARCTR-0007856.

The primary hypothesis refer to the effect of labels on healthy choices. My hypothesis is that coarse labels generate healthier and better choices with respect to the control treatment, but most detailed labels are less effective than coarse ones. Healthier choices are defined in accordance with the number of calories per serving of the selected products. The variables of interest are the average amount of calories and the probability of choosing a low-calorie food option. Better choices are defined in accordance with self-reported preferences. Subjects are asked to rank products based only on the calorie amount. This allows to separate subjects that prefer high-calorie products, low-calorie, or intermediate values. Results are consistent with the hypothesis, as reported in Figure 8 and Table 2.

Secondary hypotheses refer to the effect of labels on the search process. My hypotheses are that more detailed labels lead to lower information collection (fewer consultations of the back-of-package information) and faster responses. The results show that the introduction of a label reduces search, but more detailed labels do not affect the search process (Figure 9.a). Responses are not faster under detailed labels; the opposite is true, with detailed labels being associated with longer response time (Figure 9.b).

The hypotheses are nested in the conjecture that providing more precise information might not help consumers in their decision process. I discussed in the introduction that the experiment aims to test precision overload; if decision makers find detailed information more confusing than coarse one, then precise labels can lead to worse choices.

2.3 Experimental Procedure

CloudResearch Prime Panels. Subjects are recruited online through the CloudResearch Prime Panels platform, an aggregation of double opt-in subject panels.²⁴ To increase the representativeness of the sample CloudResearch applied demographic quotas (purposive sampling) for gender, age, race, and ethnicity that are matched to the U.S. Census.²⁵

²⁴CloudResearch is the platform used to recruit participant from online panels. The subject panels are samples of volunteers who have agreed to complete online surveys and receive some form of compensation in exchange. Subjects are recruited by navigating to a website and signing up to take surveys. For this study, I required subjects to use a computer (not a mobile device) to make sure all the choice interface was visible in a single page.

²⁵Demographic quotas. Age: 24% 18-29, 24% 30-44, 24% 45-59, 28% 60-99. Gender: 50% Female, 50% Male. Hispanic: 16% Hispanic, Latino, or Spanish origin. Race: 4% Asian, 12% Black, 82% White. Region: 21% MidWest, 17% NorthEast, 38% South, 24% West. One limitation of the sampling design is that the

Sentry Screening. All subjects who enter a survey on Prime Panels are first run through CloudResearch’s Sentry data validation system. Sentry is a short (1 minute) pre-study screening that ensures subjects are attentive, engaged, and ready to participate. People who fail Sentry are routed away from the survey. Research shows that these sorts of short, pre-study validation measures are an effective way to increase data quality obtained from market research panels like Prime Panels (Chandler et al., 2019; Litman et al., 2020).

Compensation. The study took on average 23 minutes to complete.²⁶ The cost per complete for this study is \$5.25, that is divided between CloudResearch, the panel, and the subject. Subjects also have the chance to receive one of the products selected during the study.²⁷ Subjects on Prime Panels are compensated in a variety of ways. Some receive direct monetary compensation. Others participate in studies for a donation to charity or for points that they can redeem for prizes. Compensation is determined at the panel level, meaning that each individual panel decides how much and in what form subject compensation is administered (panels do not share compensation information with CloudResearch).

Subjects. 856 subjects completed the study. Table 1 contains the distribution of the subjects according to the main demographic characteristics and shows that the four treatments are demographically balanced according to gender, age, race, ethnicity, and region.

3 Results

The main results of the experiment highlight how label precision affects choices and the search process. Consistent with the primary hypothesis, I show that coarse labels are more effective than detailed labels in reducing the number of calories of the selected product. Secondary hypotheses on the search process are only partially validated: the introduction of

survey is only offered in English.

²⁶Average completion time for each section of the study: instructions and comprehension questions (4 minutes), choice task (11 minutes), preferences over labels (1 minute), preferences over calories (2 minutes), final questionnaire (5 minutes).

²⁷The cost per complete (\$5.25) corresponds to how much the experimenter paid to CloudResearch for a complete response to the study. At the end of the study, subjects can leave their email address in order to have the chance to receive a product from the study as bonus. The bonus can be from the choice task (each subject had one chance out of ten to receive one product they selected, plus the remaining part of the \$5 endowment after paying the price displayed by the product in the study) or the calorie preference task (one subject from the study received a product based on the preferences expressed). The products used in the study have prices typically between \$4 and \$6 and were presented in the study with displayed values between \$2.50 and \$3.50 to increase their attractiveness.

	US Population	All Treatments	A	B	C	D
N. subjects		856	216	213	214	213
Female	0.51	0.51	0.56	0.46	0.50	0.49
Age 18-29	0.18	0.21	0.21	0.22	0.21	0.22
Age 30-44	0.18	0.26	0.27	0.29	0.28	0.22
Age 45-59	0.19	0.23	0.22	0.22	0.21	0.27
Age 60+	0.21	0.30	0.31	0.27	0.31	0.30
Asian	0.06	0.06	0.05	0.05	0.07	0.06
Black	0.13	0.13	0.12	0.15	0.14	0.13
White	0.76	0.74	0.77	0.75	0.71	0.75
Hispanic	0.19	0.17	0.18	0.17	0.17	0.15
MidWest	0.21	0.20	0.19	0.20	0.18	0.23
NorthEast	0.17	0.17	0.14	0.19	0.20	0.15
South	0.38	0.38	0.41	0.39	0.39	0.33
West	0.24	0.24	0.26	0.21	0.22	0.28

Table 1: Distribution of respondents based on target demographic characteristics: all subjects (column *All Treatments*) and separately by treatment (columns A, B, C, and D).

front-of-package labels reduces the probability of consulting the back of the package, but the probability is not different across labels with different precision. We see the opposite effect for response time, with slower responses for detailed labels, suggesting that precise information might be more difficult to process. When asked explicitly, subjects report that they would prefer to receive coarse labels. Towards the end of the section I look at the results into more details, by analyzing effect in different choice sets (small and large differences between products’ calories) and across groups of subjects (e.g., participants with high and low food literacy).

3.1 Treatment Effect on Choice

How do calorie labels affect food choices? Figure 8 shows that all the label treatments (Binary, Coarse, and Detailed) generate a small but statistically significant shift towards low-calorie products. We observe a U-shaped effect of label precision on calories per serving (left) and probability of selecting a high-calorie product (right).²⁸ The average amount of calories per serving decreases from 143.1 (47% high-calorie products chosen in the control treatment in which the label is Absent) to 141.4 (45%, Binary label) and 139 (42%, Coarse).

²⁸High-calorie products are defined in relative terms with respect to the current choice set; these products have a number of calories per serving above the median value in the current choice set. Results are robust to different definitions of high-calorie products based on absolute amount of calories per serving (100, 150, or 200 calories per serving).

The effect of precision on choice is not monotonic, with higher calories in the Detailed treatment (141.6 calories per serving, 45% high-calorie products chosen.)²⁹

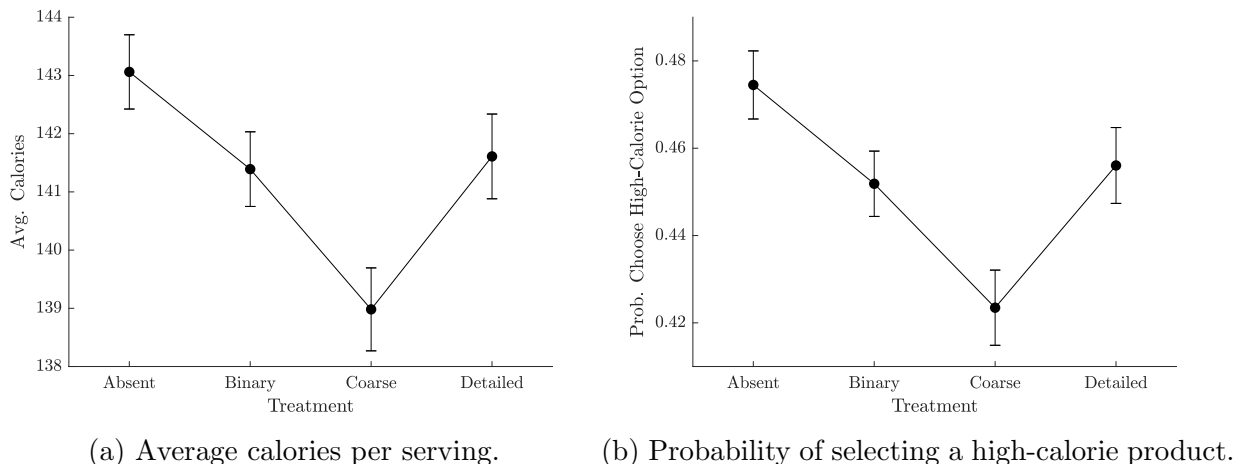


Figure 8: Treatment effect on food choices. SE clustered at the subject-category level. Left: average number of calories per serving. Right: probability of choosing high-calorie product (number of calories below the median in the current choice set).

A 3% decrease in calories is similar in magnitude to results observed in the empirical literature.³⁰ The reader should keep in mind that choices are made between similar products (cereals vs. cereals, snacks vs. snacks) and that the ranges of calories differ across trials.³¹ The feasible amount of calories given these constraints lies between 106 and 186 calories.³² The U-shaped effect contains three separate - and statistically significant - effects when we consider pairs of adjacent treatments and study how choices change when we introduce labels and refine their granularity. I will mostly focus on the effect of precise labels in the Coarse-Detailed comparison, as this is a striking results at odds with the predictions of the standard updating models. The comparison Binary-Coarse treatments highlight that increasing granularity can sometimes be beneficial, and in fact results are not statistically different between Binary and Detailed treatment.

²⁹All the differences between treatments A and B, B and C, and between C and D are statistically significant at the 5% level.

³⁰Abaluck (2011) estimates a 50-90 calories decrease in the US after the introduction of mandatory BOP labels, and Barahona et al. (2020) estimate a 6.5% decrease in calories in Chile after the introduction of mandatory warning labels. 50-90 calories correspond to 2.5-4.5% of a 2,000 calories diet. The author estimate via revealed preference that the welfare benefits of the labels lead to a \$25-40 annual gain.

³¹Low-variance choice sets have products with similar characteristics (e.g., 60 vs. 100 calories per serving), high-variance choice sets have larger differences (e.g., 100 vs. 200 calories.)

³²A respondent who chooses the lowest-calorie product in every trial would show an average amount of calories per serving equal to 106 calories. A random respondent who chooses randomly in every trial would reach 146.

OLS Regression. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	143.06*** (0.64)	146.12*** (0.53)	144.93*** (0.67)	145.40*** (1.09)	141.18*** (1.83)
Treatment B	-1.67* (0.90)	-2.56*** (0.79)	-2.56*** (0.79)	-2.85*** (0.82)	-3.22** (1.30)
Treatment C	-4.08*** (0.96)	-5.17*** (0.83)	-5.16*** (0.82)	-5.04*** (0.85)	-5.82*** (1.42)
Treatment D	-1.45 (0.97)	-2.42*** (0.88)	-2.41*** (0.87)	-2.14** (0.88)	-2.89** (1.39)
Flip Package		-14.22*** (0.95)	-14.10*** (0.92)	-13.55*** (0.96)	-13.01*** (1.27)
Treatment $B \rightarrow C$	-2.41** (0.96)	-2.61*** (0.87)	-2.60*** (0.85)	-2.18** (0.88)	-2.61* (1.35)
Treatment $C \rightarrow D$	2.63*** (1.02)	2.75*** (0.94)	2.75*** (0.93)	2.89*** (0.94)	2.94** (1.45)
Treatment $B \rightarrow D$	0.22 (0.97)	0.14 (0.91)	0.15 (0.90)	0.71 (0.91)	0.33 (1.34)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	68451	68451	68451	64851	22792
Adjusted R ²	.001	.014	.428	.427	.505

Table 2: Treatment effect on calories per serving (OLS regression). SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Tables 2 and 3 show that the results are robust to several specifications. Model 1 shows the differences across treatments. Model 2 introduces endogenous search by controlling for collection of information on the back of the package (discussed in more details in the next section). Model 3 adds choice set controls: the minimum and maximum number of calories among the products available, product category (cereals or snacks) and choice set size. Model 4 introduces demographic controls, with categorical variables for gender, age bracket, race, ethnicity, income, and education. Model 5 controls for preferences by restricting the analysis to monotonic preferences for low-calorie products. Observations that respect the monotonic preferences constraints are defined using subject-category level data on preferences over calories. Model 5 restriction provides a stronger test that address concerns about heterogeneous preferences. A possible explanation for the U-shaped effect is that subjects do not always prefer fewer calories; instead, they might find a product more appealing exactly because it has more calories than expected. I take this concern seriously by considering only observations in which, according to their own responses, subjects would prefer to move towards products with lower calories.³³ In lay terms, a monotonic observation is a trial in

³³Here is a formal definition of how I restrict the analysis to monotonic observations. Each label is

OLS Regression. DV: Probability of choosing a high-calorie product.					
	(1)	(2)	(3)	(4)	(5)
Constant	0.474*** (0.008)	0.509*** (0.006)	0.508*** (0.008)	0.513*** (0.013)	0.465*** (0.023)
Treatment B	-0.023** (0.011)	-0.033*** (0.010)	-0.033*** (0.010)	-0.035*** (0.010)	-0.035** (0.017)
Treatment C	-0.051*** (0.012)	-0.063*** (0.010)	-0.064*** (0.010)	-0.062*** (0.010)	-0.071*** (0.019)
Treatment D	-0.018 (0.012)	-0.029*** (0.011)	-0.030*** (0.011)	-0.024** (0.011)	-0.038** (0.018)
Flip Package		-0.160*** (0.012)	-0.165*** (0.012)	-0.157*** (0.012)	-0.162*** (0.017)
Treatment $B \rightarrow C$	-0.028** (0.011)	-0.031*** (0.010)	-0.031*** (0.010)	-0.027** (0.011)	-0.036** (0.017)
Treatment $C \rightarrow D$	0.033*** (0.012)	0.034*** (0.011)	0.034*** (0.011)	0.038*** (0.011)	0.032* (0.018)
Treatment $B \rightarrow D$	0.004 (0.011)	0.003 (0.011)	0.003 (0.011)	0.011 (0.011)	-0.003 (0.017)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	68454	68451	68451	64851	22792
Adjusted R ²	.001	.015	.025	.024	.033

Table 3: Treatment effect on probability of selecting a high-calorie product (OLS regression). SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. High-calorie products have a number of calories per serving above median in the current choice set.

which the subject prefers low-calories over high-calories over the whole range of values that the products can take.

3.2 Treatment Effect on Search

How does label precision affect information acquisition? Figure 9 shows the difference in search and response time across treatments. During the study, subjects can consult the nutrient facts in the BOP by flipping the package. The left panel shows that subjects check the nutrient facts sporadically, and even less when the FOP contains a label. In the Absent

associated to a range of feasible calorie values. For each product in the choice set, I consider what would be the intervals if participants were in the coarse treatment and I consider the preferences over calories in that interval. An observation fulfill the monotonicity requirement if, for all the intervals, the subject declares monotonic preferences with low-calorie products preferred to high-calorie products. This means that I always include observations of subjects that declare monotonic preferences low-to-high calories for the whole range of possible values, I always exclude the opposite pattern (high-to-low calories) and, for single-peaked preferences, I consider only observations that are on the right-hand side of the peak.

treatment they flip the package 21% of the times, whereas in the three label treatments they flip it only 14% of the times. The difference is large and statistically significant when I compare treatment A against each of the other three (all $p < 0.01$) but not when we compare the label treatments with each other (all $p > 0.82$). Table 5 show that this pattern is robust to different specifications that introduce controls for the choice set and demographic characteristics of the respondents.

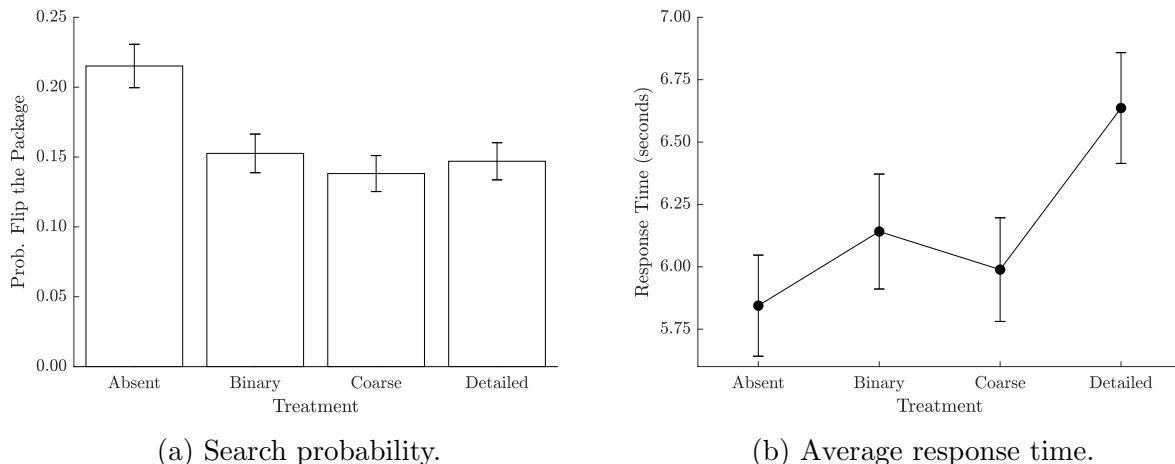


Figure 9: Process data across treatments. SE clustered at the subject-category level. Left: the probability of consulting back-of-package information decreases when any front-of-package label is introduced. Right: the average response time is significantly longer only in the Detailed treatment.

The right panel shows that the response time for each choice is on average 6 seconds, and is 0.6 seconds higher in the Detailed treatment. The difference is statistically significant when I compare treatment D against each of the other three (all $p < 0.01$) but not when we compare the first three treatments with each other (all $p > 0.57$). The longer response time associated with the detailed treatments suggests that the most precise labels could be confusing or difficult to understand for the subjects, and provides a suggestive evidence in favor of precision overload.

How does information acquisition interact with choices? Figure 10 separates choices based on the search behavior. The figures show the effect of two types of information: the exogenous FOP label manipulation and the endogenous information acquisition. Flipping the package is associated with low-calorie choices; the calories per serving in the Absent treatment are 130 and 147 ($p < 0.001$), with and without search respectively, and high-calorie choices probabilities are 33% and 51% ($p < 0.001$).³⁴ Differences remain large and

³⁴A trial is classified as “with search” if the subject checked the BOP for the selected product. The results are robust to different specifications (search for at least one product, search for all the products).


significant across treatments. This is not surprising if we consider that the few subjects who consult the nutrient facts might be those who care the most about the calorie content.

This conjecture is confirmed by comparing the search behavior across groups; Table 7 shows that search is more frequent among high-income, high-education, and high-food literacy subjects, and in general among subjects who display healthier choices in the study. Another striking feature is that search behavior appears highly concentrated, with three patterns. First, search is concentrated across subjects; only few participants flip the package, but they do it consistently. 14% of the subjects flip the package in 40 trials or more (out of 80), and 66% of the subjects flip the package in less than 5 trial. Second, search is concentrated across products; once one product in the choice set is flipped, typically all the other products available are also flipped. The conditional probabilities of flipping all the remaining packages are 97% and 75% in trials with two and four products respectively. Third, search is concentrated across time; when multiple products are flipped, they are flipped early and almost simultaneously. Trials in which all products are flipped have an average response time of 14.2 seconds; on average the first product is opened after 2.1 seconds and the last product after 4.4 seconds.


The U-shaped treatment effect is robust to the search behavior. When we consider the effect of different labels conditional on information acquisition we see the familiar pattern, with fewer calories when we move towards Coarse labels and higher calories when we introduce Detailed labels. Suppose that having access to BOP information was sufficient to make an informed choice; in that case we would expect choices not to differ across treatments, conditional on observing the endogenous information acquisition. Instead, the pattern persists even in the observations in which search occurs. A possible interpretation is that subjects do not fully absorb the information on the BOP and make use of the front label. This would be consistent with the behavioral models I discuss in Section 6, in which consumers have noisy mental representations of the information acquired.

3.3 Preferences over Types of Labels

What kind of labels do subjects prefer? After the choice task, I present a scenario in which FDA is evaluating the implementation of mandatory FOP calorie labels. I ask the subjects to rank the four possible types of labels (absent, binary, coarse, and detailed) based on their preferences as consumers. Figure 11 shows that choice probability for each label. Standard information models would predict that the most detailed label should always be preferred. Instead subject responses show that not only binary and coarse labels are selected by a large fraction of the respondents, but coarse labels are the overall preferred option. Coarse labels



(a) Average calories per serving.



(b) Probability of selecting a high-calorie product.

Figure 10: Treatment effect on food choices, conditional on search behavior. SE clustered at the subject-category level. Left: Average number of calories per serving. Right: probability of choosing high-calorie products. Observations are divided based on whether the subject had consulted the back-of-package information for the product they selected.

have the highest probability of being ranked above other ones (71%, $p < 0.001$, followed by the binary labels with 60%) and the highest probability of being ranked first (34%, $p = 0.20$, followed by the detailed labels with 31%). Table 4 considers the whole ranking and contains the choice probability for each pairwise comparison. Coarse labels are the Condorcet winner, receiving more than 60% of the preferences in each pairwise comparison. The Condorcet winner of a voting problem is the option that would win the two-options election against each of the other candidates in a plurality vote. Coarse and Detailed labels are selected as first option by a similar fraction of subjects, but coarse labels are much more likely to be ranked second (47% of the times).

Are preferences over labels affected by the treatment? I elicit the preferences after the choice task, and subjects in different treatments are exposed to one label more than the other three. Figure 12 shows that familiarity³⁵ improves the ranking of all types of labels and in particular it boosts the likelihood of the same label being ranked first, but the coarse label remains the Condorcet winner in all the treatments.³⁶ The label observed in the choice task is ranked first in 33% of the cases, while we would observe just one quarter if there was no such effect ($p < 0.001$).

Table 5 shows that the result is robust to specifications that take into account treatments

³⁵The effect of familiarity observed in the data is consistent with what is discussed in the literature, see e.g. Hansen and Wänke (2009).

³⁶The coarse label is the most selected in all the treatments and the Condorcet winner, but it is not always the most frequently ranked first. In treatments B, C, and D, the familiar label is ranked first with the highest probability.

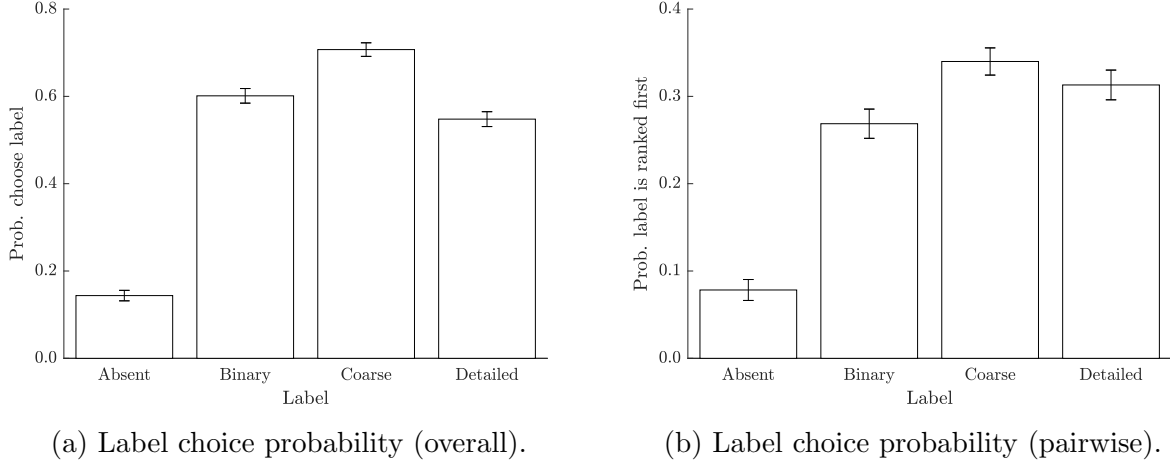


Figure 11: Preferences over types of labels. Left: overall probability of ranking a label above other ones (including all the pairwise comparisons). Right: probability of ranking a label as first among the four options.

Probability that label X...	...is preferred to label Y.			
	Absent	Binary	Coarse	Detailed
Absent		0.12	0.11	0.20
Binary	0.88		0.37	0.55
Coarse	0.89	0.63		0.61
Detailed	0.80	0.45	0.39	

Table 4: Choice probability for each pair of labels determined from the ranking distribution.

and demographic controls. Later in the section, Table 7 highlights the heterogeneous label preferences across groups; more detailed labels are preferred by subjects with high food literacy index, high education, high income, and preference for low-calorie products.

Subjects prefer coarse labels over detailed ones and also over not receiving any additional information. The result is not driven by two separate groups with opposite preferences (all or nothing); instead there is a significant number of people who place it as first option, plus other subjects (on both extremes) that are satisfied by coarse labels. This result is related but deeply different from what we know about information avoidance (see e.g. Golman, Hagmann and Loewenstein 2017 and Reisch, Sunstein and Kaiser 2021). Value of information is commonly defined through a combination of instrumental and hedonic effects. The first one is directly linked to the ability of making better choices thanks to the additional knowledge, while the second refers to the utility, or disutility, generated by receiving the information. For example, receiving good news can lead to immediate utility (or anticipation of future utility as in Loewenstein 1987) even when this has no direct effect on future outcomes. The evidence




Figure 12: Preferences over types of labels across treatments. The colors indicate the treatments that the subjects completed, the x-axis indicate the type of label. Left: overall probability of ranking a label above other ones (including all the pairwise comparisons). Right: probability of ranking a label as first among the four options.

	Prob. Search		Label Preference	
	(1)	(2)	(3)	(4)
Constant	0.22*** (0.02)	0.16*** (0.03)	0.60*** (0.02)	0.55*** (0.03)
Treatment B	-0.07** (0.03)	-0.07** (0.03)	0.02 (0.03)	0.01 (0.03)
Treatment C	-0.08*** (0.03)	-0.08*** (0.03)	0.04 (0.03)	0.04 (0.03)
Treatment D	-0.07** (0.03)	-0.07** (0.03)	0.06* (0.03)	0.06* (0.03)
Demographic Controls		✓		✓
Observations	68451	68451	856	856
Adjusted R ²	.011	.046	.006	.035

Table 5: Treatment effect on search and label preference. SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Search probability: probability of consulting the back-of-package information. Label preference: preferred calorie label (explicit elicitation) normalized between 0 (prefer Absent label) and 1 (prefer Detailed label).

of preference for coarse labels is not fully explained by this framework, that would allow heterogeneity across subjects (some of which would like to know more) but not necessarily that they want to know something in between. Preference for simple sources of information has also been shown in stylized laboratory experiment. Ambuehl and Li (2018) and Novák, Matveenko and Ravaioli (2021) report that participants are willing to pay more and choose

more often “simple” signals that can lead to degenerate posterior beliefs. This definition of simplicity would still favor granular signals, but the experiments above did not manipulate the number of possible signal realizations. The result of the experiment presented in this paper fits well the idea of precision overload, with decision makers disliking information that is too precise and potentially confusing. The model discussed in Section 5 will provide a formal explanation of how both choices and preferences over types of information in the experiment are consistent with this effect.

3.4 Response to Label Information

Why are detailed labels less effective than coarse ones? In principle, there are two extreme ways in which the detailed labels could reduce the choice of low-calorie items: either through a fixed or a gradual response to the calorie content. In the first case, choices of the product with relatively fewer calories would drop, regardless of the differences in the calorie content across products. In the second case, choices would change more for large calorie differences, but not for small differences.

The pair of graphs in Figure 13 show a significant difference in the slope of the response to the calorie difference between two products available, without significant changes to the intercept. The two graphs show that the direction of the effect is the same when we consider the differences between true (left) or the inferred product calories (right).³⁷ Inferred calories are calculated based on the information provided by the FOP label only. For example, in a trial with Binary labels the implied difference can be either 0 calories (when both products display the same label) or 100 calories (when one product displays Low and the other displays High).³⁸ The figure on the right shows the change in the choice probability in favor of the low-cal product with respect to the Absent treatment. In general, we observe that larger implied difference lead to more frequent choice of healthy products. On one extreme, when two products display the same label (implied difference equal to zero), choices appear less healthy, and this can be explained by the search reduction discussed in Figure 9. We are interested in comparing the effect of Coarse and Detailed labels. The figure offers a visual decomposition of the treatment effect between small and large differences; response is not significantly different for values below 50 calories, but the response is weaker for

³⁷The difference between implied calorie differences Δc is defined as $\mathbb{E}[c_j] - \mathbb{E}[c_k]$ where j is the target product and k is the product with the highest amount of calories in the current choice set. Expectations are taken with respect to the information provided by the label.

³⁸For choice sets with four items, differences are calculated with respect to the product with the highest number of calories. Items do not always have an exact difference of 100 calories: it could be less (e.g., 120 vs. 160 calories) or more (e.g., 80 vs. 220 calories), but in both these scenarios the products would display two different labels.

larger differences. We can interpret this result in terms of lower sensitivity to the label information. A Wald test confirms that the difference across treatments is mostly driven by the change in the slope ($\beta_C = -0.00112$, $\beta_D = -0.00076$, $p = 0.022$) and not in the intercept ($\alpha_C = -0.0485$, $\alpha_D = -0.0446$, Wald test $p = 0.28$).

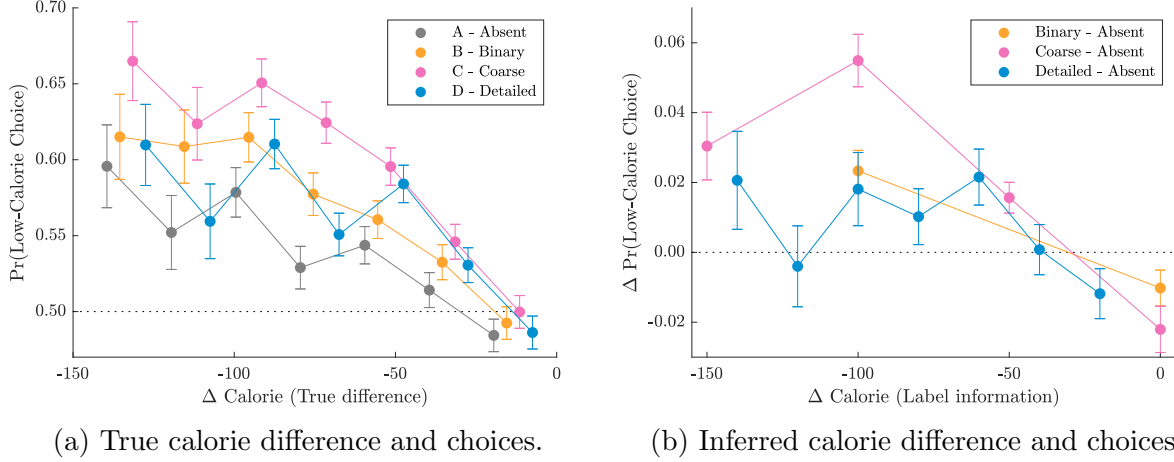


Figure 13: Response to calorie information in different treatments. SE clustered at the subject-category level. Choice probability conditional on the difference between the calorie content of two products. Left: true calorie content. Right: change in choice probability, conditional on calorie content inferred from the label.

A complementary approach is to decompose the effect across products with different calorie content. Here I consider each product in isolation, instead of by comparing the calories with other products in the choice set. Figure 14.a shows that low-calorie products below 150 calories are chosen significantly more often in the Coarse treatment compared with either the Binary and Detailed treatments, with the largest differences in the middle range of values, 100 to 200 calories, while extreme calories do not display significant differences. This suggests that even in the control treatment subjects have a good sense of which products lie at the extremes of the interval, and label is able to provide information mostly for intermediate values. The Coarse treatment choices systematically lie between the Absent and the Detailed treatment. Figure 14.b highlights the gradual information disclosure through nested labels. Moving from no label to Binary labels, and from Binary to Coarse labels increases the likelihood to select low-calorie products, following a decreasing pattern for each segment. Instead, the introduction of Detailed labels (blue) shows two patterns. First, the downward trend emerges for three out of four segments. This means that in most of the cases the adoption of detailed labels favors the relatively low-calorie products within each coarse category. Second, the choice probabilities show an overall trend reverse, with products below 150 calories being selected less often, reducing the general effectiveness of labels. This pattern is associated

with the weaker impact on choice of Detailed labels.

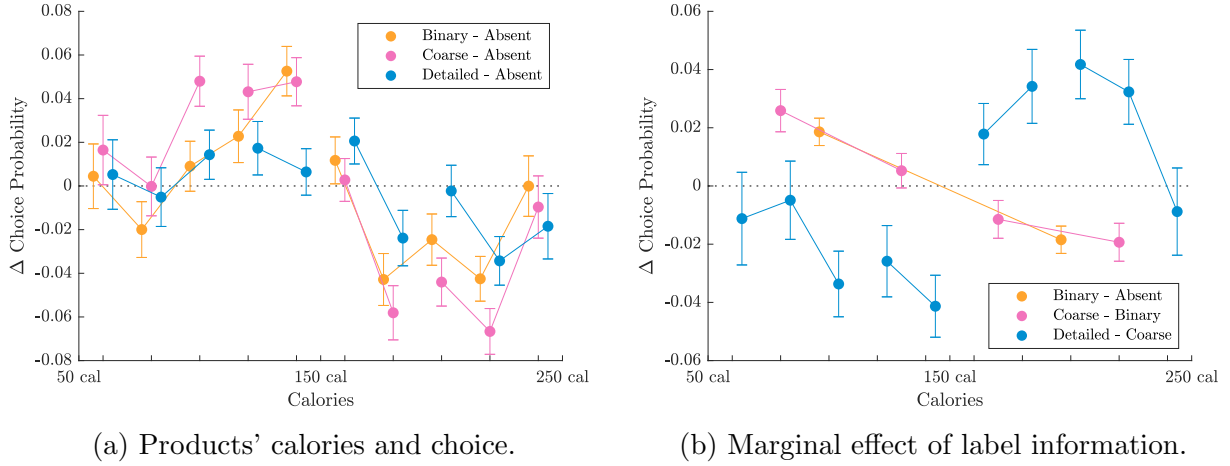


Figure 14: Response to label information in different treatments. SE clustered at the subject-category level. Left: change in choice probability, compared with the Absent treatment, conditional on true calorie content. Right: marginal change in choice probability, compared with the adjacent treatment, conditional on the inferred calorie content.

3.5 Heterogeneous Treatment Effect

Are the results different across sociodemographic groups of respondents? I collect sociodemographic data, food habits, and calorie preferences for each subject and I compare the direction and size of the effect across groups. Table 6 contains the main results. For each group, I show three summary statistics: the baseline behavior (treatment Absent), the effect of coarse labels (difference between Coarse and Absent treatments), and the *marginal* effect of detailed labels (difference between Detailed and Coarse treatments). Results are expressed in number of calories (left) and probability of choosing a high-calorie product (right).

Two important patterns emerge from the table, and can be observed in Figure 15. First, the direction of the effect is robust across groups. We systematically observe a positive effect of coarse labels (lower calories) followed by a negative marginal effect of detailed labels (higher calories). Second, the magnitude and statistical significance of the effect are not equal across groups.

The comparison across groups highlight that food behavior and socioeconomic status mediate the label effect. For food behavior, we observe similar patterns for calorie preferences and food literacy index (FLI). Subjects who prefer low calories or have high FLI have higher baseline and larger response to the introduction of coarse labels. For socioeconomic status, coarse labels are more effective for subjects with low education and low income. These groups

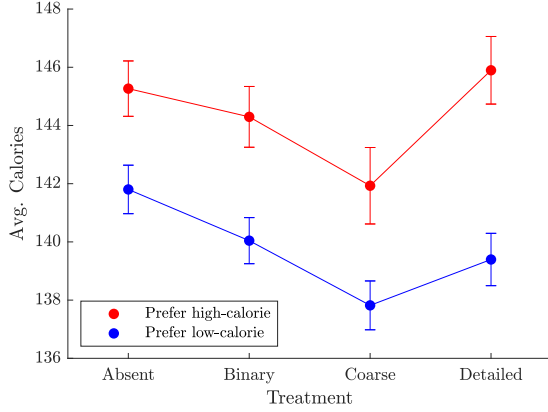
Group	Avg. calories			Pr. choose a high-calorie product		
	A	$A \rightarrow C$	$C \rightarrow D$	A	$A \rightarrow C$	$C \rightarrow D$
All	143.06 (0.64)	-4.08*** (0.96)	2.63*** (1.02)	0.474 (0.008)	-0.051*** (0.012)	0.033*** (0.012)
Prefer low calories	141.80 (0.83)	-3.98*** (1.18)	1.58 (1.23)	0.461 (0.010)	-0.051*** (0.014)	-0.020 (0.015)
Prefer high calories	145.27 (0.95)	-3.34** (1.62)	3.97** (1.75)	0.489 (0.011)	-0.040** (0.020)	0.049** (0.022)
FLI high	141.41 (0.89)	-4.45*** (1.33)	3.43** (1.44)	0.460 (0.011)	-0.058*** (0.016)	0.045*** (0.017)
FLI low	144.71 (0.90)	-3.38*** (1.35)	1.60 (1.41)	0.489 (0.011)	-0.040** (0.017)	-0.017 (0.017)
Education high	143.66 (0.82)	-2.77** (1.26)	0.34 (1.45)	0.486 (0.010)	-0.037** (0.015)	0.005 (0.017)
Education low	142.46 (0.97)	-5.19*** (1.42)	4.60*** (1.42)	0.463 (0.012)	-0.062*** (0.017)	0.057*** (0.017)
Income high	142.84 (0.79)	-3.11** (1.30)	1.54 (1.49)	0.474 (0.010)	-0.041*** (0.016)	0.019 (0.018)
Income low	143.37 (1.05)	-5.04*** (1.44)	3.57** (1.40)	0.476 (0.013)	-0.061*** (0.018)	0.044*** (0.017)

Table 6: Heterogeneous treatment effect. SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. For each group (row) the table reports the treatment effect on calories (left) and probability of choosing a high-calorie product (right). Summary statistics indicate baseline values (Absent treatment), the effect of coarse labels (Absent to Coarse treatment), and the marginal effect of detailed labels (Coarse to Detailed treatment). High-calorie products: calories per serving above median in the current choice set.

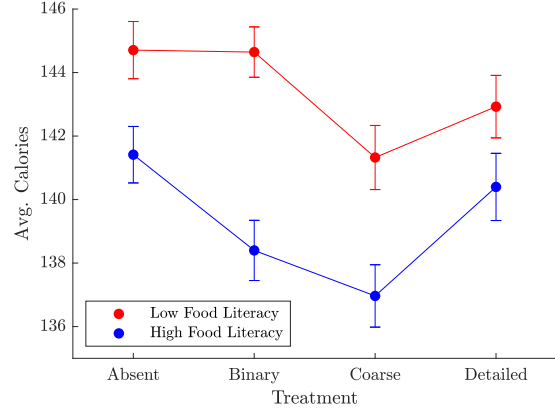
display a larger effects both in the comparison Absent-Coarse (positive) and Coarse-Detailed (negative effect).

These differences are consistent with the results highlighted in experiments that involve nudges; Mrkva et al. (2021) show that consumers with low SES, domain knowledge, and numerical ability are impacted more by nudges. This lead to a reduction of disparities when “good nudges” (designed to increase the choice of superior options) are introduced, but the opposite is also true when under “bad nudges.” In the context of my study, this suggests that detailed (complex) labels can be interpreted as bad nudges³⁹ and lead to higher disparity in choices.


³⁹Mrkva et al. (2021) study the effect of nudges that introduce a default option in a series of experiments in which a correct answer exists. They classify good and bad nudges based on the ex post outcome, i.e. if the default option is indeed the correct answer of the problem. In my experiment there is no explicit correct or wrong answer, and the intervention affects the information



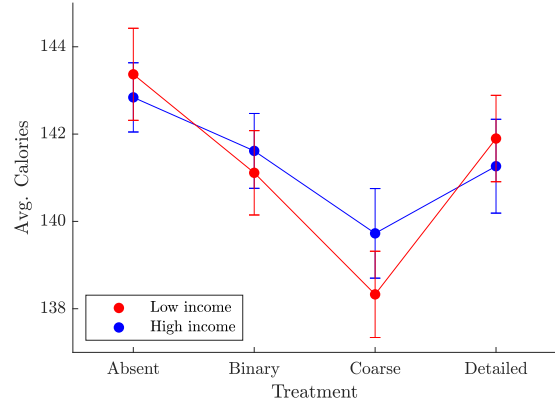
(a) Label choice probability (overall).



(b) Label choice probability (pairwise).



(c) Label choice probability (overall).



(d) Label choice probability (pairwise).

Figure 15: Heterogeneous treatment effect on calories per serving. Results are presented separately based on (a) calorie preferences, (b) food literacy, (c) education, and (d) income. Calorie preferences (a) are classified based on the preferred calorie amount (above or below 150 calories). The other three figures (b, c, and d) classify respondents based on the median split. Standard Errors clustered at the subject-category level.

4 Bayesian Updating Model

The descriptive evidence from the experiment suggests that subjects might underreact to detailed labels' information, but this is not sufficient to rule out that their behavior is still consistent with the Bayesian updating model. In this section I test whether the experimental dataset provides stronger evidence against the Bayesian updating model and in favor of the claim that subjects underreact to detailed labels. I do this by characterizing the choice task in the experiment as a choice problem under uncertainty about the products' characteristics. By taking this approach I am able to derive testable predictions about how consumers' beliefs are related across treatments. I use the experimental dataset to estimate the label fixed effects

Group	Pr. Search	R. Time	Calorie Pref.	Label Pref.	Food Literacy
All	0.163 (0.010)	4.7 (0.2)	84.5 (1.9)	0.629 (0.011)	63.4 (0.7)
Prefer low calories	0.177*** (0.012)	4.8* (0.2)	58.3*** (1.3)	0.632 (0.012)	65.5*** (0.8)
Prefer high calories	0.124*** (0.017)	4.2* (0.3)	160.2*** (2.1)	0.621 (0.022)	57.4*** (1.6)
FLI high	0.214*** (0.016)	5.3*** (0.2)	79.4*** (2.5)	0.679*** (0.014)	79.9*** (0.5)
FLI low	0.109*** (0.012)	4.0*** (0.2)	90.1*** (2.8)	0.576*** (0.016)	45.6*** (0.8)
Education high	0.170 (0.015)	4.8 (0.2)	81.2* (2.6)	0.649* (0.015)	67.8*** (0.9)
Education low	0.157 (0.014)	4.6 (0.2)	87.4* (2.7)	0.612* (0.015)	59.6*** (1.1)
Income high	0.163 (0.014)	4.5 (0.2)	82.9 (2.5)	0.642 (0.015)	66.8*** (0.9)
Income low	0.163 (0.015)	4.9 (0.2)	86.2 (2.8)	0.616 (0.015)	59.9*** (1.1)

Table 7: Heterogeneous responses across groups. For each group (row) the table reports the average and SE for search (probability of flipping the products’ package), response time (expressed in seconds), preference over calories (peak of the preferences, with 60 for the lowest and 240 for the highest amount of calories), preference over types of labels (ranked based on precision and normalized in the range 0-1, with 0 for the Absent label and 1 for the Detailed labels), and Food Literacy Index (between 0, low, and 100, high). T-test for the differences between the two groups: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

and test the predictions. The estimates will confirm what the descriptive evidence suggest; labels have larger weights on decisions in the coarse treatments, and subjects underreact to detailed label information.

4.1 Setup: Discrete Choice under Uncertainty

Preferences. I follow closely the setup used by Barahona et al. (2020) to characterize consumers who learn from label. The model consists of a distribution of consumers indexed by $i \in \mathcal{I}$ and a set of J products indexed by $j \in \mathcal{J}$. In each trial $t \in \mathcal{T}$ the consumer faces a choice set \mathcal{J}_t and selects one product among those available. Trials can differ in the set of products offered \mathcal{J}_t and in the prices of the products p_{jt} . Later I will introduce product labels, that can also differ across trials.

I assume that the utility derived by individual i when purchasing product j can be

separate into three components:

$$u_{ijt} = \underbrace{h_i(c_j)}_{\text{health consequences}} + \underbrace{\delta_{ij}}_{\text{experience/taste}} + \underbrace{\alpha_i p_{jt}}_{\text{price}} \quad (1)$$

Products have c_j calories and the consumer has a mapping $h_i(c_j)$ from calories to utility that we can interpret as the part of utility related to health consequences. The second part is δ_{ij} , that represents the direct utility from consuming the product. The third part is $\alpha_i p_{jt}$ that captures the disutility from paying price p_{jt} , with α indicating the price elasticity.

Labels. I assume consumers do not observe the true calorie content c_j , but have beliefs about it based on initial knowledge and information conveyed through labels. A label message $l_{jt} \in \mathcal{L}_t$ provides information about product j in trial t according to a labeling regime (\mathcal{L}_t, τ_t) . In the experiment, the labels l_{jt} remain the same across all the trials within a treatment. The labeling regime (\mathcal{L}_t, τ_t) is a pair of sets: a sorted set of label messages \mathcal{L} (e.g., Low-Calorie and High-Calorie) and a sorted set of thresholds values τ that generate the partition of the calorie space into coarse labels (e.g., 0, 200, and 1000 calories).⁴⁰

Product j is assigned a label l_{jt} according to the labeling regime and based on its calorie content c_j .

$$l_{jt} = l_m \Leftrightarrow \tau_{m-1} \leq c_j < \tau_m \quad (2)$$

In lay terms, the product receives the m -th label message if and only if its calorie content lies in the m -th cell of the calories partition.

The consumer infers the calorie amount, and the associated health consequences, from the label and chooses the product that maximizes the expected utility

$$\mathbb{E}_c[u_{ijt}] = \mathbb{E}_c[h_i(c_j)|l_{jt}] + \delta_{ij} + \alpha_i p_{jt} \quad (3)$$

We are interested in the role of $\pi_{h|l_{jt}} := \mathbb{E}_c[h_i(c_j)|l_{jt}]$,⁴¹ that represents the expected utility that depend on the label message (label utility thereafter). I assume that before observing any label the consumer is uninformed and holds prior beliefs follow the true distribution of calories across products $c \sim F_C$.

I assume agents observe accurately the label and interpret its meaning correctly. When

⁴⁰When it is not confusing I will use the term label to refer to the label message l , and regime to refer to the labeling regime (\mathcal{L}, τ) . Precise and coarse labels indicate nested labeling regimes, with the precise regime being nested (more detailed) in the coarse regime.

⁴¹To indicate beliefs and expected values I introduce a different notation from Barahona et al. (2020) and I use $F_{H|l}$ to indicate the posterior distribution given the label information, and $\pi_{h|l}$ to indicate the expected utility.

a label is observed the consumer updates own calorie and health beliefs to the distribution $F_{C|l}$ and $F_{H|l}$ respectively, defined as

$$f_{C|l}(c) := f_C(c) \frac{Pr(l|c)}{Pr(l)} \quad \text{and} \quad f_{H|l}(h) := \int_c f_{C|l}(c) \cdot \mathbb{1}(h(c) = h) dc$$

and depend on the a priori probability of encountering the label $Pr(l)$ and the conditional probability $Pr(l|c)$ which is 0 for all the values of c that are associated with different messages.

In this setup I assume that consumers observe the front-of-package label but do not consult (or remember) the exact information on the nutrition facts panel. If the exact information was observed, the consumer would not learn anything new from the label. A more general model can separate the consumers into two categories: attentive consumers, who know already the nutrient characteristics, and inattentive consumers, who learn from the labels. This would be consistent with what I observe in the experiment. Few subjects systematically engage in the search process to consult the nutrition facts, but the vast majority is inattentive and learns from the front-of-package only. The Discussion section contains an extension in which consumers have imperfect memory and learn both from the front and back of the package.

4.2 Testable Predictions

Nested Labeling Regimes. In the experimental design I introduce treatments that differ from each other in the granularity of the labeling regime. Each treatment is a refinement of the previous one in the sense that each message is divided into multiple messages that convey the same information and more. I characterize formally this setup as a partition of the calorie space and I define when a label is coarser than another.

Consider two labeling regimes with thresholds τ and τ' . The labeling regime with τ is *nested* in the labeling regime with τ' if $\tau' \subset \tau$ - i.e. the set of values that determine the coarse labels is a subset of those that determine the nested (precise) labels.⁴²

For each coarse label l'_c we can define the set $\mathcal{L}_{d|l'_c}$ of detailed labels that are feasible under l'_c

$$\mathcal{L}_{d|l'_c} := \{l_d \in \mathcal{L}_d : Pr(l_d|l'_c) > 0\}$$

where $Pr(l_d|l'_c)$ is the probability that a product with label l'_c in the coarse regime can receive

⁴²We can define a partial ordering of the labels \succ_p based on the precision of their cutoffs, and following the notation from above $\tau \succ_p \tau'$ structure generates a strict partition of the latter. This relation generates a surjection $\mathcal{L} \rightarrow \mathcal{L}'$: for each precise label $l \in \mathcal{L}$ there is one and only one corresponding coarse label $l' \in \mathcal{L}'$, but there can be multiple precise labels $l_1, l_2 \in \mathcal{L}$ associated to the same coarse label $l' \in \mathcal{L}'$. For example $\mathcal{L}' = \{A, B\}$ (coarse) and $\mathcal{L} = \{A+, A-, B+, B-\}$ (precise).

the detailed label l_d , calculated according to the calorie distribution F_C .

Relation Between Nested Labeling Regimes. We have all the ingredients to establish the relation between the expected label utility $\pi_{h|l}$ in different regimes. Consider a product that displays a specific coarse label $l'_c \in \mathcal{L}_c$ and with expected label utility

$$\pi_{h|l'_c} := \mathbb{E}_c[h(c)|l'_c] = \int_c h(c) d\hat{F}_c(c|l'_c) \quad (4)$$

where $\hat{F}_c(c|l)$ is the subjective calorie distribution, conditional on observing the label l , and may differ from the correct distribution $F_c(c|l)$. Similarly, we can use $\hat{Pr}(l_d|l'_c)$ to indicate the subjective conditional probability.⁴³ We want to express this value in terms of the detailed label utilities, considering the counterfactual scenario in which the labels are more granular. It is useful to use the law of iterated expectations to rewrite

$$\pi_{h|l'_c} = \mathbb{E}[h|l'_c] = \mathbb{E}[\mathbb{E}[h|l'_c, l_d]] = \sum_{l_d \in \mathcal{L}_d} \hat{Pr}(l_d|l'_c) \cdot \mathbb{E}[h|l'_c, l_d]$$

We are considering detailed labels that are partitions of the coarse ones, and this allows to simplify the expression in two ways. First, following the definition of feasible set we have $Pr(l_d|l'_c) = 0 \quad \forall l_d \notin \mathcal{L}_{d|l'_c}$. Second, the detailed labels are sufficient statistics for the posterior beliefs that leads to $\mathbb{E}[h|l'_c, l_d] = \mathbb{E}[h|l_d]$. We can now rewrite the value decomposition:

$$\pi_{h|l'_c} = \mathbb{E}[h|l'_c] = \mathbb{E}[\mathbb{E}[h|l'_c, l_d]] = \sum_{l_d \in \mathcal{L}_{d|l'_c}} \hat{Pr}(l_d|l'_c) \cdot \pi_{h|l_d} \quad (5)$$

The value of the coarse label $\pi_{h|l'_c}$ is the convex combination of the *feasible* detailed labels $\pi_{h|l_d}$, weighted by the probability of encountering each detailed label.⁴⁴ This setup allows to derive two types of predictions about the relation between the label weights in different regimes.

The first prediction requires that the agent operates Bayesian updating but does not require further assumptions on the belief distributions. We can derive bounds for the coarse label weights from the detailed ones. The utility of the coarse label lies between the lowest

⁴³The subjective conditional probability of the partition from a coarse to a detailed regime is $\hat{Pr}(l_d|l'_c) = Pr(l_d(c_j) = l_d|l'_c(c_j) = l'_c)$. The value is zero for each detailed label outside the partition cell associate with the coarse label l'_c .

⁴⁴The intuition can be presented with a simple numerical example. Consider a detailed regime with four labels $\{l_{d1}, l_{d2}, l_{d3}, l_{d4}\}$ (quartiles) and a coarse regime with two labels $\{l_{c1}, l_{c2}\}$ (below and above median) such that l_{d1} and l_{d2} are a refinement of l_{c1} . Given this partition of the distribution we have $Pr(l_{d1}|l_{c1}) = Pr(l_{d2}|l_{c1}) = \frac{1}{2}$. Suppose the expected utility for the detailed labels is $\pi_{h|l_{d1}} = 5$ and $\pi_{h|l_{d2}} = 1$, then $\pi_{h|l_{c1}} = \sum_{l_d \in \mathcal{L}_{d|l_{c1}}} Pr(l_d|l_{c1}) \cdot \pi_{h|l_d} = \frac{1}{2} \cdot 5 + \frac{1}{2} \cdot 1 = 3$.

and highest values associated to the feasible detailed label, namely $\pi_{h|l'_d} := \min\{\pi_{h|l_d}\}_{l_d \in \mathcal{L}_{d|l'_c}}$ and $\pi_{h|l'_d} := \max\{\pi_{h|l_d}\}_{l_d \in \mathcal{L}_{d|l'_c}}$. This weak prediction represents a necessary but not sufficient condition to test whether an agent updates beliefs after receiving a signal.

The second prediction also requires that the agent has correct beliefs about the probability distribution. This assumption allows to pin down the exact value of the coarse label as a weighted average of the coarse ones, with weights represented by the conditional probabilities.

Testable Predictions. To recap, we have two predictions that we can test using the experimental dataset.

Prediction 1. Weak restrictions on label weights (inequalities).

Under the assumptions of Bayesian updating and any beliefs about the calorie distribution, we have the inequality prediction for the value of the coarse label $\pi_{l'_c}$:

$$\pi_{h|l'_d} \leq \pi_{h|l'_c} \leq \pi_{h|l'_d} \quad (6)$$

with boundaries defined as $\pi_{h|l'_d} := \min\{\pi_{h|l_d}\}_{l_d \in \mathcal{L}_{d|l'_c}}$ and $\pi_{h|l'_d} := \max\{\pi_{h|l_d}\}_{l_d \in \mathcal{L}_{d|l'_c}}$.

Prediction 2. Strong restrictions on label weights (equality).

Under the assumptions of Bayesian updating and correct calorie distribution, we have the equality prediction for the value of the coarse label $\pi_{l'_c}$:

$$\pi_{h|l'_c} = \sum_{l_d \in \mathcal{L}_{d|l'_c}} Pr(l_d|l'_c) \cdot \pi_{h|l_d} \quad (7)$$

Preferences over Labels. Another prediction of the Bayesian updating model is that agents prefer more detailed signals, as they are weakly more informative in the Blackwell sense (Blackwell, 1953). The additional information provided by the detailed label can allow to make better choices. A consumer could be indifferent between two products when only coarse labels are available (based on taste and price only) but if the product have different detailed labels the consumer would break the tie in favor of the preferred one. This can be easily shown using Jensen's inequality: the coarsening of the detailed labels cannot benefit, and possibly harms, the selection of the expected utility-maximizing product.

Consider a choice set with J products indexed $j \in 1, 2, \dots, J$, and compare the labeling regime C (relatively coarse) and its partition D (detailed). The discrete choice problem is the same one described at the beginning of this section. The agent can observe p and δ , and form expected utilities $\pi_{h|l} = \mathbb{E}_h[u|l]$. The expected value under the coarse regime can be

expressed as convex combination of values in the detailed regime. Under coarse labeling the health consequences cannot be separated among products with the same label

$$EU_c^* = \max_j \{ \sum_{l_d \in \mathcal{L}_{d|l_{jc}}} Pr(l_d|l_{jc}) \cdot v_j|l_d \}_j \quad (8)$$

whereas the differences can appear under detailed labels, and the maximization uses the granular information

$$EU_d^* = \sum_{l_{1d} \in \mathcal{L}_{d|l_{1c}}} \sum_{l_{2d} \in \mathcal{L}_{d|l_{2c}}} Pr(l_{1d}, l_{2d}|l_{1c}, l_{2c}) \cdot \max_j \{ v_j|l_{jd} \}_j \quad (9)$$

and by Jensens's inequality $EU_d^* = \mathbb{E}[\max_j(v_j)|l_d] \geq \max_j(\mathbb{E}[v_j|l_d]) = EU_c^*$.

Prediction 3. Preference for granular labels.

In a choice between nested labeling regimes, the rational decision maker prefers the most detailed label available as it generates the highest expected utility. This prediction holds under the assumptions of Bayesian updating and known meaning of the calorie-label mapping, and for any beliefs about the calorie distribution.

I showed in Section 3.3 that participants in the study deviate from this prediction. Only 31% indicate the detailed labels as most preferred ones in the ranking task. In the pairwise choices, both binary and coarse labels are preferred to the detailed ones by the majority of the subjects.

4.3 Estimation

I use the experimental dataset to estimate the label effect on choice in each treatment and test the predictions of the Bayesian updating model. The experiment provides the exogenous variation in the granularity of the label across subjects. For example, a product with 60 calories per serving is presented, according to the treatment, without label (Absent), or with a label that reports *Low Calorie* (Binary, range 50-150 calories), *Very low calorie* (Coarse, range 50-110 calories), or *3% calorie Daily Value* (Detailed, range 50-70 calories). By updating the beliefs, this product should become more appealing for consumers who prefer lower calories.

In order to test the predictions on the expected utility I estimate the demand model

$$u_{ijt} = \underbrace{\lambda_{l_{jt}}}_{\text{label}} + \underbrace{\delta_j}_{\text{experience/taste}} + \underbrace{\alpha p_{jt}}_{\text{price}} + \underbrace{\varepsilon_{ijt}}_{\text{Preference shock}} \quad (10)$$

The empirical model in Equation 10 is conceptually analogous to the theoretical model in Equation 1, but it relies on various homogeneity restrictions that are motivated by the features of the dataset. First, there is no explicit outside option in the study so I cannot separate the baseline product utility and the baseline health/label utility.⁴⁵ For clarity in the empirical model I express the first component of the utility as $\lambda_{l_{jt}}$ instead of $\pi_{h|l_{jt}}$ and I use the restrictions that the label utility derived from not observing any label is normalized to zero, and the average label utility within each treatment is also normalized to zero. Second, each subject is assigned to one treatment only, so it is not possible to separately estimate subject-level and treatment-level coefficients. To address this limitation I drop the subject-specific coefficients i and I assume homogeneous utility from labels λ_l , products δ_j , and price elasticity α . Different preferences across subjects and across trials are modeled using the preference shock term ε_{ijt} that I assume takes a Type 1 Extreme Value distribution. Finally, each product displays only one label in each treatment. There is no deception in the experiment, so a 60 calories product is never presented as high-calorie. If each product was presented with multiple types of labels in each treatment we would be able to estimate the label-product specific fixed effect, that captures how much subjects learn about each product through the label. That would be represent a stronger test, as it would not rely on the assumptions that, absent the label, subjects have the same initial beliefs for all the products.

The predictions introduced in the previous section can be now expressed with the notation of the empirical model.

$$\lambda_{l'_d} \leq \lambda_{l'_c} \leq \lambda_{l'_a} \quad \forall \lambda_{l'_c} \quad (11)$$

$$\lambda_{l'_c} = \sum_{l_d \in \mathcal{L}_{d|l'_c}} Pr(l_d|l'_c) \cdot \lambda_{l_d} \quad \forall \lambda_{l'_c}. \quad (12)$$

A possible concern is that, even without labels, subjects hold different beliefs about the calorie content for different products.⁴⁶ For example, consumers might have precise initial beliefs about high-calorie products but poorly calibrated beliefs for low-calorie ones. This does not affect the weak predictions in Equation 11; the bounds still hold if the agents has different subjective beliefs for each product and treatment, as long as they implement

⁴⁵For each product category, cereals and snacks, I normalize the value of the product with the highest amount of calories to 0.

⁴⁶A different concern is that the tests are presented for the aggregate estimates. The fact that the experiment uses a between-subject design does not allow to test the restrictions at the individual level. Choice probabilities from logit (or probit) models are nonlinear functions of the label fixed effect, so I cannot express the aggregate weights as a linear function of individual weights.

Bayesian updating.

4.4 Estimation Results

I estimate the parameters $(\lambda_l, \delta_j, \alpha)$ using a maximum likelihood estimator, and I estimate the standard errors for the same coefficients using the bootstrap method.⁴⁷ The estimates of the label fixed effects λ_l are shown in Figure 16.




Figure 16: Estimates of the label fixed effect. Values estimated by maximum likelihood and bootstrap standard errors. ML label effect expressed in WTP λ_l/α , bootstrap SE. Blue lines indicate the groups of feasible detailed labels associated with every coarse label.

The weak test on the Bayesian updating model restrictions requires the coarse label estimates to be within the range of feasible detailed label estimates.⁴⁸ The prediction is systematically violated for all the four coarse-detailed labels comparisons (all $p < 0.01$). I defer a test of the full model restriction to the next subsection.

The label estimates λ_l are scaled by the price elasticity α to facilitate the interpretation.⁴⁹ Based on the estimates, subjects in the study are willing to pay about \$1.80 more for the same product when they observe the lowest-calorie binary label, \$3.40 for the lowest-calorie coarse label, but only \$2 more when they observe the lowest-calorie detailed label, with symmetric results for the highest values (binary -\$1.80, coarse -\$4.60, detailed -\$2.90).

⁴⁷For the bootstrap estimation of the standard errors I use 1001 simulated dataset generate by random sampling with replacement from the original one.

⁴⁸We can also test the restrictions binary-coarse and binary-detailed. The prediction is systematically rejected in the first case, but cannot be rejected in the second one. Binary and detailed labels' estimates have similar magnitude, but both of them have smaller values than the coarse ones.

⁴⁹Prices vary in the experiment independently from other characteristics. This allows the identification of the price elasticity parameter.

4.5 Estimation of the Restricted Models and Model Selection

The Maximum Likelihood estimates show that the coefficients for the label weights in different treatments violate even weak predictions on the relation between treatments. Here I consider an alternative approach to test the predictions by estimating different models, with and without restrictions. I run a model selection exercise to formally test the restrictions generated by the Bayesian updating model. I consider three nested model specifications.

1. Unconstrained model.
2. Constrained model with inequalities (prediction 1).
3. Constrained model with equalities (prediction 2).

The equalities and inequalities for Models 2 and 3 follow the restrictions presented in Equations 11 and 12 and provide a formal test for the restrictions.

Table 8 contains the maximum likelihood parameter estimates for different model specifications. The estimate of the price elasticity α is 0.052, with no significant differences across models, a low coefficient compared to what usually found in the empirical literature.⁵⁰ The table shows the estimated label fixed effects λ_i for the two most extreme labels in each treatment, expressed in terms of willingness to pay and with the coefficients shown in Figure 16 for model 1. The restrictions lead to very different estimates in the other models, with larger weights of the detailed labels (in all the models) or lower weights of the coarse labels (in model 3). Finally I show the summary statistics, average and standard deviation, for the product fixed effects. The product with the highest number of calories in each category is the implicit outside option of the choice model with $\delta_j = 0$. Positive average parameters confirm that even without labels products with lower calories are preferred, with large dispersion of the evaluations within each category.

Table 9 contains information about the model fit: maximized value of the log-likelihood of the data LL and the number of degrees of freedom dof . The Likelihood Ratio Test for model 3 rejects the equality restriction with respect to the unrestricted model 1 (6 dof, $p < 0.001$). The test cannot be run between model 1 and 2, even if the second is nested, because the two models have the same number of degrees of freedom.⁵¹ For this reason, I use the Bayes Information Criterion (BIC) statistic to compare all the models. The BIC is

⁵⁰In their study on response to the Chilean warning label reform, Barahona et al. (2020) estimate $\alpha = 0.05$ for products expressed in the amount of 100 grams, while the products used in the experiment are in the range 16-32 ounces, corresponding to 450-900 grams.

⁵¹A test of the equality restriction between model 2 and model 3 also rejects the equality restriction (6 dof, $p < 0.001$).

	(1)	(2)	(3)
Price Elasticity			
α	0.0520 (0.0068)	0.0520 (0.0081)	0.0519 (0.0089)
Label FEs			
λ_{B1}/α	1.76 (0.39)	1.96 (0.48)	2.00 (0.47)
λ_{B2}/α	-1.76 (0.39)	-1.96 (0.48)	-2.00 (0.47)
λ_{C1}/α	3.43 (0.68)	3.63 (0.76)	2.61 (0.70)
λ_{C4}/α	-4.62 (0.74)	-4.16 (0.87)	-2.73 (0.70)
λ_{D1}/α	1.96 (1.18)	3.63 (0.94)	3.77 (1.12)
λ_{D10}/α	-2.93 (1.14)	-4.16 (1.08)	-3.65 (1.15)
Product FEs			
$\delta_{cereals}/\alpha$ - Avg.	1.07	0.90	1.26
$\delta_{cereals}/\alpha$ - St.Dev.	4.10	4.16	4.13
δ_{snacks}/α - Avg.	0.43	0.25	0.61
δ_{snacks}/α - St.Dev.	3.65	3.63	3.65
Observations	68,454	68,454	68,454

Table 8: ML estimates in different model specifications and bootstrap standard errors. Price elasticity α ; label fixed effects λ_l for the two labels at the extremes of the range in each treatment (e.g., B1 indicates the first label of treatment Binary, that is the one associate with lower calories); average and standard deviation of the product fixed effects δ_j for the two categories of products used in the study.

defined as

$$BIC = -2LL + dof \cdot \log N_{obs}.$$

The BIC penalizes the use of additional free parameters and allows to compare how well different models are fitting the data. The unrestricted model (model 1) is the preferred one under the BIC, as it leads to the lowest BIC value (138,869). Models 2 and 3, that capture the two types of restrictions provided by the Bayesian updating model, have similar BIC (138,896 and 138,893) and the higher log-likelihood generated by model 2 is offset by the additional degrees of freedom by relaxing the equality restrictions. I calculate the Bayes factor for the models in the table with respect to the unrestricted model. The Bayes factor is the multiplicative factor by which the relative posterior probability that $m2$ rather than

$m1$ is the correct model of the data is increased by the observations in the dataset (see Burnham and Anderson (2002)). For any two models $\mathcal{M}1$ and $\mathcal{M}2$ the Bayes Factor K is defined as

$$\log K_2 = \frac{1}{2}[BIC(\mathcal{M}1) - BIC(\mathcal{M}2)].$$

A Bayes factor of -12 points provides strong evidence against model 3. The data increase the relative posterior probability of the unrestricted model being the correct one by a factor of more than 150,000.⁵² The results are analogous for model 2, that considers a weaker restrictions by introducing the inequality bounds (weak prediction from the Bayesian updating model) and has a -13.7 Bayes factor.

	Model	-LL	dof	BIC	Bayes Factor
1	Unrestricted	69,145	52	138,869	
2	Restricted (inequalities)	69,159	52	138,896	-13.66
3	Restricted (equalities)	69,190	46	138,893	-11.97

Table 9: MLE and model selection. LogLikelihood, degrees of freedom, Bayes Information Criterion, and Bayes Factor for the models discussed in text. Bayes Factor computed with respect to the unrestricted model.

The evidence against the Bayesian updating model opens a follow-up question on what model would be consistent with the experimental results. In the next section I explore a generalization of the updating model that introduces precision overload and includes endogenous accuracy in the decision process.

5 Limited Attention Models and Precision Overload

In this section, I introduce a generalized version of the discrete choice model that includes limited attention. This model allows for violations of the Bayesian decision theory predictions and can capture the main stylized facts from the experiment.

The three main facts from the experiment that I want to explain are U-shaped label effect, diminishing sensitivity to information, and preference for coarse labels. First, increasing the granularity of the labeling regime leads to a U-shaped effect on calories. Second, subjects are more sensitive to the calorie messages of coarse labels compared with detailed labels. Third, in the experiment coarse labels with four possible messages are preferred to other labels with

⁵²A Bayes factor lower than -10 (higher than +10) is typically considered strong evidence against (in favor of) the alternative model, as it corresponds to a ratio of the likelihood equal to e^K . This result multiplied by the ratio of the prior probabilities of the two model being correct is equal to the ratio of the posterior probabilities of the models given the observations available.

more and fewer levels. The experimental data reject the Bayesian framework, as detailed labels are less informative than coarse ones. Classic attention models like salience (Bordalo, Gennaioli and Shleifer, 2012) and rational inattention (Sims, 2003; Caplin and Dean, 2015) also fail to capture the main results.

I generalize the Bayesian framework by relaxing the assumption that agents have perfect accuracy in observing and processing the information contained in the label messages. Consumers may have a limited amount of cognitive resources that can be used to understand the information contained in the label. I consider an encoding/decoding model (Paninski, Pillow and Lewi, 2007; Polania, Woodford and Ruff, 2019) in which the label information is retrieved with noise. I assume that part of the cognitive resources are depleted in the process of recognizing the label, before the inference process. The key assumption of the model is that the features of the label, for example the level of detail, determines the resources available to reduce the noise. Precision overload emerges when detailed labels are harder to process, reduce attention resources, and lead to larger noise. I conclude the section with a numerical exercise that shows how this approach can capture the three main features of the experimental dataset.

Before proceeding, I consider two classic limited attention models in the context of coarse label information. Attention — and lack of it — have been extensively studied and discussed in the psychology literature and have recently been incorporated in economics models. Salience and rational inattention provide two leading examples; they have been used to explain a variety of empirical and experimental features of human behavior, yet they are not able to capture the experimental results presented in this paper.

5.1 Coarse Labels in a Salience Model

In a salience model of choice (Bordalo, Gennaioli and Shleifer, 2013) the decision maker’s attention is drawn to salient attributes of products. The model introduces a context-dependent value representation, in which the true evaluation is distorted in favor of salient features.

In the framework adopted to interpret the choice problem, products have various attributes: price, calories, and taste. By introducing salience, large calorie differences across products generate higher attention, leading a larger weight of the calorie attribute in the product evaluation.⁵³ The consumer might find the calorie amount of a product salient in a choice set where calories are dispersed across products, and not salient in a different choice set where products have similar calorie content.

⁵³I use calories as an attribute to refer to the health consequences from the consumption of calories, that is part of the utility function discussed in the discrete choice model.

A salience model that distorts attention towards large calorie differences would not be able to capture the lower sensitivity to detailed label information. The adoption of detailed labels would, on average, increase the salience of the calories and distort the evaluation in favor of the low-calorie product, the opposite of what I observe in the experiment. A granular labeling regime creates both larger and smaller calorie differences across products, but the higher weight on large differences would exceed the lower weight on small differences. Consider an extension of the choice model that introduces salience in favor of values far from the reference. Salience being absent, the decision maker would observe a pair of products and calculate the expected value for each of them

$$\mathbb{E}[v_j|l_j] = \mathbb{E}[h|l_j] + \delta_j + \alpha p_j$$

Consider the choice between two products $j = 1$ (low-calorie) and $j = 2$ (high-calorie). The agent makes a choice based on Δv , the difference in the evaluation of the two products

$$\Delta v := \mathbb{E}[v_1 - v_2|\bar{l}] = \mathbb{E}[\Delta h|\bar{l}] + \Delta\delta + \alpha\Delta p \quad (13)$$

where \bar{l} indicates the vector of labels displayed by the products, and Δ indicates the difference in the expected values between the two products for each attribute.

In a simple salience model, the weight of the calorie is distorted based on how extreme the values are within the choice set. In the case with two products, we can restrict to the analysis of the differences. The difference in values now becomes

$$\Delta v = S(\mathbb{E}[\Delta h|\bar{l}]) + \Delta\delta + \alpha\Delta p \quad (14)$$

where $S(x)$ is a function that captures the effect of salience. For this simple analysis I am focusing on the feature that more extreme values increase even further the weight of the attribute.⁵⁴ A function with this property is convex for $x > 0$ and concave for $x < 0$. I consider the probability of choosing the low-calorie product, by rearranging the order of the products such that $\mathbb{E}[\Delta h|\bar{l}] \geq 0$. By doing so I can focus on the convex part of the salience function.

In the experiment, I manipulate the granularity of the label. This corresponds to a change in the distribution of inferred health differences $\mathbb{E}[\Delta h|\bar{l}]$. The distribution changes for two

⁵⁴For the purpose of this analysis I am making two simplifications with respect to the model discussed in Bordalo, Gennaioli and Shleifer (2013), yet preserving the general spirit of the model. First, I consider absolute differences, not relative differences. I am abstracting from the feature that values are perceived in relative terms. Second, I consider the salience of the calorie attribute, without changing the weights of the other attributes such as taste and price, that face no information change across treatments.

main reason. First, two products can have the same coarse label l^c (zero inferred difference) and different detailed labels l^d (strictly positive difference). This leads to an unambiguous increase in the magnitude and salience of the health attribute.

$$l_1^c = l_2^c \Rightarrow \mathbb{E}_c[\Delta h|\bar{l}^d] \geq \mathbb{E}_c[\Delta h|\bar{l}^c] \Rightarrow S(\mathbb{E}_c[\Delta h|\bar{l}^d]) \geq S(\mathbb{E}_c[\Delta h|\bar{l}^c]) \quad (15)$$

Second, two products can have different coarse labels, and therefore also different detailed labels. On average, the health consequences are the same

$$l_1^c \neq l_2^c \Rightarrow \mathbb{E}_{l^d}[\mathbb{E}[\Delta h|\bar{l}^d]|\bar{l}^c] = \mathbb{E}[\Delta h|\bar{l}^c]$$

The salience function $S(\cdot)$ is convex in the positive support, so I can apply Jensen's Inequality:

$$l_1^c \neq l_2^c \Rightarrow \mathbb{E}_{l^d}[S(\mathbb{E}[\Delta h|\bar{l}^d])] \geq S(\mathbb{E}_{l^d}[\mathbb{E}[\Delta h|\bar{l}^d]|\bar{l}^c]) = S(\mathbb{E}[\Delta h|\bar{l}^c]) \quad (16)$$

This indicates that, for products that display different coarse labels, moving to a more granular labeling regime will increase, on average, the perceived calorie difference. By combining Equations 15 and 16 I obtain, that, on average, the detailed labels generate larger sensitivity to the calorie information

$$\mathbb{E}_{l^d}[S(\mathbb{E}[\Delta h|\bar{l}^d])] \geq S(\mathbb{E}[\Delta h|\bar{l}^c]) \quad (17)$$

and this contradicts the experimental results, that show a decrease in the sensitivity to calories.

A more general version of the model could assume that the consumer evaluates multiple attributes of the product (e.g., nutrients, price, taste) and the calorie attribute is salient when the consumer is prompted to think about it. For example, the consumer would normally not consider the calorie amount, but the fact that the products display a calorie label on the front of the package is sufficient to make this attribute salient. This setup would be more similar to the concept of salience as discussed in Chetty, Looney and Kroft (2009) in the case of salient taxes. This version of the salient model would also not provide an explanation for the experiment results, as the effect occurs in the comparison across labels with different coarseness.

5.2 Coarse Labels in a Rational Inattention Model

Rational inattention models like the ones used by Sims (2003) and Caplin and Dean (2015) look at attention as a costly resource that can be allocated to increase the accuracy of the beliefs. Even when all the information required to determine the best action are available,

the agent can make costly mistakes because being more accurate requires cognitive resources that are also costly. In the food choice problem, the fully attentive consumer would be able to consider all the information provided about price, taste, and healthiness of each product and pick the one that best matches own taste. Instead, the inattentive consumers would have a noisy representation of the information available.

The noisy framework follows closely the Bayesian framework, with the difference that the posterior health beliefs might differ from the original ones because of the noise in the representation of the true state.

$$\underbrace{c}_{\text{Calories}} \rightarrow \underbrace{l}_{\text{True label}} \rightarrow \underbrace{\hat{l} \sim G_l}_{\text{Noisy label}} \rightarrow \underbrace{F_{c|\hat{l}}}_{\text{Beliefs (calories)}} \rightarrow \underbrace{F_{h|\hat{l}}}_{\text{Beliefs (health)}} \rightarrow \underbrace{\mathbb{E}[h|\hat{l}]}_{\text{Expected health}}$$

Each product displays a label l based on its true number of calories c (unobserved), and the label information allows to refine the beliefs over the calorie amount. In the previous setting without uncertainty, the agent is fully accurate about l and updates own beliefs about the calorie $F_{c|l}$.⁵⁵ In the general case with uncertainty, the agent might not recall the label accurately. The true label is replaced with a noisy label $\hat{l} \in \hat{\mathcal{L}}$ generated according to the process G_l . The agent is aware that \hat{l} might not be accurate, knows the distribution of the noisy retrieval G_l , and makes optimal use of the noisy label information to update the beliefs about calories (posterior beliefs distribution $F_{c|\hat{l}}$) and about the health (posterior beliefs distribution $F_{h|\hat{l}}$). After observing \hat{l} , the expected health-utility component is $\mathbb{E}[h|\hat{l}]$. In the standard Bayesian framework, introducing a more granular labeling regime reduces the uncertainty about the calorie content and generate more accurate beliefs about the health consequences. This is not necessarily true in the generalized model, as the labeling regime might also affect the noise in the rest of the inference process.

The discrete choice problem introduced in Section 4 can be expressed in the general way to allow for noise in the representation of the label. The decision maker selects the product j out of the choice set \mathcal{J} that maximizes the expected utility based on the noisy labels available.

$$EU^*(\mathcal{J}, \{\hat{l}_j\}_j) := \max_{j \in \mathcal{J}} \mathbb{E} \left[h(c_j) + \delta_j + \alpha p_j |\hat{l}_j| \right] = \max_{j \in \mathcal{J}} \left[\delta_j + \alpha p_j + \int_l \int_c h(c) dF_{c|\hat{l}} dF_{l|\hat{l}} \right] \quad (18)$$

The conditional distribution $F_{l|\hat{l}}$ indicates what is the probability that the true label is l when the agent recalls \hat{l} . It follows Bayes rule from the noisy label distribution $G_l(\hat{l})$ and

⁵⁵In the general case the beliefs distribution $F_{c|\hat{l}}$ is a mixture of the posterior distributions $F_{c|l}$ under different realizations of the perceived label.

the true distribution of labels across products F_l , with $f(l|\hat{l}) = \frac{g_l(\hat{l})}{\sum_l g_l(\hat{l})} \cdot f_l$.

Does this framework provide a good explanation for the results of the experiment? Two sets of results are at odds with the prediction of a rational inattention model. First, we should not find differences across treatments when the subjects consult the back-of-package nutrient information. Second, participants should prefer detailed information to coarse ones.

The first prediction refer to the choice behavior conditional on the back-of-package information being observed. In all treatments, the nutrition facts on the back-of-package are a summary statistic for the front-of-package labels. In a rational inattention model, the posterior beliefs distribution depends on the available information about the state of the world. It follows that the distribution of beliefs over calories will be the unchanged if the agent observes the back-package information alone, or jointly with any other (coarser) signal. In the study I observe the choice probability for products in different treatments, that I assume depend on consumers' preferences and information received.⁵⁶ If the participants who observe the back-of-package information across treatments have the same preferences, we should expect them also to display the same choice distribution.

The experimental results displayed in Figure 10 show that choice probabilities in trials where subjects collect back-of-package information are different across treatments, a violation of the prediction above.

The second prediction refers to the preference for receiving detailed information. In rational inattention model, the attention cost depends on the change in the action distribution as a response to the additional information. The attention cost that these models use refers to the accuracy of the information *used* in the decision process, not to the information *available*. No direct cost is associated with searching or having access to more detailed information. The agent can always drop the level of detail in excess and coarsen the information with no additional cost. The introduction of a more detailed information structure would weakly increase the net expected utility of a rationally inattentive agent, and the agent would weakly prefer receiving the most detailed label available. This implies that the agents in the choice task should consult the back-of-package information, that is always available across treatments, and should prefer detailed labels when asked explicitly to indicate own preferences. Both of these predictions fail in the experiment: 66% of the participants consult the back-of-package information in less than 5% of the trials, and 69%

⁵⁶In principle it is not guaranteed that the preference distribution is identical across treatments. Participants might self-select in the collection of detailed information. Based on the observed behavior in the experiment, this might be relevant for the comparison of treatment Absent with the other treatments that contain the front-of-package label, but not across the three label treatment. This is because there are no significant differences in the fraction of participants who consult the back-of-package information across the three label treatments.

of the participants prefer a type of label different from the most detailed one.

The discrete choice model can be extended in order to allow for multiple types of costs associated with the collection and use of new information.⁵⁷ The first result suggests that consumers might face a search cost associated with flipping the package and consulting the back-of-package information. A natural extension of the rational inattention model would include both costs, for the collection and use of information. Costly search is a common assumption in models in which uncertainty can be gradually resolved (e.g., Weitzman 1979). But costly search would not be sufficient to characterize the second result. If the information on the front of the package is equally costly, regardless of its precision, the decision maker should prefer the most granular and informative one. In the next section I maintain a framework with multiple information costs. Instead of considering the search cost related to the information collection, I use a framework with costly encoding and decoding of information to introduce a cost associated with the collection of more precise information.

5.3 Statistical Model of Encoding and Decoding

I consider a model of encoding and decoding of the label information with two types of attention cost. The information is stored in memory, and then recalled to evaluate the product, and both these processes are noisy. The decision maker has limited cognitive resources that they can use to increase the accuracy. A richer labeling regime requires more resources to achieve an accurate encoding, with fewer resources available to reduce the noise in the retrieval phase, and this can generate the precision overload effect.

The decision maker faces noise in two separate phases of the process: first in the *storing* phase, when the label information is encoded in the memory, and later in the *retrieval* phase, when the decision maker recalls the information. Both phases can be noisy, and the Bayesian agent is aware of both natures of noise. The noisy storing represents an additional phase in the evaluation process, with \tilde{l} being the information stored in memory according to the distribution $G_{l,\tilde{l}}$.

$$\underbrace{l}_{\text{True label}} \rightarrow \underbrace{\tilde{l} \sim G_{l,\tilde{l}}}_{\text{Stored label}} \rightarrow \underbrace{\hat{l} \sim G_{\tilde{l},\hat{l}}}_{\text{Retrieved label}} \rightarrow \underbrace{F_{c|\hat{l}}}_{\text{Beliefs (calories)}} \rightarrow \underbrace{F_{h|\hat{l}}}_{\text{Beliefs (health)}} \rightarrow \underbrace{\mathbb{E}[h|\hat{l}]}_{\text{Expected health}}$$

Each phase is associated with a cognitive cost that depends on the degree of accuracy of the storing and retrieval. The agent has a finite amount of cognitive resources K that

⁵⁷Multiple costs can also be used to characterize the collection of different pieces of information about the product. For example, the agent might face an attention cost to collect information about the price, the calories, and other nutritional facts of the product. I will continue assuming that the agent has immediate access to the prices without attention costs, whereas calorie information are costly.

represent the attention budget and is shared between the two phases.

$$\underbrace{\beta_{\text{Encoding}} \cdot C_{\text{Encoding}}(G_{l,\tilde{l}})}_{\text{Cost for recognition}} + \underbrace{C_{\text{Decoding}}(G_{\tilde{l},\hat{l}})}_{\text{Cost for retrieval}} \leq \underbrace{K}_{\text{Budget}} \quad (19)$$

I assume that the two types of costs have the same functional form, but they might have different relative costs according to the parameter β_{Encoding} . In traditional rational inattention models $\beta_{\text{Encoding}} = 0$, the decision maker has perfect accuracy in the encoding phase, and selects the signal structure $G_{l,\tilde{l}}$ that maximizes the expected utility. For the functional form of the attention costs function I assume they are proportional to the mutual information $I(l, \tilde{l})$ and $I(\tilde{l}, \hat{l})$ between the two random variables

$$C_{\text{Encoding}} = I(l, \tilde{l}) = \sum_l \sum_{\tilde{l}} p(l, \tilde{l}) \log_2 \left(\frac{p(l, \tilde{l})}{p(l)p(\tilde{l})} \right) \quad (20)$$

and in the same way for the mutual information $C_{\text{Decoding}} = I(\tilde{l}, \hat{l})$ in the decoding phase. The recognition cost C_{Encoding} increases in the number of possible labels $N_L = |\mathcal{L}|$. When each label is equally likely, the highest encoding cost is $C_{\text{Encoding}}(G_{l,\tilde{l}}(\gamma = 1)) = I(l, l) = \log_2(L_N)$ and it represents the upper bound for the cost encountered under different distributions. The use of the binary logarithm provides an intuitive interpretation for the maximum attention capacity required in the encoding stage: doubling the number of labels requires β additional attention units.

The decision maker faces the discrete choice problem described in Equation 18 and maximizes the expected utility under the attention constraint of Equation 19. The model requires further assumptions about how the agent chooses the encoding and decoding distributions $G_{l,\tilde{l}}$ and $G_{\tilde{l},\hat{l}}$ within the set of feasible distributions. A full-fledged optimization model would determine the optimal distributions that maximize expected utility given the agents' preferences and the budget constraint. Without making further assumptions on the subjective preferences, in the next section I present a simple parametrization that prioritizes accuracy in the encoding phase.

5.4 Costly Information Recognition and Precision Overload

I introduce two assumptions on the process of noise reduction that allow to focus on a tractable model specification. First, I restrict the encoding and decoding mappings to a family of one-parameter distributions. The stored and retrieved labels are defined over the same set of possible label realizations $\tilde{\mathcal{L}} = \hat{\mathcal{L}} = \mathcal{L}$. The parameter γ indicates the accuracy of the representation, ranging between $\gamma = 1$ (perfectly accurate encoding or decoding) and

0 (completely noisy). The agents store the label accurately with probability γ ,⁵⁸ otherwise they are equally likely to store any other label.

$$p(\tilde{l} = l|l) = \gamma + \frac{1 - \gamma}{N_L}; \quad p(\tilde{l} = l'|l) = \frac{1 - \gamma}{N_L}, \quad \forall l' \neq l$$

where $N_L := |\mathcal{L}|$ is the number of existing labels.

Second, the decision maker prioritizes accuracy in the encoding phase. The cognitive resources are allocated first to achieve the highest encoding accuracy, and the remaining resources are then exhausted to increase the accuracy of the decoding process. The encoding accuracy $\gamma_{\text{Recognition}}$ maximizes the function

$$\gamma_{\text{Recognition}} = \max_{\gamma \in [0,1]} \gamma \quad \text{s.t.} \quad C_{\text{Encoding}}(G_{l,\tilde{l}}(\gamma)) \leq \frac{K}{\beta_{\text{Encoding}}}. \quad (21)$$

The remaining budget for the decoding phase is

$$K_{\text{Retrieval}} := K - \beta_{\text{Encoding}} C_{\text{Encoding}}(G_{l,\tilde{l}}(\gamma_{\text{Recognition}}))$$

and the decoding accuracy is determined from

$$\gamma_{\text{Retrieval}} = \max_{\gamma \in [0,1]} \gamma \quad \text{s.t.} \quad C_{\text{Decoding}}(G_{\tilde{l},l}(\gamma)) \leq K_{\text{Retrieval}}. \quad (22)$$

Finally, the agent makes Bayesian inference to determine the posterior distribution $F_{c|\tilde{l}}$ conditional on the retrieved signal. The introduction of the recognition cost can lead to precision overload, as more granular labels require more resources in the encoding phase and leave fewer resources available in the decoding phase. In the case of food choices, this leads to less accurate information about the calorie content, more noise, and lower sensitivity to the calorie content.

Is there evidence, from the experiment or the literature, supporting the precision overload effect? In both cases we have some indirect evidence. In the experiment, response time data suggest that detailed labels make the decision more difficult compared to the other conditions. Figure 9 shows that subjects take on average 10% more time ($p < 0.05$) to make a choice in Treatment Detailed, but there is no significant difference in the other three treatments (average response time 6 seconds per choice set). This effect is more broadly connected with the perception literature that analyzes the limits in the human ability to recognize and categorize stimuli (e.g., Pollack 1952).

⁵⁸The same accuracy function is applied to the encoding and the decoding phase, with different parameters $\gamma_{\text{Recognition}}$ (encoding) and $\gamma_{\text{Retrieval}}$ (decoding).

5.5 Numerical Exercise

A numerical exercise illustrates how the encoding-decoding framework with costly recognition can capture the three main experimental results. Consider a product choice scenario in which consumers maximize expected utility with respect to prices and calories (with linear preferences over calorie amount) according to the utility

$$EU_j = -\mathbb{E}[c|\hat{l}_j] - \alpha \cdot p_j \quad (23)$$

and choose one out of a pair of products available in each choice set.⁵⁹ The accuracy of the information recognition and retrieval are endogenous and depends on the budget constraint expressed in Equation 19 with mutual information costs as in Equation 20.

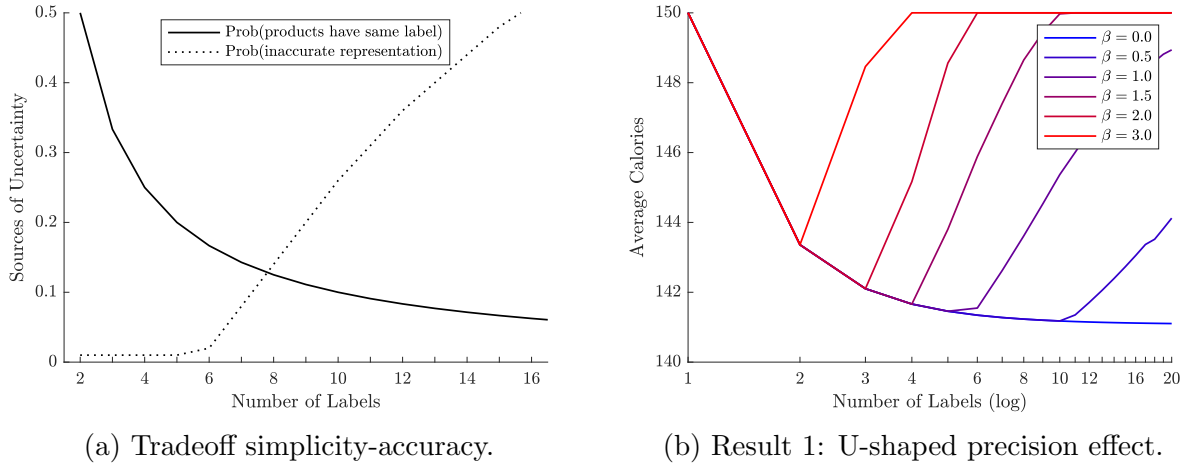



Figure 17: Numerical exercise. Left: Tradeoff between sources of uncertainty; parameters $K = 5$, $\beta = 1$. Right: Average calories per serving; budget $K = 5$, variable recognition cost β .

The numerical exercise allows to perform a comparative statics and study how higher precision (more labels) affects choices. The overall effect of precision depends also on the parameters K (budget) and β (relative cost of recognition), with $\beta = 0$ characterizing a standard inattention model without precision overload. Figure 17.a shows the tradeoff between the two effects of higher precision. The solid line indicates the uncertainty reduction, since fewer pairs of products will display the same label (making it difficult to identify what is the healthier option). The dotted line captures the precision overload: the decision maker can correctly identify up to five labels, and mistakes gradually increase with larger numbers.

⁵⁹I assume that calorie and prices are independent, and they are distributed according to a uniform prior. In the simulations I use $c \sim \text{Unif}(50, 250)$ calories, $p \sim \text{Unif}(4, 6)$ dollars, price elasticity $\alpha = 200$.

The three main results are shown in Figures 17.b, 18.a and 18.b. First, the U-shaped precision effect. When the retrieval noise becomes too large compared to the precision of the label, the number of calories increases even for agents who prefer fewer calories. Reducing the recognition cost parameter delays the saturation point and diminishes the severity of the overload problem, and when $\beta = 0$ calories diminish monotonically in the precision of the label.

The second result is the diminishing sensitivity to information. Figure 18.a shows that the probability of choosing the low-calorie product, after controlling for the inferred calorie difference, is lower when agents face more detailed labels. Also in this case the effect would disappear without precision overload: if we remove the recognition cost we see that choice probabilities are the same across labeling regimes, and only depend on the x-axis difference.



(a) Result 2: diminishing sensitivity to information. (b) Result 3: preference for coarse labels.

Figure 18: Numerical exercise. Left: Response to implied calorie differences; budget $K = 5$, recognition cost $\beta = 1$. Right: Average calories per serving; variable budget K , recognition cost $\beta = 1$.

The third result is that subjects in the experiment prefer coarse labels with intermediate level of precision. Figure 18.b shows that agents with precision overload also prefer intermediate level of precision. This is an immediate consequence of the first result under the assumption that the cognitive constraint is binding.⁶⁰ For each value of K we can determine the optimal number of labels $N_K^*(\beta)$ as a function of the marginal recognition cost $\beta > 0$.⁶¹

⁶⁰The attention model introduced by Lipnowski, Mathevet and Wei (2020) can generate the U-shaped effect on quality of choice but it still predicts preference for more granular information. This is because decision makers maximize the net utility as difference between value of the chosen option and attention cost.

⁶¹A fully attentive agent face diminishing benefit from increasing the number of labels (Figure 17.a). A boundedly rational agent faces a marginal loss from the decrease in accuracy, and the optimal number of labels is weakly lower than the value that equalize benefits (for the optimal agent) and loss. A fully attentive agent in this setup has infinite cognitive resources K , and the marginal attention cost β is irrelevant. An

The optimal — and preferred — number of labels also depends on the cognitive budget constraint K , and lower K leads to 1) weakly worse choices for every level of N_L and 2) a weakly lower optimal number of label. Figure 18.b compares two agents, identical in all but the budget constraints, across a range of labeling regimes. The agent look similar in the way that both benefit from more precise labels up to a certain point, and gradually make worse choices after. But they differ in the location of the constrained optimal number of labels, with the low-budget agent preferring three categories whereas the high-budget agent prefers ten. In the experiment, subjects with low food literacy show even stronger preference for coarse labels (two or four categories), and the same pattern emerges for subjects with low socioeconomic status (low income or low education). A possible interpretation is that these subjects are not necessarily different in terms of preferences, but might find it more difficult to understand and make good use of the detailed information.

6 Discussion

This section discusses alternative interpretations for the consumers’ choice problem and how this framework can be nested in a broader strategic environment that includes policymakers and companies. First, I focus on consumers. A different interpretation of labels as information on the products’ health consequences (instead of their calorie content) leads to my results. Modeling precision overload as noise in the understanding of the mapping also leads to similar implications. I discuss the role of endogenous attention as information acquisition of products’ characteristics and consideration sets. After, I discuss the labeling problem from the perspective of the policymaker. The adoption of coarse labels depends on the alignment of consumer and regulators’ objectives and the consumers’ ability to process detailed labels. Finally, I consider the firms’ strategic response to the introduction of labels. Label coarseness affects the incentive to reformulate the product composition.

6.1 Alternative Interpretations of Label Information

Interpretive labels, contextual information, and normative connotation.

Labels can provide information about the product characteristics (as in the main model), about the health consequences (as discussed in the paragraph above), and more, and these approaches are not mutually exclusive. Ikonen et al. (2020) classifies labels as reductive or interpretive, and I discuss this distinction in details in the Appendix. In general, reductive

attentive agent with $\beta = 0$ but finite K would be able to process the signals with perfect accuracy only up to a finite number.

labels contain the same numerical information as the Nutrient Facts panel, and interpretive labels provide an evaluation of the information using colors, stars, or health symbols. Interpretive labels may have more than one type of interpretation, and all of them may help consumers on top of informing about the nutrition content. First, they may provide contextual interpretation about the typical range of values - are 100 calories high or low compared to other products? Second, they may provide a normative interpretation of the nutrient information - are 100 calories a healthy amount compared to a recommended diet? The consumer may look for recommendations from the regulator about which products are healthy. Finally, interpretive labels may be simpler to understand because they are coarser than the exact nutrient amount. This is the main function discussed in this paper. The instructions of the experiment aim to address the previous two to focus on the coarseness dimension. Nevertheless, the normative connotation of categorical labels might be part of the difficulty to compare coarse and detailed labels and contribute to the experimental results.

Labels as signals about health consequences.

Sections 4 and 5 focus on a specific choice problem in which consumers use the label to update their beliefs about the calorie content of products. Using a different interpretation of the role of labels, consumers can learn directly about the health consequences of consuming a product instead of indirectly about the calorie content. Previously I described consumers who do not know the exact calorie content c of the products but understand the mapping calorie-utility. Here I consider consumers who do not observe the mapping $h(c)$ and learn directly about health consequences. What if labels provide direct information on health instead of on calories? In the Appendix, I consider the setup in which consumers update $F_{h|l}$ directly from the label instead of indirectly through the beliefs $F_{c|l}$. This setup is motivated by the fact that consumers might not know the exact meaning of 100 calories but can understand that Low-Calorie is associated with a range of health consequences. Even if the interpretation is different, the general results for the Bayesian agents are the same as in Section 4. First, more detailed labels provide the same information and more by partitioning messages and refining the posterior beliefs. Second, consumers who understand how labels are created would prefer the most detailed messages. Third, the discrete choice model can be estimated using the same procedure as before, and this allows to test the necessary condition on Bayesian updating. Finally, consumers with a recognition cost may prefer coarser labels to invest more cognitive resources in the accurate interpretation of the label signal.

Labels as signals of health-calorie mapping.

What if label give indirect information about the health consequences? Instead of providing

a signal about $h(c)$, labels could inform the consumers about the health function $h(\cdot)$ itself. In Section A.3 I discuss this alternative approach that follows the approach introduced in Kamenica (2008). The consumers might be uncertain the shape of the function, or about one parameter that determines the mapping. I discuss in the Appendix why this approach alone would also not be sufficient to generate lower sensitivity to calorie information under detailed labels.

We face an interesting case if we combine the two different interpretations of labels as signal for the state (calorie content) and signal for the mapping (meaning of the calorie content). As discussed above, labels in the real world typically aim to cover both these roles. A possible concern is that the treatments I use in the experiment are *perceived* by the subjects as if they were covering different functions, despite the instructions, comprehension questions, and information presented in each page of the choice task. Suppose the subjects interpret that numerical labels as signals of the calorie content, and the categorical labels as signals about the calorie mapping, or even recommendations with a normative — not descriptive — meaning. From an experimental perspective, this suggests that further tests are required to disentangle the effects (manipulate precision and type of label separately) and test the mechanism. As a suggestive evidence in support to this hypothesis, I consider the additional data collected in the study. After the choice task, I elicit preferences over calories. The preferred calorie content is not significantly different between treatments Absent and Detailed, but I observe a small yet significant difference when I compare Absent with the treatments Binary and Coarse. In both cases, subjects' preferences are skewed more towards low-calorie products, suggesting that the label information affected their calorie preferences.

This possible setup suggests a follow-up question about what type of information regulators want to provide to the consumers through labels. The experimental results suggest that regulators might face a tradeoff. Coarse labels might be more effective when consumers are not sure about own preferences (e.g., for health- and environment-related choices). Detailed labels, instead, might be preferable when consumers know how to interpret them, for example when consequences can be expressed in currency.

Precision overload and noisy mapping calorie-label.

Non-Bayesian consumers may find it difficult to understand the meaning of a label message, instead of having a noisy recall of its content. In Section 4, I described precision overload in the noisy recalling process. Consumers observe a label l and with some positive probability they mistakenly recall a different l' . This causes the recalled information to be less reliable and eventually discounted in the evaluation process. I present now a different interpretation of the precision overload phenomenon as noise in the interpretation (instead of the recall)

of the label. The different formulation maintains the main property of tradeoff between simplicity and accuracy of the message conveyed by the label. Consider the different use of the label message for consumers with full and limited attention. Fully attentive consumers know that a label l_k is assigned to products with calorie content $\tau_{k-1} \leq c < \tau_k$, where τ represent the thresholds that create the calorie-label mapping. Consumers with limited attention may have imperfect beliefs about the thresholds used to generate the labels. The noisy agent evaluates the products using a stochastic vector of thresholds with $\hat{\tau}_k \sim F_{\tau_k, \sigma}$. The distribution depends on the correct thresholds τ_k and a noise parameter σ that indicates the uncertainty about the label meaning, with larger noise leading to less accurate beliefs over τ_k . For each set of values $\hat{\tau} = \{\hat{\tau}_k\}_k$ the agent has $\lambda_{l, \hat{\tau}} = E_{\hat{\tau}}[h(c)|l]$. The expected utility for choosing a product is now $E_{\tau}[u_j] = \int_{\tau} h_{\tau}(c) dF_{\tau} + \delta_j + \alpha p_j$. The precision overload setup can be used to determine the accuracy of the thresholds beliefs based on the attention resources available. The accuracy cost $C_{Out}(F_{\tau, \sigma})$ depends on the accuracy of the beliefs over τ , with lower noise σ requiring more resources. One tractable way to define the cost is the following. The accuracy cost is a function of the $C_{Out}(F_{\tau, \sigma}) = \gamma \cdot \log_2(\frac{1}{\sigma})$: additional γ units of attention are required to double the accuracy (inverse of the noise parameter). Also, the beliefs are normally distributed around the true values $\hat{\tau}_k \sim N(\tau_k, \sigma)$ and in this case the noise parameter is simply the standard deviation of the distribution. This example allows to characterize the simplicity-accuracy tradeoff in a similar way as the numerical exercise. A larger number of labels depletes part of the cognitive resources and reduces the accuracy of the beliefs over the label mapping. Bayesian agents respond to lower accuracy: they discount the content of the label and are less responsive to it. If the number of label is sufficiently high with respect to the cognitive budget available, the accuracy effect dominates the information of the label, generating worse choices.

6.2 Endogenous Attention: Information Acquisition and Consideration Sets

Consumers can decide whether to acquire information in addition to the label. In the Bayesian updating model I considered consumers who are uninformed about the calorie content of the product. In the experiment the exact calorie content is available even without labels as the agent can flip the package. The updating problem can be nested in a search process in which the consumer chooses whether to acquire information, and then makes a choice between the products available. Information comes from three sources: basic product characteristics, front-package label, back-package nutrition facts. The consumer always pays attention to the basic characteristics (brand, price), and can learn more by consulting the

other two. I assume that the consumers can pay a cost χ_F to observe the front-of-package labels (if any)⁶² and a different cost χ_B for the back-of-package nutrient labels.⁶³ The search process involves the choice between the information sets $\{I_F, I_B, I_F \cup I_B, \emptyset\}$: one piece of information, both, or none. In the experiment I observe a decline in the collection of back-of-package information when any front-of-package label is introduced, but not between the label treatments. This is suggesting two empirical features of the search process. First, there is endogeneity in the search process: when no label is available, consumers have a larger incentive to consult the nutrient panel. Second, consumers use a “satisficing” search rule: they search if they are sufficiently uninformed but they do not show a significant response to the margin of information of the label. The observation that search probabilities are similar across label treatments suggests that this channel is not crucial to capture the experimental results discussed in this paper. This satisficing search rule suggests that companies may be able to manipulate the search process by providing poorly informative messages. Experimental studies on confusing health claims show that consumers are willing to pay more for oranges that are sold as gluten-free (but clearly all oranges are). If consumers stop searching after receiving some detailed information, this creates an incentive for low-quality products to display true but irrelevant information and inhibit the search process.

Consideration sets offer a different way to think about endogenous attention. Consumers might not consider all the products available. The choice process can be divided into two steps. First, restrict attention to a choice subset. Second, choose within the consideration set. An open question is how the consideration set is formed and its effect on the final choice. Consideration sets models like the ones in Manski (1977) and Abaluck and Adams-Prassl (2021) assume that each product has a probability of being considered. A different approach would consider how consumers build a consideration set by focusing on few salient attributes, and labels could provide a tool to guide the creation of the consideration set. Suppose the consumer operates an initial screening based on the calorie label and excludes all the products that are above a certain cutoff, e.g., remove all the high-calorie products. In this setting, coarse and precise labels can differ in the way they guide the screening process. The precision overload for large choice sets can be interpreted as a factor for the consideration set formation. The consumer can pay a screening cost $C_{\text{Screening}}(N_{\text{Labels}})$ that depends on the complexity of the label message, or randomly select a subset of items for the consideration

⁶²In the Bayesian updating model in Section 4 $\chi_F = 0$ and $\chi_B = \infty$, and the consumers always consults the front-of-package but never the back. In the precision overload setup we can define the front-of-package cost as the recognition cost $\chi_F = C_{In}(F_{i|I})$.

⁶³I assume for simplicity that the agent fully acquires the front or back information after paying the cost instead of paying an individual cost for each product. In the experiment, the search process is highly clustered (once one product is flipped, typically all of the others also are). Sequential search models like Weitzman (1979) characterize the optimal search order and the stopping choice.

set. Coarse labels would require little screening cost, as it is easy to focus on the low-calorie and exclude the high-calorie products. Detailed labels would be more complex and costly to use, and leave fewer attention resources to evaluate each option accurately. Online platforms often use coarse categories to classify products based on price or quality and facilitate the screening process. For instance, Amazon allows to filter based on suggested price range or ratings expressed in stars. This choice tool operates the screening in lieu of the consumer based on the desired criteria.

6.3 Optimal Labeling and Consumer Welfare Evaluation

The policymakers can consider what is the best labeling regime in terms of coarsening of information. This problem belongs to a well-studied family of principal-agent problems with optimal information disclosure. When principal and agents' objectives are aligned, full information is optimal (Blackwell, 1953; Kamenica and Gentzkow, 2011) but this is not true anymore if incentives are misaligned or if agents are boundedly rational. If the principal want to persuade the agent to choose some options, then they have incentive not to provide full information. Coarsening naturally emerge from the strategic interaction and it harms the consumer. Rayo and Segal (2010) show that principal's utility is typically maximized by using partial information disclosure. In a model of choice over prospect, sellers pool together appealing *baits* and unappealing *switches* to convince the consumer to accept more offers than they would under full information. The regulator who wants to promote healthy food choices can apply a similar mechanism. By pooling together very- and mildly-unhealthy products they can persuade the consumers to stop consuming the latter category when they would normally just avoid the first one.

A benevolent regulator can also implement coarse labels when the consumers are boundedly rational. Limited attention and myopic intertemporal preferences are two cases in which the information design can mitigate the consumers' mistakes. In this paper I consider the case of precision overload as a source of mistakes. When the information is too detailed, consumers may be confused and make more mistakes. I showed in Section 5 that the quality of choices can improve in some cases by reducing the number of labels.

6.4 Companies' Response to Labeling Regime Changes

Manufacturers can respond to the introduction of a label by changing the characteristics of their products (Akerberg and Crawford, 2009; Wollmann, 2018). In the short run, they can change the prices. In the long run, they can introduce new products or change the ingredients of the old ones. Barahona et al. (2020) show that both these responses occurred in Chile

after the introduction of warning labels, and as a result this increased the effectiveness of the labeling reform. In particular, manufacturers responded by bunching the calorie content below the threshold to avoid receiving the infamous warning label. Introducing a binary label incentivizes companies just above the threshold to change the recipe as long as the quality gain for being awarded a better label is larger than the production cost for reaching the requirements. Moving to a more detailed label has an ambiguous effect on the firms' incentives to reformulate. Suppose that the regulator introduces a healthy-label for product below 100 calories in addition to a warning-label for products above 200 calories. On the one hand, it reduces the gain for high-calorie products (e.g., 210 calories) to move to the adjacent partition cell, because beliefs are more accurate. Consumers would know that the product is in the 100-200 calorie range instead of the 0-200 range. On the other hand, it can incentivize reformulation for low-calorie products, that can move from e.g., 110 calories to less to be awarded the new label. A formal way to include the companies' response in the model is to consider an endogenous choice set of products. A labeling regime $L = (\mathcal{L}, \tau)$ leads to a choice set \mathcal{J}_L that, in equilibrium, is the result of the strategic response of the manufacturers.⁶⁴ Consumers' demand depends on both the products available in the choice set \mathcal{J}_L and the information provided through the labels' messages $l \in \mathcal{L}$. The regulator would evaluate the response to the labeling regime through both the supply and demand channels.

How would precision overload affect the manufacturers' response? In a model with inattentive consumers, companies would face a different tradeoff between the reformulation cost and the benefit of displaying a more desirable label. Compared with the full attention benchmark, companies face the same production cost to change the recipe, but inattentive consumers may discount the quality improvement when labels are more detailed. This diminishes the incentive to reformulate, since part of the consumers might not recognize the product change.

7 Related Literature

This paper contributes to three separate streams of literature on food labels, consumer inattention, and coarse information.

Food labels play an important role in informing consumers and driving food choices. A

⁶⁴Under the assumption that companies do not retire or introduce products, but only vary its recipe and price, the vector of products' characteristic can be written as (δ_j, p_j^L, c_j^L) where the reformulation can change the calorie content c across regimes but not its taste δ_j . This is the approach used by Barahona et al. (2020) to model the optimal reformulation problem for manufacturers after the introduction of a warning label.

vast and interdisciplinary literature has been studying the effect of food labels on consumer behavior using a combination of laboratory (Bix et al., 2015; Ono and Ono, 2015; Crosetto et al., 2020) and field studies (Dubois et al., 2020; Shangguan et al., 2019). Experimental evidence and surveys indicate that front-of-package labels can help consumers to recognize which products are healthier⁶⁵ (Ikonen et al., 2020; Drichoutis, Lazaridis and Nayga Jr, 2006). There is no clear evidence on what is the effect of more granular information on food choices, and this is the primary question I am studying in this paper. The closest paper in the literature addressing a similar question is Crosetto, Muller and Ruffieux (2016). The authors use a 3×2 design to vary label types and time pressure. For time pressure, they ask to complete the task with and without time constraint. For the type of labels, they use coarse colors (traffic lights to indicate high, medium, or low content), precise numbers (number of calories, grams of sugar, etc.) or a combination of the two (both number and traffic light color). They assign the task to assemble a dietary plan that respects specific nutritional objectives and report that detailed labels lead to higher accuracy in the task under free time, but results are not significantly different under limited time. Differently from their design, I use coarse-categorical and detailed-numerical labels but I maintain the same color palette to convey the message, I focus on one single nutrient (calories), and I analyze the effect on consumer choices (instead of a task with an objective correct answer).

The 2016 Chilean nutrient label reform allowed to study the effect of mandatory warning messages using secondary data (Barahona et al., 2020; Pachali et al., 2020). The empirical investigation shows a significant effect on consumer behavior through a combination of direct effect of the label and indirect effects through manufacturers’ response, namely price change and ingredients reformulation. Chilean warning labels are binary; for each nutrient, a product either displays the warning or not. In principle, it is not clear whether the effects of the reform and their impact on consumer welfare would be larger by providing more detailed labels.⁶⁶ My experimental results suggest this might not be the case, and that low-SES groups might be the ones that benefit the most from coarse information.

A separate concern comes from the evidence that food labels could affect consumer behavior even when they provide no real content or confuse consumers. Wilson and Lusk (2020) show that consumers respond to meaningless nutritional messages, e.g. expressing higher willingness to pay for gluten-free orange juice despite the fact that all orange juices

⁶⁵The definition of healthy product varies across papers. It includes product with higher/lower content in one desirable/undesirable nutrient (e.g., higher in fiber or lower in sugar) or higher score according to one particular scale (e.g., NutriScore).

⁶⁶Policymakers and information designer face multidimensional problems about what information to present, and how. Field experiments show that the effect of information can be ambiguous on behavior, as shown for example by Ferraro and Price (2013) (reduce water consumption) and Benneer et al. (2013) (avoid high-arsenic wells).

are gluten-free. This kind of evidence points out that labels' role might be a combination of information and persuasion, and that consumers might be confused about their real meaning. The precision overload model provides a general setup that is able to capture this effect: consumers might have imperfect ability to understand the message on the label and misuse confusing information.

Another relevant stream of literature studies the role of inattention in consumer behavior. When the goal of the regulator is to inform consumers, making the information available might not be sufficient if consumers are not attentive. We have experimental and observational evidence that consumers do not use all the information available and make costly mistakes. For example, Chetty, Looney and Kroft (2009) show that consumers underreact to taxes that are not salient. Evidence of inattention in store purchases also relate to base prices and promotions (Dickson and Sawyer, 1990), quantity surcharge (Clerides and Courty, 2017), and expiration dates (Hansen, Misra and Sanders, 2021). In my study the front-of-package label provide a coarser version of the calorie amount available on the back-of-package. If consumers were attentive and motivated to acquire that piece of information, we would not expect any response to the experimental manipulation. The results show that inattention might be enhanced by excess *precision* of the message. I define precision overload as the state in which precision inhibits the decision maker's ability to optimally determine the best possible decision: excessive precision can be confusing and lead to worse choices. In the experiment I do not have an objective definition of correct choices, as preferences are subjective, but I evaluate the difference across information treatment based on the policy-maker's perspective (who want to reduce calories), the consumer perspective (based on the declared preferences over calories), and the general preference over types of labels (whether consumers prefer coarse or detailed labels, or none). Precision overload is analogous to other effects related to limited ability to process complex problems, such as choice overload (large number of options to choose from, Iyengar and Lepper 2000; Brown and Jeon 2019) and information overload (large number of attributes to consider to evaluate each option, Lee and Lee 2004; Ursu 2018; Roetzel 2019).

The paper is motivated by the remark that coarse labels are widely used both by private companies that are profit-maximizer and public authorities that aim to protect and inform consumers. The theoretical literature has discussed various reasons according to which coarse information could be desirable when consumers are fully able to process the information (e.g., cheap talk models) or face a limited ability to update that depends on the actual use of information in the decision process (e.g., rational inattention models). The novelty in the model lies in considering how the complexity of the information can negatively affect the agents' ability to use it. Coarsening commonly emerges in principal-agent settings in which

incentives are not aligned and sellers (signal senders) want to inflate the perceived quality of the product from the perspective of the consumers (receivers). This is the case in cheap talk models (Crawford and Sobel, 1982; Chakraborty and Harbaugh, 2010) and settings with verifiable disclosure (Milgrom, 2008; Rayo and Segal, 2010) in which pooling of multiple states is optimal. The idea of withholding information to benefit the sender is also discussed in the Bayesian persuasion literature (Kamenica and Gentzkow, 2011). Bloedel and Segal (2018) consider a persuasion model in which receivers are rationally inattentive and senders find it optimal to provide coarse information even when the goal misalignment is arbitrarily small. Lipnowski, Mathevet and Wei (2020) introduce an attention management model in which a benevolent principal can choose to withhold information to make the rationally inattentive agent focus on fewer issues. The agents would still prefer receiving detailed information, but higher welfare can be achieved under coarse signals. Other examples of coarsening emerge from voluntary certification (Lizzeri, 1999; Harbaugh and Rasmusen, 2018); even a benevolent authority might introduce a coarse and voluntary certification in order to incentivize sellers to opt in.

The precision overload model provides a new explanation for the emergence of coarse labels even outside strategic settings. I consider consumers whose cognitive resources can be depleted because of the complexity (precision) of the information, regardless of its subsequent use. This means that agents do not have free information disposal; in particular they cannot coarsen information for free. If they could, having more granular information would always be desirable as it can give the sender the freedom to choose between detailed and coarse signals.⁶⁷

The removal of the free information disposal assumption leads to a violation of Blackwell's theorem (Blackwell, 1953; Crémer, 1982): a detailed label, that is statistically sufficient to represent the signal of the coarse label, might be less informative than the coarse label itself. The detailed signal is statistically sufficient because it generates the coarse signal by pooling together messages that belong to the same partition cell. Blackwell information criterion instead relies on the expected utility from implementing the optimal action conditional on the signal. If the detailed message is associated to larger noise this can generate costly mistakes and reduce the expected utility. This violation can explain preference for coarse labels, that is one of the results I find in the experiment. Consumers might prefer coarse information because they are indeed able to make better choices with respect to what they themselves

⁶⁷Under free information disposal the agent that receives a signal σ about the state of the world s (calorie amount) would always be able to replace it with a weakly less informative (in the Blackwell order sense) signal σ' that is a garbling of σ . Given a (relatively) detailed labeling regime - defined according to the vector of threshold τ in the notation introduced in Section 5 - the label signals are deterministic and indicate the range of feasible calorie amounts.

want. This does not exclude that in the real world many of these effects occur at once. The adoption of coarse labels can be due to a combination of demand for simplicity and attempts to guide consumer choices towards desirable behavior, for example healthier diets. Further motivations include the strategic dimension of the problem, with manufacturers responding with price changes (Pachali et al., 2020) or reformulation of the product characteristics (Barahona et al., 2020).

8 Conclusions

In this paper, I study the role of coarse and precise information in food labeling. The project is motivated by the question: why do companies and regulators adopt coarse labels instead of more detailed ones? According to the Bayesian inference model, more granular labels would be more informative and would help consumers to make better choices, given their preferences. But if detailed labels are confusing, consumers with limited attention might prefer coarse labels that are simpler to understand. I explore this question in a food choice setting by conducting an experiment in which I manipulate labels with calorie information.

Coarse-categorical labels are effective in promoting the choice of products with fewer calories (-3%), and more detailed numerical labels are less effective (only -1%). Also, 61% of the subjects declare they prefer coarse labels with an intermediate level of precision (2 or 4 categories). The choice data suggest that subjects are more sensitive to coarse information and are learning less from detailed labels. Consistent with this conjecture, the estimates of the label effect in a simple demand model highlight that coarse labels have systematically larger weights in the decision process. I introduce a limited attention model in which agents have limited cognitive resources consumed by two types of attention cost: a retrieval cost that depends on the accuracy of the information used to make a choice, and a recognition cost that depends on the complexity of the message. This model can capture the main results from the experiment and characterizes the precision overload effect. Precise labels require more cognitive resources and reduce the ability to extract information from the label message.

I acknowledge the limitations of the experiment and the model, and I hope these results will encourage further exploration of the role of product labels – and more broadly signals in information models - in new directions. I identify three major limitations of the experimental results and indicate how the experiment can be modified to address them. First, the experiment compares the effect of coarse-categorical and detailed-numerical labels. Further studies are required to identify if the different formats of the label drive the result. High precision is typically associated with numerical information in the real world, and this alone

may represent a source of confusion for consumers. Second, the experiment involves small choice sets with similar products, for example two cereals. It is unclear a priori if the effects would be larger with larger choice sets. Consideration sets models assume that consumers with limited attention focus on a subset of the options available, and coarse labels might guide the screening phase that determines the consideration set. This question is suitable for experiments that use process data, such as mouse-tracking or eye-tracking. Third, I do not directly observe the subjective values of each option. Ordinal preferences over calories provide a robustness measure for monotonicity but do not allow to measure welfare gains. The setting can be extended to capture cardinal preferences, for example by using different tasks (e.g., delegated choice tasks with explicit goals) or stimuli (e.g., abstract options with a clear value in a simplified choice environment).

The model of precision overload provides a simple extension of the Bayesian inference model with endogenous noise. Its purpose in this paper is to show how it can capture the main stylized results of the experiment, but other questions remain open for future research. First, explore further if other attention models are compatible with the stylized results and can lead to lower accuracy when under more precise information. Second, establish what are the necessary and sufficient conditions required to characterize the precision overload effect in a dataset with stochastic choice.

Finally, further work is required to explore the welfare implications of labeling reforms and compare with other policy interventions. The experimental results suggest that the welfare analysis of the information reform might depend on the response of consumers with limited — and possibly heterogeneous — attention. Policymakers can use various tools to promote healthier food choices, including banning, taxing, or labeling products deemed harmful for most consumers. Standard economics models can benefit from the introduction of bounded rationality assumptions on consumer behavior, with implications on which tools to use to help consumers. If consumers are myopic and underestimate future negative consequences, this can motivate the introduction of sugar taxes. Instead, if consumers are simply uninformed about the health consequences, the policymakers might want to promote information provision. Standard information models would prescribe the adoption of the most precise labels, but the experiment results suggest that excessive precision might reduce the choice of low-calorie products, especially for consumers with low socioeconomic status. Designing simple and effective labels that inform consumers and help them choose what they want will require further investigation to characterize the tradeoff between simplicity and precision.

References

- Abaluck, Jason.** 2011. “What would we eat if we knew more: The implications of a large-scale change in nutrition labeling.” *Massachusetts Institute of Technology Working Paper*.
- Abaluck, Jason, and Abi Adams-Prassl.** 2021. “What do consumers consider before they choose? Identification from asymmetric demand responses.” *The Quarterly Journal of Economics*, 136(3): 1611–1663.
- Ackerberg, Daniel A, and Gregory S Crawford.** 2009. “Estimating Price Elasticities in Differentiated Product Models with Endogenous Characteristics.” Mimeo, University of Warwick.
- Alé-Chilet, Jorge, and Sarah Moshary.** 2020. “Beyond Consumer Switching: Supply Responses to Food Packaging and Advertising Regulations.” *Available at SSRN 3678744*.
- Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky.** 2019. “Should we tax sugar-sweetened beverages? An overview of theory and evidence.” *Journal of Economic Perspectives*, 33(3): 202–27.
- Ambuehl, Sandro, and Shengwu Li.** 2018. “Belief updating and the demand for information.” *Games and Economic Behavior*, 109: 21–39.
- Andor, Mark A, Andreas Gerster, and Lorenz Götte.** 2019. “How effective is the European Union energy label? Evidence from a real-stakes experiment.” *Environmental Research Letters*, 14(4): 044001.
- Anesbury, Zachary, Magda Nenycz-Thiel, John Dawes, and Rachel Kennedy.** 2016. “How do shoppers behave online? An observational study of online grocery shopping.” *Journal of Consumer Behaviour*, 15(3): 261–270.
- Balasubramanian, Siva K, and Catherine Cole.** 2002. “Consumers’ search and use of nutrition information: The challenge and promise of the nutrition labeling and education act.” *Journal of marketing*, 66(3): 112–127.
- Barahona, Nano, Cristóbal Otero, Sebastián Otero, and Joshua Kim.** 2020. “Equilibrium Effects of Food Labeling Policies.” *Available at SSRN 3698473*.
- Bennear, Lori, Alessandro Tarozzi, Alexander Pfaff, Soumya Balasubramanya, Kazi Matin Ahmed, and Alexander Van Geen.** 2013. “Impact of a randomized controlled trial in arsenic risk communication on household water-source choices in Bangladesh.” *Journal of environmental economics and management*, 65(2): 225–240.
- Bix, Laura, Raghav Prashant Sundar, Nora M Bello, Chad Peltier, Lorraine J Weatherspoon, and Mark W Becker.** 2015. “To see or not to see: Do front of pack nutrition labels affect attention to overall nutrition information?” *PloS one*, 10(10): e0139732.

- Blackwell, David.** 1953. "Equivalent comparisons of experiments." *The annals of mathematical statistics*, 265–272.
- Bloedel, Alexander W, and Ilya R Segal.** 2018. "Persuasion with rational inattention." *Available at SSRN 3164033*.
- Boon, Caitlin S, Alice H Lichtenstein, Ellen A Wartella, et al.** 2010. *Front-of-package nutrition rating systems and symbols: Phase I report*. National Academies Press.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2012. "Salience theory of choice under risk." *The Quarterly journal of economics*, 127(3): 1243–1285.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2013. "Salience and consumer choice." *Journal of Political Economy*, 121(5): 803–843.
- Brown, Zach Y, and Jihye Jeon.** 2019. "Endogenous Information Acquisition and Insurance Choice." *University of Michigan and Boston University Working Paper*.
- Burnham, Kenneth P, and David R Anderson.** 2002. "A practical information-theoretic approach." *Model selection and multimodel inference*, 2.
- Caplin, Andrew, and Mark Dean.** 2015. "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review*, 105(7): 2183–2203.
- Chakraborty, Archishman, and Rick Harbaugh.** 2010. "Persuasion by cheap talk." *American Economic Review*, 100(5): 2361–82.
- Chandler, Jesse, Cheskie Rosenzweig, Aaron J Moss, Jonathan Robinson, and Leib Litman.** 2019. "Online panels in social science research: Expanding sampling methods beyond Mechanical Turk." *Behavior research methods*, 51(5): 2022–2038.
- Chetty, Raj, Adam Looney, and Kory Kroft.** 2009. "Salience and taxation: Theory and evidence." *American economic review*, 99(4): 1145–77.
- Clerides, Sofronis, and Pascal Courty.** 2017. "Sales, quantity surcharge, and consumer inattention." *Review of Economics and Statistics*, 99(2): 357–370.
- Cowan, Nelson.** 2001. "The magical number 4 in short-term memory: A reconsideration of mental storage capacity." *Behavioral and brain sciences*, 24(1): 87–114.
- Crawford, Vincent P, and Joel Sobel.** 1982. "Strategic information transmission." *Econometrica: Journal of the Econometric Society*, 1431–1451.
- Cr  mer, Jacques.** 1982. "A simple proof of Blackwell's "comparison of experiments" theorem." *Journal of Economic Theory*, 27(2): 439–443.
- Crosetto, Paolo, Anne Lacroix, Laurent Muller, and Bernard Ruffieux.** 2020. "Nutritional and economic impact of five alternative front-of-pack nutritional labels: experimental evidence." *European Review of Agricultural Economics*, 47(2): 785–818.

- Crosetto, Paolo, Laurent Muller, and Bernard Ruffieux.** 2016. "Helping consumers with a front-of-pack label: Numbers or colors?: Experimental comparison between Guideline Daily Amount and Traffic Light in a diet-building exercise." *Journal of Economic Psychology*, 55: 30–50.
- Davis, Lucas W, and Gilbert E Metcalf.** 2016. "Does better information lead to better choices? Evidence from energy-efficiency labels." *Journal of the Association of Environmental and Resource Economists*, 3(3): 589–625.
- Derby, Brenda M, and Alan S Levy.** 2001. "Do food labels work? Gauging the effectiveness of food labels pre-and post-NLEA."
- Dickson, Peter R, and Alan G Sawyer.** 1990. "The price knowledge and search of supermarket shoppers." *Journal of marketing*, 54(3): 42–53.
- Drichoutis, Andreas C, Panagiotis Lazaridis, and Rodolfo M Nayga Jr.** 2006. "Consumers' use of nutritional labels: a review of research studies and issues." *Academy of marketing science review*, 2006: 1.
- Dubois, Pierre, Paulo Albuquerque, Olivier Allais, Céline Bonnet, Patrice Bertail, Pierre Combris, Saadi Lahlou, Natalie Rigal, Bernard Ruffieux, and Pierre Chandon.** 2020. "Effects of front-of-pack labels on the nutritional quality of supermarket food purchases: evidence from a large-scale randomized controlled trial." *Journal of the Academy of Marketing Science*, 1–20.
- Ferraro, Paul J, and Michael K Price.** 2013. "Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment." *Review of Economics and Statistics*, 95(1): 64–73.
- Golman, Russell, David Hagmann, and George Loewenstein.** 2017. "Information avoidance." *Journal of Economic Literature*, 55(1): 96–135.
- Hainmueller, Jens, Michael J Hiscox, and Sandra Sequeira.** 2015. "Consumer demand for fair trade: Evidence from a multistore field experiment." *Review of Economics and Statistics*, 97(2): 242–256.
- Hamrick, Karen S.** 2016. "Do Americans Eat Meals Anymore or Do They Just Snack?" In *The economics of multitasking*. 109–143. Springer.
- Hansen, Jochim, and Michaela Wänke.** 2009. "Liking what's familiar: The importance of unconscious familiarity in the mere-exposure effect." *Social cognition*, 27(2): 161–182.
- Hansen, Karsten, Kanishka Misra, and Robert E Sanders.** 2021. "Consumer (In) attention to Expiration Dates: A Field Study." *Available at SSRN*.
- Harbaugh, Rick, and Eric Rasmusen.** 2018. "Coarse grades: Informing the public by withholding information." *American Economic Journal: Microeconomics*, 10(1): 210–35.

- Hoadley, James E, and J Craig Rowlands.** 2014. “FDA perspectives on food label claims in the United States.” In *Nutraceutical and functional food regulations in the United States and around the world*. 121–140. Elsevier.
- Hoyer, Wayne D.** 1984. “An examination of consumer decision making for a common repeat purchase product.” *Journal of consumer research*, 11(3): 822–829.
- Ikonen, Iina, Francesca Sotgiu, Aylin Aydinli, and Peeter WJ Verlegh.** 2020. “Consumer effects of front-of-package nutrition labeling: An interdisciplinary meta-analysis.” *Journal of the Academy of Marketing Science*, 48(3): 360–383.
- Iyengar, Sheena S, and Mark R Lepper.** 2000. “When choice is demotivating: Can one desire too much of a good thing?” *Journal of personality and social psychology*, 79(6): 995.
- Julia, Chantal, Fabrice Etilé, Serge Hercberg, et al.** 2018. “Front-of-pack Nutri-Score labelling in France: an evidence-based policy.” *Lancet Public Health*, 3(4): e164.
- Kamenica, Emir.** 2008. “Contextual inference in markets: On the informational content of product lines.” *American Economic Review*, 98(5): 2127–49.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. “Bayesian persuasion.” *American Economic Review*, 101(6): 2590–2615.
- Kuksov, Dmitri, and J Miguel Villas-Boas.** 2010. “When more alternatives lead to less choice.” *Marketing Science*, 29(3): 507–524.
- Lee, Byung-Kwan, and Wei-Na Lee.** 2004. “The effect of information overload on consumer choice quality in an on-line environment.” *Psychology & Marketing*, 21(3): 159–183.
- Lipnowski, Elliot, Laurent Mathevet, and Dong Wei.** 2020. “Attention management.” *American Economic Review: Insights*, 2(1): 17–32.
- Litman, Leib, Zohn Rosen, Cheskie Ronsezweig, Sarah L Weinberger, Aaron J Moss, and Jonathan Robinson.** 2020. “Did people really drink bleach to prevent COVID-19? A tale of problematic respondents and a guide for measuring rare events in survey data.” *MedRxiv*.
- Lizzeri, Alessandro.** 1999. “Information revelation and certification intermediaries.” *The RAND Journal of Economics*, 214–231.
- Loewenstein, George.** 1987. “Anticipation and the valuation of delayed consumption.” *The Economic Journal*, 97(387): 666–684.
- Mannell, Ashley, Patricia Brevard, Rodolfo Nayga, Pierre Combris, Robert Lee, and Janet Gloeckner.** 2006. “French consumers’ use of nutrition labels.” *Nutrition & Food Science*.
- Manski, Charles F.** 1977. “The structure of random utility models.” *Theory and decision*, 8(3): 229.

- Mayer, Robert N, Debra L Scammon, and Linda F Golodner.** 1993. "Healthy confusion for consumers." *Journal of Public Policy & Marketing*, 12(1): 130–132.
- Milgrom, Paul.** 2008. "What the seller won't tell you: Persuasion and disclosure in markets." *Journal of Economic Perspectives*, 22(2): 115–131.
- Miller, George A.** 1956. "The magical number seven, plus or minus two: Some limits on our capacity for processing information." *Psychological review*, 63(2): 81.
- Moorman, Christine.** 1996. "A quasi experiment to assess the consumer and informational determinants of nutrition information processing activities: The case of the nutrition labeling and education act." *Journal of Public Policy & Marketing*, 15(1): 28–44.
- Mrkva, Kellen, Nathaniel A Posner, Crystal Reeck, and Eric J Johnson.** 2021. "Do nudges reduce disparities? Choice architecture compensates for low consumer knowledge." *Journal of Marketing*, 0022242921993186.
- Norwich, Kenneth H.** 1993. *Information, sensation, and perception*. Academic Press San Diego.
- Novák, Vladimír, Andrei Matveenko, and Silvio Ravaioli.** 2021. "The Status Quo and Belief Polarization of Inattentive Agents: Theory and Experiment."
- Ono, Makoto, and Akinori Ono.** 2015. "Impacts of the FoSHU (Food for Specified Health Uses) system on food evaluations in Japan." *Journal of Consumer Marketing*.
- Pachali, Max J, Marco JW Kotschedoff, Arjen van Lin, Bart J Bronnenberg, and Erica van Herpen.** 2020. "Do warning labels make products more differentiated and lead to higher prices?" *Available at SSRN 3642756*.
- Paninski, Liam, Jonathan Pillow, and Jeremy Lewi.** 2007. "Statistical models for neural encoding, decoding, and optimal stimulus design." *Progress in brain research*, 165: 493–507.
- Poelman, Maartje P, S Coosje Dijkstra, Hanne Sponselee, Carlijn BM Kamphuis, Marieke CE Battjes-Fries, Marleen Gillebaart, and Jacob C Seidell.** 2018. "Towards the measurement of food literacy with respect to healthy eating: the development and validation of the self perceived food literacy scale among an adult sample in the Netherlands." *International Journal of Behavioral Nutrition and Physical Activity*, 15(1): 54.
- Polania, Rafael, Michael Woodford, and Christian C Ruff.** 2019. "Efficient coding of subjective value." *Nature neuroscience*, 22(1): 134–142.
- Pollack, Irwin.** 1952. "The information of elementary auditory displays." *The Journal of the Acoustical Society of America*, 24(6): 745–749.
- Rayo, Luis, and Ilya Segal.** 2010. "Optimal information disclosure." *Journal of political Economy*, 118(5): 949–987.

- Reisch, Lucia A, Cass R Sunstein, and Micha Kaiser.** 2021. "What do people want to know? Information avoidance and food policy implications." *Food Policy*, 102: 102076.
- Rhodes, Donna G, Meghan E Adler, John C Clemens, and Alanna J Moshfegh.** 2017. "What we eat in America food categories and changes between survey cycles." *Journal of Food Composition and Analysis*, 64: 107–111.
- Roetzel, Peter Gordon.** 2019. "Information overload in the information age: a review of the literature from business administration, business psychology, and related disciplines with a bibliometric approach and framework development." *Business research*, 12(2): 479–522.
- Rousseau, Sandra.** 2015. "The role of organic and fair trade labels when choosing chocolate." *Food Quality and Preference*, 44: 92–100.
- Sandoval, Luis A, Carlos E Carpio, and Marcos Sanchez-Plata.** 2019. "The effect of 'Traffic-Light' nutritional labelling in carbonated soft drink purchases in Ecuador." *PloS one*, 14(10): e0222866.
- Schuldt, Jonathon P.** 2013. "Does green mean healthy? Nutrition label color affects perceptions of healthfulness." *Health communication*, 28(8): 814–821.
- Shangguan, Siyi, Ashkan Afshin, Masha Shulkin, Wenjie Ma, Daniel Marsden, Jessica Smith, Michael Saheb-Kashaf, Peilin Shi, Renata Micha, Fumiaki Ima-mura, et al.** 2019. "A meta-analysis of food labeling effects on consumer diet behaviors and industry practices." *American journal of preventive medicine*, 56(2): 300–314.
- Shannon, Claude Elwood.** 1948. "A mathematical theory of communication." *The Bell system technical journal*, 27(3): 379–423.
- Shen, Meng, Lijia Shi, and Zhifeng Gao.** 2018. "Beyond the food label itself: How does color affect attention to information on food labels and preference for food attributes?" *Food Quality and Preference*, 64: 47–55.
- Sims, Christopher A.** 2003. "Implications of rational inattention." *Journal of monetary Economics*, 50(3): 665–690.
- Ursu, Raluca M.** 2018. "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions." *Marketing Science*, 37(4): 530–552.
- Wartella, Ellen A, Alice H Lichtenstein, Ann Yaktine, and Romy Nathan.** 2012. "Committee on Examination of Front-of-Package Nutrition Rating Systems and Symbols (Phase II)." *Washington: Institute of Medicine*.
- Weitzman, Martin L.** 1979. "Optimal search for the best alternative." *Econometrica: Journal of the Econometric Society*, 641–654.
- Wilson, Lacey, and Jayson L Lusk.** 2020. "Consumer willingness to pay for redundant food labels." *Food Policy*, 97: 101938.

Wollmann, Thomas G. 2018. “Trucks without bailouts: Equilibrium product characteristics for commercial vehicles.” *American Economic Review*, 108(6): 1364–1406.

A Appendix: Bayesian updating model

A.1 Bayesian agents and U-shaped response of label precision

I provide an example of how using more precise labels can generate a U-shaped response even with strictly monotonic preferences. In order to observe this effect we need to observe a case in which the consumer chooses the (relatively) low-calorie option under coarse label and switches to the high-calorie option under detailed label regime.

Consider a fully rational consumer who always prefers product with low calories and strongly dislikes very high calories, above \tilde{c} 220 calories. Their simple utility function is taking into account only calories and price expressed in cents:

$$u = h(c) - p = 1000 - c - 1[c > \tilde{c}] \cdot 10 \cdot (c - \tilde{c}) - p$$

Take a choice set with an expensive low-calorie product (j , 120 calories, \$4) and a cheaper high-calorie product (k , 200 calories, \$3). Under the coarse labeling regime the expected utility for choosing each product is $E[u_j] = 1000 - 130 - 400 = 470$ and $E[u_k] = 1000 - 287 - 300 = 413$, and product j is selected. This is because the high-calorie label is pooling together products above and below \tilde{c} . Under the detailed labeling regime we have $E[u_j] = 1000 - 120 - 400 = 480$ and $E[u_k] = 1000 - 200 - 300 = 500$, and product k is selected. The same effect can emerge for any arbitrary value of \tilde{c} that represents a discontinuity in a piece-wise linear utility, and more broadly if we have nonlinear effect of calories on the utility.

A.2 Alternative setup: Labels as signals of health-utility

Consider a setup similar to the one presented in Section 4, with the difference that consumers do not know the mapping $h(c)$ and learn about the health consequences of a product directly from the label message.

This is consistent with the idea that consumers do not know how to use the numerical information about the calories: they do not know how healthy is a snack with 100 calories instead of 200. But they understand that high-calorie is worse than low-calorie by a certain factor, that can differ across consumers.

I assume that the utility derived by individual i when purchasing product j can be

separate into three components:

$$u_{ijt} = \underbrace{\beta_i \cdot h_j}_{\text{health consequences}} + \underbrace{\delta_{ij}}_{\text{experience/taste}} + \underbrace{\alpha_i p_{jt}}_{\text{price}}$$

Products have some health consequences that the regulator can observe and quantify with a numerical value h_j (possibly a vector of values) and consumers can have different preferences about how much they value health with respect to taste and price. In this setting labels can be easily interpreted as recommendations for consumers. Consumers with preference coefficient β should renounce to Δ dollars in order to select a product that, on expectations, has $\frac{\Delta \cdot \beta}{\alpha}$ more units of health. The health consequences can be inferred from the label

$$\mathbb{E}_h[u_{ijt}] = \beta \cdot \mathbb{E}_h[h_j | l_{jt}] + \delta_{ij} + \alpha_i p_{jt}$$

where the consumer has a prior belief about h_j and uses Bayes' rule to update own beliefs after observing l_{jt} . This naturally extends to the special case of homogeneous preferences $\beta_i = \beta$ for all the consumers, and the tradeoff health-taste instead (or jointly) of health-price. Similar to the case of inference I want to show that a more detailed labeling regime is more informative and improves the consumers' expected utility. This can be easily shown using Jensen's inequality: the coarsening of the detailed labels cannot benefit, and possibly harms, the selection of the expected utility-maximizing product.

Consider two product $j \in 1, 2$, and compare the labeling regime C (relatively coarse) and its partition D (detailed). The agent can observe p and δ , and in the detailed labels they form expected utilities $v_l = E_h[u|l]$, for simplicity of notation. The expected value under the coarse regime can be expressed as convex combination of values in the detailed regime as discussed in Section 4.

$$v_{j|l'_c} = \sum_{l_d \in \mathcal{L}_{d|l'_c}} Pr(l_d | l'_c) \cdot v_{j|l_d}$$

Under coarse labeling the agent maximizes the expected utility according to the label displayed for each product, and health consequences cannot be separated among products with the same label

$$EU_c^* = \max_j \{ \sum_{l_d \in \mathcal{L}_{d|l_j c}} Pr(l_d | l_j c) \cdot v_{j|l_d} \}_j$$

whereas the differences can appear under detailed labels, and the maximization uses the granular information

$$EU_d^* = \sum_{l_d \in \mathcal{L}_{d|l_1 c}} \sum_{l'_d \in \mathcal{L}_{d|l_2 c}} Pr(l_d, l'_d | l_1 c, l_2 c) \cdot \max_j \{ \sum_{l_d \in \mathcal{L}_{d|l'_c}} Pr(l_d | l'_c) \cdot v_{j|l_d} \}_j$$

and by Jensens's inequality $EU_d^* = \mathbb{E}[\max_j(v_j)|l_d] \geq \max_j(\mathbb{E}[v_j|l_d]) = EU_c^*$.

A.3 Alternative setup: Labels as signals of health-calorie mapping

Consider a setup similar to the one presented in Section 4, with the difference that consumers do not know the mapping $h(c)$ and learn about the health consequences of a product indirectly from the label message. As in Kamenica (2008) uninformed consumers do not know a parameter β that determines the mapping of a product characteristic (calorie, or quality) into their utility. Kamenica (2008) calls it the global preference parameter and allows consumers to have individual preferences that depend on two components (a global and an individual parameter). Consumers have a set of utility functions $h_\beta(c)$, with true utility $h_{\beta^*}(c)$ based on the true β^* . The consumer has a prior belief distribution $\beta \sim F_\beta$ and updates own beliefs based on the labeling regime chosen by the regulator.

This is consistent with the idea that consumers do not know how to use the numerical information about the calories: they do not know how healthy is a snack with 100 calories instead of 200. But they understand that the regulator decided to separate high-calorie and low-calorie products according to their knowledge of the consumers' preferences.

The benevolent regulator knows F_β and β^* and chooses a labeling regime based on its value. A labeling regime is characterized by the vector τ of thresholds that determine the calorie-label mapping. If the consumers were able to observe τ they would immediately infer the true β^* . Instead, they jointly observe products' calories c and labels l . This allows to infer τ and as a consequence refine the beliefs over β .

The health consequences component of the utility functions is $h_\beta(c_j)$. I assume that the consumer has exact beliefs about c_j . In this setting labels can be easily interpreted as educational tools to learn how to interpret the labels. This indirectly implies that the consumer also learn about the product that displays the label.

The expected health utility after observing a product j with label l_j is $\mathbb{E}_\beta[h_\beta(c_j)|l_j]$. Given a set of thresholds $\tau = \{\tau_1, \tau_2, \dots, \tau_N\}$ that determine the labeling regime, adding an additional partition that leads to $\tau' = \{\tau_1, \tau_2, \dots, \tau_N, \tau_{N+1}\}$ weakly increases the accuracy of the beliefs over β . The reason is that the consumer either observes the same signal as before (with no impact on the beliefs) or observes one out of two separate signals that lead to more precise posterior beliefs.

B Appendix: Limited attention models

B.1 Saliency model and sensitivity to information

In Section 5 I discuss why a saliency model a' la Bordalo, Gennaioli and Shleifer (2012) is not able to capture the experiment result. One of the experimental results is sensitivity to information. I observe that subjects appear less sensitive to the calorie content in the detailed treatment compared with the coarse treatment. This is at odds with what the saliency model predicts. The model assumes that the consumer can have a true utility function (without saliency distortion)

$$u_j = q_{1j} + q_{2j}$$

based on the value of two quality attributes for each product j . The consumer makes a choice using the distorted utility function that includes saliency

$$u_j = \delta_{1j} \cdot q_{1j} + \delta_{2j} \cdot q_{2j}$$

where the δ_{aj} weights attached to each attribute a depends on the level of the attribute for the product j compared to the other products in the choice set. The weights depend on a saliency function $\sigma(q_{aj}, \bar{q}_a)$ that depends on the value of the attribute for the product compared to the rest of the distribution (average \bar{q}_a). This function satisfies conditions that capture key features of sensory perception, in particular response to differences and diminishing sensitivity (Weber's law).

In this context it is important to highlight that the most salient attribute receives a larger weight, and that the saliency of an attribute depends on the distance of its value from the average. We can consider the choice between food items as an evaluation of taste t and healthiness (inferred calories c). The taste attribute is equally observed across treatments, but the calorie attribute varies because of the coarseness of the label. I follow the example in Bordalo, Gennaioli and Shleifer (2013) and use a saliency function satisfying homogeneity of degree zero to obtain a simple statement on which attribute should be salient for each product (Proposition 1, p. 809). In a choice set without a dominant product, the attribute c is salient for the relatively *low-calorie* product j iff $c_j/\bar{c} < t_j/\bar{t}$.

The adoption of a more granular label has no effect on taste t , but it can impact the calorie attribute through c/\bar{c} . When we consider a binary choice set we have three scenarios:

- The two products have the same coarse label, and also the same detailed label. No change in $c_j/\bar{c} = 1$.
- The two products have the same coarse label, and different detailed labels. Then c_j/\bar{c}

decreases, and it is more likely that the calorie-healthiness attribute is *more* salient (if the change is large enough), and would never lead to lower salience for the attribute.

- The two products have different coarse labels, and different detailed labels. The result is ambiguous and depends on how the partition divides the range and where the calorie amount lie on the range. The calorie-healthiness attribute becomes less salient only if c_j/\bar{c} increases, and the ratio increase is large enough. A positive change in the ratio occurs with probability $Pr(c_{j,\text{detailed}} - c_{j,\text{coarse}} - c_{k,\text{detailed}} + c_{k,\text{coarse}})$ but this does not guarantee that the change is sufficient to generate a salience change.

C Appendix: Nutrition Labels, Institutional background

Food choices provide a natural setting to study the effect of label precision. The shelves of the supermarkets are filled with products like low-fat yogurts, high-fiber cereals, reduced-calorie chocolate, and sugar-free snacks, all examples of coarse messages. In this section I provide some background about nutrition labels and their role in food choice, and I discuss the relevance of studying this topic in light of the recent discussion about this topic in many countries, including United States and Chile. Readers who want to jump to the experimental design can simply skip this section. I am confident that their personal experience as consumers is sufficient to understand the role of nutrition labels. Coarse labels are used in different domains but food labels are of particular interest for two reasons: a policy and an information one. First, many countries have recently discussed the adoption of different kinds of labels as part of the effort to reduce obesity (more on this in Section 2.2) and as an alternative to other policy tools like sugar tax or soda tax. Second, coarse labels provide information on product characteristics that are already described on the package. Packaged foods typically indicate exact information on nutritions and ingredients in a Nutrition Facts panel, and front-of-package labels are not revealing information that was previously missing, but present it in a way that is simpler or more salient for the consumer. This leads to a natural null hypothesis: coarse labels should not affect consumer behavior at all since the information is already provided. Consumers might not be sufficiently motivated to consult the nutrition fact, and in that case the coarse labels might be providing relevant information. Marketing studies indicate that grocery shoppers spend on average 12 seconds purchasing from a category when in store (Hoyer, 1984; Dickson and Sawyer, 1990) and 10 seconds for online grocery (Anesbury et al., 2016), and they rarely pay attention to the nutrients,⁶⁸ and

⁶⁸In a survey conducted in France, Mannell et al. (2006) report that only 45% of the consumers read the nutrition labels, with the majority of them reading labels only occasionally. Derby and Levy (2001) use the

what is written on the front is more prominent for consumers' choices. With consumers making fast choices, understanding how they respond to package messages is important to design effective policies.

In the rest of this section I provide more background about nutrition labels. First, I discuss the different types of labels that are commonly used. After, I provide a brief overview of the current regulation in the United States. Finally, I discuss some mandatory and voluntary nutrition labels recently introduced in Latin American and European countries.

C.1 Typologies of Nutrition Labels

Nutrition labels can have different content and purposes, and it might be confusing to refer to the entire universe of labels without clarifying what I am focusing on in the study and which other relevant features are not part of this investigation. In the experiment I will focus on a very specific type of labels: interpretive, nutrient-specific, mandatory, front-of-package labels, and I am about to clarify what this means. A possible multidimensional classification of labels can consider four features: interpretation (simplification of the numerical value), synthesis (summary of multiple nutrient), obligatoriness (mandatory to display), and position (on the front of the package).⁶⁹

Interpretation: Reductive vs. Interpretive. Reductive labels contain the same numerical information as in the Nutrient Facts panel (NFP), and interpretive labels provide an evaluation using for example colors, stars, or health symbols. Typical examples include Facts Up Front as reductive label (show few nutrient information on the FOP) and warning messages as interpretive label (indicate that the calorie content is above a recommended threshold).

Synthesis: Nutrient-specific vs. Summary. Nutrient-specific labels refer to one or more individual nutrient, and summary labels aggregate multiple nutrients in a single scale. Typical examples include multiple traffic lights as nutrient-specific labels (indicate for each nutrient if it is high or low) and Nutri-score as summary label (combine calories, sugar, fat, and other nutrients into a single value).

Obligatoriness: Voluntary vs. Mandatory. Public authorities regulate what the product package must, must not, and can contain. Many countries, including the United States, require information about the main nutrition facts in a detailed numerical format and let the manufacturers free to add claims as long as they are not confusing for the consumers.

data of the US Diet and Health Survey and highlight that one third of respondents declare they had changed their decision to by a product because of nutrition label information.

⁶⁹Ikonen et al. (2020) provide a classification of reductive, interpretive, and summary labels and compare their effectiveness in a meta-analysis.

This regulation is called “truth in advertising” and gives the manufacturer large flexibility on how to emphasize positive properties of the product. When standard claims are regulated by public authorities, some terms or combinations of terms could not be used freely. For example, FDA regulates the use of “low calorie” but terms like “low sugar” are not allowed on product packages.

Position: Front-Of-Package vs. Back-Of-Package. Regulators rarely indicate where the label should be displayed. One notable exception is the warning label introduced in Chile, that is mandatory and must be shown on the front of the package to guarantee higher visibility.

C.2 Nutrition Labels in the United States

In the United States the Food and Drug Administration regulates food labeling with a combination of mandatory information (nutrition facts panel) and truth-in-advertising regulation of nutrition claims.⁷⁰ The nutrition facts panel located on the back-of-package is mandatory since the introduction of the Nutrition Labeling and Education Act (NLEA) in 1990 (Moorman, 1996; Balasubramanian and Cole, 2002; Abaluck, 2011), that became fully effective in 1994. The NLEA also introduced other tools to simplify the comparison across products, including the food ingredient panel, standard serving sizes, and definitions for claims such as *low-fat* and *light*. The objective of this regulation is to prevent ambiguous claims that could confuse consumers (Mayer, Scammon and Golodner, 1993; Hoadley and Rowlands, 2014). Requirements and standard terms faced minor adjustments since the NLEA; FDA required additional information in the nutrition labels since 2003, and redesigned the label in 2016. In 2010 the Institute of Medicine and FDA created a commission to collect and analyze evidence related to the impact of food Front-Of-Package labels on consumer behavior (Boon et al., 2010; Wartella et al., 2012). The initial motivation was to provide guidelines for a new standardized FOP label for US consumers, but the project was prematurely terminated in 2012 without providing a final answer. In the same year various US manufacturers voluntarily adopted the Facts Up Front label, that displays some key nutrition information on the front of the package. The FDA commission wrote two reports that discuss the plurality of functions of front-of-package rating systems and symbols and support the goal to target the “general population,” with emphasis on fighting obesity. But they also highlight that given the current knowledge they recommend a system equivalent to the EnergyStar (voluntary and binary) used by the Environmental Protection Agency to encourage consumer

⁷⁰FDA regulation of food labeling is part of the Code of Federal Regulations, Title 21, Chapter 1, subchapter B, Part 101, available on [FDA website](#).

to purchase energy-efficient products. According to the report, the system should use “*one simple, standard symbol translating information from the Nutrition Facts panel (NFP) on each product into a quickly and easily grasped health meaning, making healthier options unmistakable*” (Wartella et al., 2012). It appears that the need for simplicity and coarseness is justified by the concern that consumers would be confused, and make mistakes, by adopting more complex rating systems.

C.3 Coarse Food Labels in the World

Despite the attempt of the FDA commission, the United States have not reached a standardized front-of-package label. Other countries started from the very same premises and recently introduced labels that aim to guide consumers towards healthier options. Some interesting insights come from Chile, Ecuador and France, that between 2013 and 2017 introduced standardized labels. In all three cases the labels are coarse, with a number of possible messages between two and five.

Chile: Warning labels. Chile’s food labeling and advertising law was approved in 2012 and enacted in 2016. It requires food and beverage packages to display warning labels if products exceed the thresholds for calories, sugars, saturated fat, or sodium.⁷¹ It is the manufacturers’ responsibility to implement the requirement and display the label on the front of the package, and products can display up to four warning labels (if they exceed the thresholds for all the four nutrients), but the message is binary: for each nutrient the product can either display or not display the label. The law also prohibited the sale of products with labels in schools and restricted the advertising to children under fourteen. Early studies on the effect of this reform show large and persistent effect (Pachali et al., 2020; Barahona et al., 2020; Alé-Chilet and Moshary, 2020), with an overall 6.5 reduction in calories consumed.

Ecuador: Traffic lights. Ecuador’s traffic lights labeling was approved in 2013 and enacted in 2014. Traffic light labels indicate the amount of sugar, fat, and salt and can take four values (high, medium, low, none) expressed through a combination of color and text. Sandoval, Carpio and Sanchez-Plata (2019) show that the reform did not have a significant effect on soft drink purchases. The difference with the successful Chilean reform might be due to the position of the label; many Ecuadorean manufacturers placed the labels on the side or back of the products, making them less visible than the Chilean one. Another possible concern, more related to this study, is that Ecuadorean labels might be more complex to process; consumers might be unsure about the meaning of a medium-label whereas in the

⁷¹Chilean warning labels thresholds are expressed with respect to 100 grams servings for food, and 100 milliliters for beverages. The thresholds were gradually reduced in 2018 and in 2019.

case of Chilean labels there are only two possible outcomes (label or no label).

France: Nutri-score. In 2017 the French government recommended a voluntary nutritional rating system. The scoring takes into account calories, fat, sugar, and other nutrients per 100 grams of product and generates a letter/color combination between A/green and E/red. The Nutri-score rating system has also been recommended by the European Commission, WHO, and public authorities in Belgium, Germany, Netherlands, Spain, and Switzerland. The score is determined as a coarsening of a numerical value, calculated according to the UK Food Standards Agency nutrient profiling system. The labels are voluntary but laboratory and field studies have shown that Nutri-score outperformed other labeling systems in driving consumers towards healthier food options (Dubois et al., 2020; Julia et al., 2018).

D Appendix: Experimental design, implementation, and recruitment

D.1 Timeline of the experimental session

Participants are recruited through CloudResearch Prime Panels. Before entering the study they have no information about the nature of the task they will completed. They only know it is a 20-minutes study that they can complete on their computer. Once they click on the study, participants are asked to complete demographic questions, to meet the quotas, and screening questions, to exclude respondents that show little attention. Participants are asked to read carefully the instructions and answer three comprehension questions. They can proceed with the study only after they answer correctly all the questions. 31% of the respondents abandon the study during the instructions or comprehension questions, and 2% of the respondents abandon the study after successfully passing the comprehension questions. After answering the questions, participants are randomly assigned to one of the four treatments. They receive one page with treatment-specific instructions about the meaning of the labels they will encounter in the study. The rest of the session contains the product choice task, additional choice tasks, and a final questionnaire.

D.2 Implementation

Products and Calories Distributions. The study contains choice sets comprised of two or four items. The products used are cereals and salty snacks that can be typically purchased in grocery stores. Within each trials, all the product belong to the same category.

These product categories have the advantages of being largely popular food (Hamrick, 2016; Rhodes et al., 2017) and displaying a large heterogeneity in the nutrient characteristics. Using publicly available data of the US Department of Agriculture, I extract the nutritional information on calories per serving and I include that in the study as part of the instructions (to explain calorie heterogeneity across products) and task (to explain how each label is associated with a particular range of calorie values).

Randomization. Randomization is operated at the subject level and at the trial level. Subjects are randomly assigned a treatment and an order of the blocks of choice trials. Trials are grouped into blocks to avoid that the same item appears several times at short distance, and to maintain the same choice set size for several rounds in a row. For each trial, I randomize the position of the products on the screen to remove position concerns.

Programming. The study is conducted using the online platform Qualtrics and custom JavaScript code. Qualtrics does not have specific features for mousetracking, and other features, for example the ranking questions, have limited flexibility. I developed various JavaScript tools that facilitate product visualization and data collection and can be used for a variety of online studies.

D.3 Product Choice Task

The choice task is comprised of 80 trials. The trials are equally divided between product categories (snacks or cereals), choice set sizes (2 items or 4 items), and price levels (low-calorie products more or less expensive). The trials are also designed to cover a large range of combinations of low and high-calorie products.

D.4 Questionnaire

The final questionnaire contains questions about food literacy and habits.

- Allergies and dietary restrictions.
- Food and behavior:
 - Who does the majority of grocery shopping in your household?
 - How likely are you to purchase cereals/snacks?
 - How important is [price/brand/taste/healthiness/ethical/organic] when you purchase food?
 - When you think about your body, how do you consider yourself? [Underweight, healthy weight, overweight, obese]

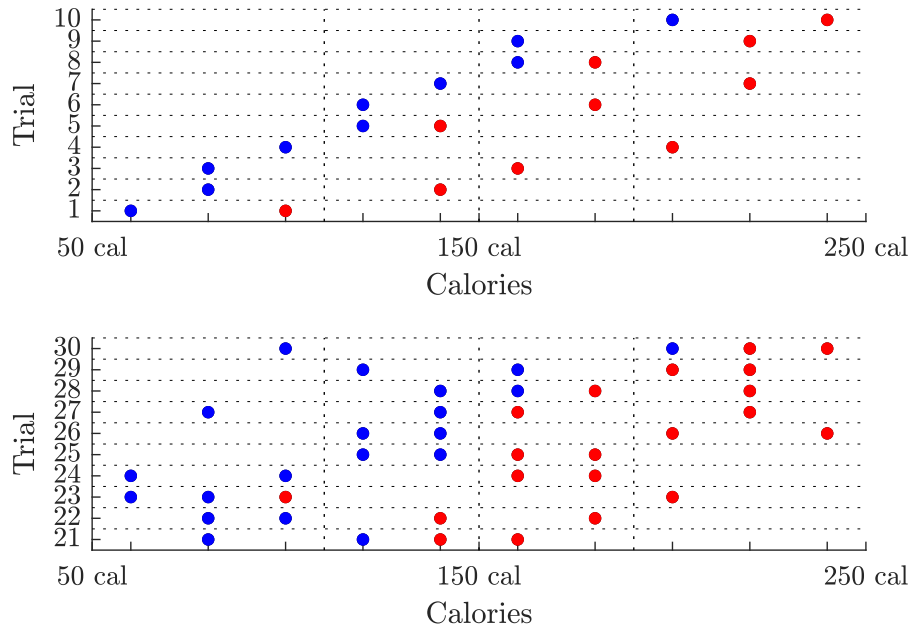


Figure 19: Calorie composition across choice trials. Trials 1-10 indicate binary choices, trials 21-30 indicate quaternary choices. The red dots indicate the high-calorie products in each choice set (above the median value), blue dots indicate low-calorie products.

- Food literacy: 8 items, Likert 5 scale, from the questionnaire validated in Poelman et al. 2018. The scores are summed and normalized to be between 0 (low food literacy) and 100 (high food literacy)
 1. Do you check the nutritional labels of products for calories, fat, sugar or salt content?
 2. Do you compare the calories, fat, sugar or salt content of different products?
 3. Do you check the price of the products you are buying?
 4. Do you check the ingredients of the products you are buying?
 5. If you have something to eat, do you reflect on what you have eaten earlier that day?
 6. If you have something to eat, do you take account of what you will eat later that day?
 7. Do you purchase healthy foods, even if they are a bit more expensive?
 8. Do you purchase healthy food, even if you have limited money?

D.5 Recruitment

The recruitment and compensation of participants for completing the study was performed by CloudResearch, previously known as TurkPrime. They provide the following information about their sampling, screening, and compensation.

Sampling. The CloudResearch Prime Panels platform is an aggregation of double opt-in participant panels. This means that participants are recruited into the panel by navigating to a website and signing up to take surveys. Prime Panels offers a convenience sample; it does not offer access to participants recruited with probability-based sampling methods. To increase the representativeness of samples gathered on Prime Panels, CloudResearch applies demographic quotas and post data collection weights. When researchers want to approximate the U.S. population, CloudResearch applies quotas for gender, age, race, and ethnicity (i.e., Hispanic or non-Hispanic) that are matched to the U.S. Census. The application of such quotas is known as purposive sampling, a technique common among researchers who use online panels. After the data are gathered, CloudResearch attempts to further increase representativeness by applying survey weights. Both of these methods are common among researchers who sample from online panels (AAPOR, 2010). When sampling is intended to match the demographics of a state or a city, CloudResearch will set quotas based on available local demographic data. One limitation of the sampling design is that the survey is only offered in English. Participants are free to respond to the survey on any device they wish.

CloudResearch applies quotas for gender, age, race, and ethnicity (i.e., Hispanic or non-Hispanic) that are matched to the U.S. Census.

- Age: 24% 18-29, 24% 30-44, 24% 45-59, 28% 60-99.
- Gender: 50% Female, 50% Male.
- Hispanic: 16% Hispanic, Latino, or Spanish origin.
- Race: 4% Asian, 12% Black, 82% White.
- Region: 21% MW, 17% NE, 38% S, 24% W.

Screening. All participants who enter a survey on Prime Panels are first run through CloudResearch’s Sentry data validation system. Sentry is a short (> 1 minute) pre-study screening that ensures participants are attentive, engaged, and ready to participate. People who fail Sentry are routed away from the survey. Research shows that these sorts of short, pre-study validation measures are an effective way to increase data quality obtained from market research panels like Prime Panels (Chandler et al., 2019; Litman et al., 2020).

Participant Compensation and Experience. Participants on Prime Panels are compensated in a variety of ways. Some receive direct monetary compensation. Others participate in studies for a donation to charity or for points that they can redeem for prizes. Compensation is determined at the panel level, meaning that each individual panel decides how much and in what form participant compensation is administered. Panels do not share compensation information with CloudResearch.

E Appendix: Experimental results, product choice

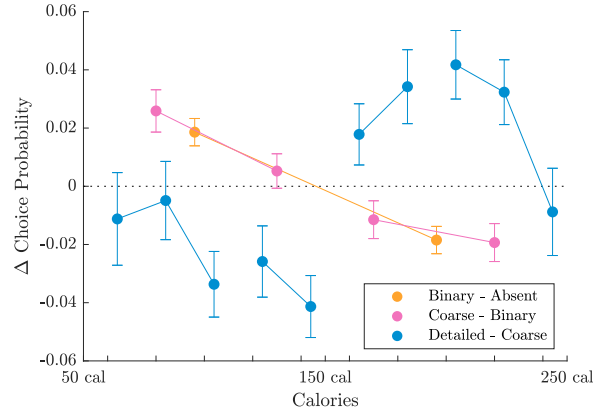


Figure 20: Information disclosure: marginal label effect. The x-axis indicate the inferred product calories, with the products grouped based on the label displayed on each treatment. The y-axis indicate the change in the probability of choosing the product with respect to the adjacent treatment. Joint segments indicate groups that were receiving the same label in the coarser label, and different labels in the more detailed label. All the segments but one show a downward trend: the gradual introduction of more detailed labels moves choices towards relatively low-calorie products. The segments differ in their absolute position, with the detailed treatment (blue) showing an overall decrease in the choice of products below 150 calories.

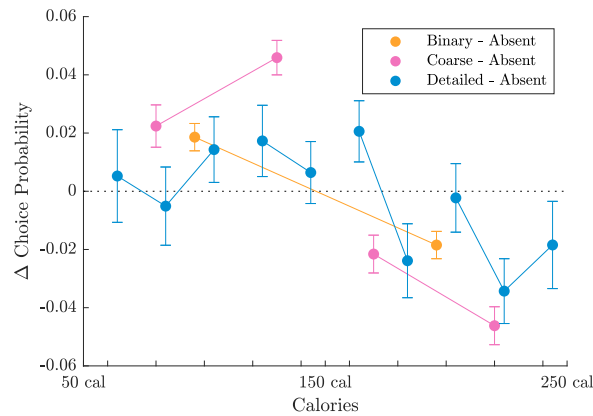


Figure 21: Information disclosure: total label effect. The x-axis indicate the inferred product calories, with the products grouped based on the label displayed on each treatment. The y-axis indicate the change in the probability of choosing the product with respect to the control treatment (Absent). Coarse labels show the larger deviations from the choices in the absent treatment.

F Appendix: Experimental results, controls




Figure 22: Control for order effect in the session. Product choice trials are grouped in four blocks of 20 trials each. The same choice set is presented twice during the session, with different prices, at least 20 trials away from each other. The figure shows whether and how search and choices change during the session. The four panels, from top to bottom, show: 1) Probability of flipping the package to consult back-of-package information. The probability is significantly higher in the first block. There are no statistical differences within each block. 2) Response time. We observe a longer response time for the first trial of each block, and an overall speeding trend, with the average response time decreasing from 8 to 5 seconds from the first to the last block. 3) Probability of choosing the relatively low-calorie product and 4) Average number of calories of the chosen product. None of these two variables shows significant effect within or across blocks, indicating that the main results of the study are robust over time.




Figure 23: Control for size of the choice set. Binary choice sets (left) and Quaternary choice sets (right). Variables of interest, from top to bottom: probability of choosing the relatively low-calorie product, average number of calories, average response time, probability of flipping the package.

G Appendix: Experimental results, heterogeneity analysis

Group	Avg Calories			Pr(low-cal)		
	A	$A \rightarrow C$	$C \rightarrow D$	A	$A \rightarrow C$	$C \rightarrow D$
All	143.06 (0.8)	-4.08*** (0.91)	2.62** (0.92)	0.53 (0.01)	0.05*** (0.01)	-0.03** (0.01)
Female	142.74 (1.12)	-4.41*** (1.32)	1.63 (1.38)	0.52 (0.01)	0.06*** (0.02)	-0.03 (0.02)
Male	143.47 (1.15)	-3.83** (1.27)	3.54** (1.21)	0.53 (0.01)	0.04* (0.02)	-0.04* (0.01)
Age 18-29	141.88 (2.06)	-2.71 (1.86)	3.69 (1.38)	0.54 (0.03)	0.03 (0.02)	-0.05* (0.02)
Age 30-44	144.51 (1.18)	-4.31** (1.71)	1.27 (1.68)	0.51 (0.02)	0.05** (0.02)	-0.02 (0.02)
Age 45-59	141.66 (1.84)	-2.52 (2.03)	3.28 (2.11)	0.54 (0.02)	0.03 (0.03)	-0.03 (0.02)
Age 60+	143.59 (1.47)	-5.95*** (1.75)	2.40 (1.86)	0.52 (0.02)	0.08*** (0.02)	-0.03 (0.02)

Table 10: Treatment effect, subjects grouped based on demographic characteristics (gender and age). SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Low-calorie products: calories per serving below median in the current choice set.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	144.71*** (0.90)	146.77*** (0.72)	145.62*** (0.89)	144.62*** (1.34)	140.33*** (2.11)
Treatment B	-0.06 (1.20)	-0.29 (1.05)	-0.28 (1.04)	-0.03 (1.07)	-1.33 (1.63)
Treatment C	-3.38** (1.35)	-4.00*** (1.19)	-3.96*** (1.18)	-3.81*** (1.20)	-2.16 (1.88)
Treatment D	-1.78 (1.33)	-2.43** (1.20)	-2.39** (1.20)	-2.18* (1.19)	-2.25 (1.71)
Flip Package		-15.29*** (1.77)	-14.24*** (1.75)	-14.37*** (1.71)	-17.88*** (2.16)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	32878	32878	32878	31438	10397
Adjusted R ²	.001	.012	.42	.422	.502

Table 11: Treatment effect - Food Literacy Index - Low.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	141.41*** (0.89)	145.26*** (0.78)	144.07*** (0.98)	145.30*** (1.76)	141.39*** (2.82)
Treatment B	-3.02** (1.30)	-4.48*** (1.16)	-4.54*** (1.15)	-5.55*** (1.19)	-5.82*** (1.91)
Treatment C	-4.45*** (1.33)	-6.01*** (1.15)	-6.05*** (1.14)	-6.14*** (1.20)	-8.98*** (1.90)
Treatment D	-1.02 (1.38)	-2.29* (1.27)	-2.32* (1.26)	-2.18* (1.29)	-4.23** (2.06)
Flip Package		-13.01*** (1.12)	-13.44*** (1.08)	-12.09*** (1.17)	-9.41*** (1.59)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	35573	35573	35573	33413	12395
Adjusted R ²	.001	.014	.438	.438	.518

Table 12: Treatment effect - Food Literacy Index - High.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	143.37*** (1.05)	146.53*** (0.86)	144.90*** (1.04)	144.95*** (1.44)	142.35*** (2.22)
Treatment B	-2.26 (1.43)	-3.38*** (1.25)	-3.37*** (1.24)	-3.11** (1.31)	-3.26* (1.94)
Treatment C	-5.04*** (1.44)	-5.95*** (1.24)	-5.95*** (1.23)	-5.83*** (1.27)	-3.56* (1.91)
Treatment D	-1.47 (1.45)	-2.97** (1.31)	-2.96** (1.30)	-2.95** (1.30)	-2.47 (1.86)
Flip Package		-13.77*** (1.46)	-13.72*** (1.43)	-14.00*** (1.50)	-14.78*** (1.69)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	33190	33190	33190	31190	10840
Adjusted R ²	.002	.013	.426	.425	.512

Table 13: Treatment effect - Household income - Low.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	142.84*** (0.79)	145.84*** (0.67)	145.07*** (0.87)	146.84*** (1.76)	140.00*** (3.06)
Treatment B	-1.23 (1.17)	-1.94* (1.03)	-1.94* (1.02)	-2.64** (1.03)	-3.27* (1.75)
Treatment C	-3.11** (1.30)	-4.52*** (1.14)	-4.49*** (1.13)	-4.30*** (1.16)	-7.64*** (2.01)
Treatment D	-1.58 (1.34)	-1.97* (1.19)	-1.95 (1.19)	-1.40 (1.22)	-3.93* (2.15)
Flip Package		-14.65*** (1.23)	-14.46*** (1.19)	-13.26*** (1.23)	-11.36*** (1.85)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	35261	35261	35261	33661	11952
Adjusted R ²	.001	.014	.43	.43	.503

Table 14: Treatment effect - Household income - High.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	142.46*** (0.97)	145.70*** (0.76)	144.44*** (0.92)	143.80*** (1.37)	139.92*** (2.20)
Treatment B	-0.16 (1.31)	-1.20 (1.10)	-1.21 (1.08)	-0.73 (1.12)	-0.09 (1.77)
Treatment C	-5.19*** (1.42)	-5.75*** (1.18)	-5.75*** (1.17)	-4.93*** (1.21)	-4.86** (1.98)
Treatment D	-0.60 (1.38)	-1.68 (1.20)	-1.68 (1.19)	-0.88 (1.18)	-0.93 (1.84)
Flip Package		-16.15*** (1.37)	-16.20*** (1.34)	-16.21*** (1.40)	-17.67*** (1.65)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	36546	36546	36546	34466	11850
Adjusted R ²	.002	.018	.43	.43	.515

Table 15: Treatment effect - Education - Low.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	143.66*** (0.82)	146.42*** (0.74)	145.31*** (0.95)	150.54*** (1.71)	145.18*** (2.80)
Treatment B	-3.25*** (1.25)	-3.96*** (1.14)	-3.95*** (1.15)	-5.14*** (1.17)	-6.94*** (1.82)
Treatment C	-2.77** (1.26)	-4.24*** (1.14)	-4.18*** (1.13)	-4.43*** (1.18)	-6.47*** (2.01)
Treatment D	-2.42* (1.37)	-3.18** (1.27)	-3.15** (1.27)	-3.86*** (1.31)	-6.14*** (2.24)
Flip Package		-11.97*** (1.26)	-11.64*** (1.22)	-10.59*** (1.22)	-7.27*** (1.70)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	31905	31905	31905	30385	10942
Adjusted R ²	.001	.01	.427	.429	.502

Table 16: Treatment effect - Education - High.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	141.80*** (0.83)	145.48*** (0.67)	144.13*** (0.84)	143.97*** (1.36)	141.28*** (1.96)
Treatment B	-1.76 (1.15)	-3.05*** (0.98)	-3.10*** (0.98)	-3.18*** (1.01)	-3.47** (1.39)
Treatment C	-3.98*** (1.18)	-5.53*** (1.01)	-5.52*** (1.00)	-5.27*** (1.05)	-6.17*** (1.51)
Treatment D	-2.41** (1.23)	-3.53*** (1.08)	-3.49*** (1.07)	-2.73** (1.08)	-2.98** (1.49)
Flip Package		-14.97*** (1.10)	-14.93*** (1.08)	-14.40*** (1.13)	-13.54*** (1.34)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	46137	46137	46137	43617	21212
Adjusted R ²	.001	.016	.436	.434	.486

Table 17: Treatment effect - Calorie preference - Low.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OLS. DV: Calories per serving of the selected product.					
	(1)	(2)	(3)	(4)	(5)
Constant	145.27*** (0.95)	147.09*** (0.85)	146.15*** (1.08)	146.54*** (1.84)	142.25*** (3.30)
Treatment B	-0.97 (1.41)	-1.23 (1.32)	-1.15 (1.30)	-1.77 (1.36)	-0.82 (2.01)
Treatment C	-3.34** (1.62)	-3.71** (1.46)	-3.69** (1.45)	-3.77** (1.48)	-2.00 (2.25)
Treatment D	0.63 (1.50)	-0.06 (1.46)	-0.07 (1.44)	-0.44 (1.44)	-2.12 (2.13)
Flip Package		-11.24*** (1.74)	-10.83*** (1.67)	-10.14*** (1.64)	-5.26*** (1.65)
Choice set Controls			✓	✓	✓
Demographic Controls				✓	✓
Restrictions					Monotonic
Observations	22314	22314	22314	21234	1580
Adjusted R ²	.001	.008	.417	.42	.744

Table 18: Treatment effect - Calorie preference - High.

SE clustered at the subject-category level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.