Stanford Data Science for Social Good: Technical Report

Using LLMs and Text Processing to Analyze Parentally-Reported Food Reactivity Case Data

Amira Mahmedjan, DSSG Fellow
Stanford University

Sean Nesamoney, DSSG Fellow
Stanford University

Michael Howes, Technical Mentor
Stanford University

Dr. Trillitye Paullin, Faculty Mentor
Hoover Institution Veteran Fellow,
Stanford University

Dr. Annie K. Lamar, Technical Adviser
University of California, Santa Barbara

***Abstract*—Food allergies among children have emerged as a significant public health issue in the United States. This study investigates parentally-reported food reactivity in children, focusing on identifying patterns, severity, and the broader social implications. Our research assesses the prevalence and impact of these reactions, explores the effectiveness of dietary modifications in alleviating symptoms, and examines the emotional, financial, and social consequences for families. By analyzing these aspects, the study informs on more effective interventions and support systems for managing food reactivity in children.**

ADOLESCENT FOOD ALLERGIES are a growing public health concern in the United States, with prevalence rates rising over the past few decades. Approximately 5.6 million children under 18 years of age in the U.S. suffer from food allergies, equating to about 1 in 13 children [1], [2]. These allergies can lead to a range of symptoms, from mild to life-threatening, and they have significant implications not only for the affected children but also for their families and healthcare providers. Understanding the scope, impact, and management of these food reactivities is crucial for developing effective interventions and support systems.

Date of publication Summer 2024

INTRODUCTION

The purpose of our study is to investigate parentally reported food reactivity in children, with a focus on identifying patterns, severity, and the social implications of these reactions. Specifically, our study aims to accomplish three goals: (1) assess the prevalence and severity of food reactivity in children as reported by parents, (2) explore the relationship between dietary modifications and the alleviation of food reactivity symptoms, and (3) evaluate the impact of food reactivity on families, including financial, emotional, and social aspects.

First, we provide a comprehensive overview of existing research on food allergies and reactivities in children, highlighting key findings and gaps in knowledge. We also detail the methods used to conduct our research, including relevant data cleaning procedures and analytical tools. We present the findings from our data analysis, focusing on the prevalence, severity, and patterns of food reactivity, as well as the outcomes of dietary modifications. In addition, we discuss the implications of our findings, particularly in the context of improving food reactivity management and support for affected families. Finally, we summarize the key findings of our study and suggest areas for future research.

DATA

The dataset used in this study was collected by Free to Feed through a voluntary survey. Out of 558 children, 410 children (approximately 73%) observed were from the United States. The data collected is largely qualitative, consisting of a combination of standardized yes/no questions and free-response entries. The survey asked respondents questions pertaining to theirs and their child's diets, symptoms, medical history, prior treatment plans, and related information (Table 1).

The respondents were primarily individuals who had personal connections with Free to Feed or were recommended to Free to Feed by their current provider. Therefore, the scope of this study is not entirely representative of all cases of food reactivity, and its findings should be considered in the context of the backgrounds and demographics of its participants.

Table 1. Summary of data collected from survey

Number of Observations	558
Number of Questions	155
Percentage of Yes/No Questions	70%
Percentage of Free-Text Questions	15%
Percentage of Other Questions	15%

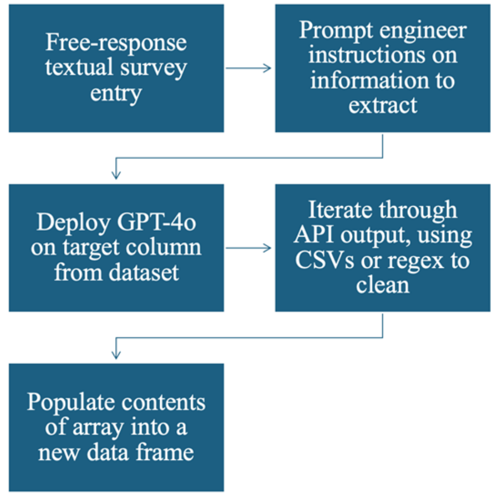

Figure 1. Workflow of the implementation of GPT-4o [3] to extrapolate key information from the dataset.

DATA CLEANING

Due to the structure of the survey, many participants' responses were inconsistent and unstandardized. As a result, manual cleaning was required in order to eliminate redundant or unnecessary data, extract target information (such as symptoms and trigger foods), and ensure the accuracy and reliability of the dataset. This process involved addressing missing values, normalizing varied formats, and standardizing the terminology used by respondents. The meticulous cleaning allowed for a more cohesive and analyzable dataset, which was crucial for drawing meaningful insights and making informed decisions regarding food reactivity in children.

In addition to using manual cleaning techniques, we employed large language models, specifically the GPT-4 OpenAI API [3], to improve efficiency and streamline data organization. The API was utilized to automatically identify and correct inconsistencies within the dataset, enhance the precision of symptom and trigger food categorization, and facilitate the standardization of responses across various entries. This integration of advanced language processing capabilities allowed for a more systematic approach to data cleaning, significantly reducing the time and effort required while maintaining the integrity and robustness of the dataset.

The effectiveness of GPT-4o relied on the meticulous engineering of prompts, which were critical for obtaining accurate and relevant responses. To maximize its effectiveness, prompts were crafted thoughtfully, specifying which aspects of the data should be extrapolated and delineating the desired presentation format. This involved providing clear and detailed instructions to guide the model in generating meaningful insights, ensuring that the outputs were not only aligned with analytical objectives but also presented in a coherent and actionable manner.

GPT-4o outputted the desired results in a one-dimensional array, using commas to separate distinct entries. Using customized regular expression sequences, the arrays were iterated through, trimming unnecessary characters from strings and appending information to a new data frame. This process was repeated with different target variables from the overall dataset. As a result, data specific to the three project objectives were cleaned, extracted, and made available for thorough analysis procedures.

RESULTS

Symptom-Trigger Mapping

The first aim of our project is understanding the initial representation of food reactivity in children. To address this aim, we first analyzed the column where parents answered the question “What is / was this child reactive to?”. The parents were then asked to describe each food trigger and corresponding symptoms their child has. The food triggers were then classified by GPT-4o into a standardized list provided to us by Free to Feed for ease of analysis.

The most commonly reported trigger is dairy (including cow’s milk), followed by soy, eggs, and tree nuts.These results are summarized in Figure 2.

The survey also asked parents to report the symptoms associated with each food trigger. The most commonly reported physical symptoms are eczema, mucus and blood in the stool, and vomiting. The majority of children in the dataset grapple with severe reactions to their respective food triggers. These results are summarized in Figure 3.

Almost all children started facing food reactivity in a range between 0-39 weeks (less than a year of age). 287 children (around 51%) were reported to first experience food reactivity between 0-9 weeks of age. These results are summarized in Figure 4.

Finally, we investigate parentally-reported sensitivity levels. Parents were asked to indicate whether their child has high (cannot consume even trace derivatives), moderate (can consume derivatives, but not obvious forms even at the baked level), low (can consume lower levels of broken down forms, but not more whole versions), or very low (can consume whole versions of protein, but not over a certain volume) sensitivity levels. The results showed that the majority of survey participants reported either high or medium sensitivity levels for all common most common food triggers.

Dietary Modifications and Interventions

The second aim of our study is to explore the relationship between dietary modifications and the alleviation of food reactivity symptoms. To address this aim, we focused on two intervention techniques reported in the dataset: elimination diets and reintroduction diets. Elimination diets are a diagnostic tool used by families to pinpoint specific food triggers that are causing certain physiological symptoms.

In this study, elimination diets were conducted by lactating parents, meaning the child was still breastfeeding. By excluding individual foods from a diet, parents can identify which triggers are associated with specific symptoms by using the presence or absence of child reactivity as a proxy. The reintroduction diets are dietary modifications that include reintroducing a known trigger through direct child consumption. Reintroductions can sometimes result in the child’s reaction becoming milder or even disappearing entirely, which are successful reintroductions. Failed reintroductions may result in the persistence or worsening of symptoms, potentially leading to more severe allergic reactions in the future.

To best understand the role of elimination diets in food reactivity intervention, we investigated how long diets lasted before parents observed improvement in their child’s symptoms. To do so, we plotted the distribution of the durations of all elimination diets recorded in the dataset. The survey used in this research asked participants to indicate the length of time (weeks) that the diet was implemented before symptoms had alleviated.

The majority of elimination diets implemented in the study lasted over 20 weeks before parents observed improvements in their child’s symptoms Figure 4. Such findings suggest that the effectiveness of elimination diets rely on extensive trial and error cycles. Nearly 80 diets lasted for nearly one week. It is clear that the durations of elimination diets fall on either extreme, with symptoms alleviating almost immediately or after several weeks.

To further assess food trigger reintroductions, we consider the success rate of reintroduced triggers and

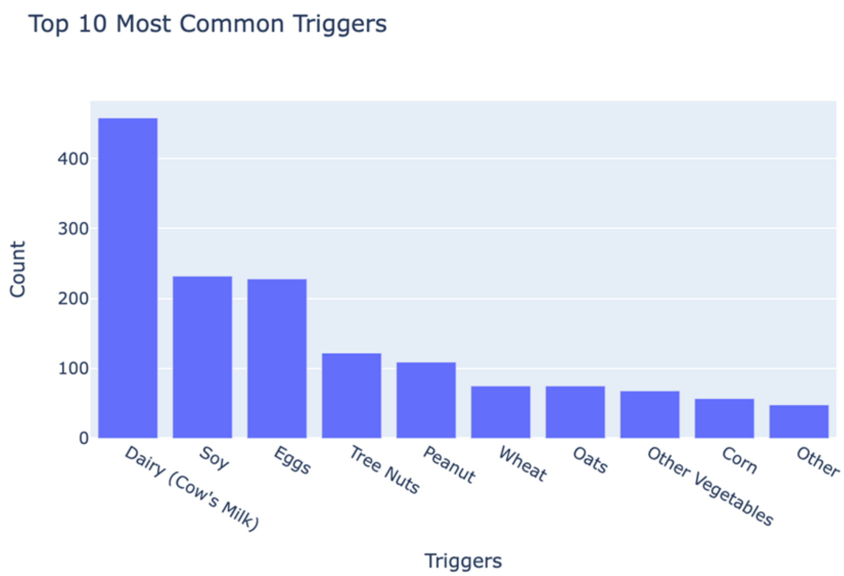

Figure 2. Top ten most common food triggers in food-reactive children as reported by parents.

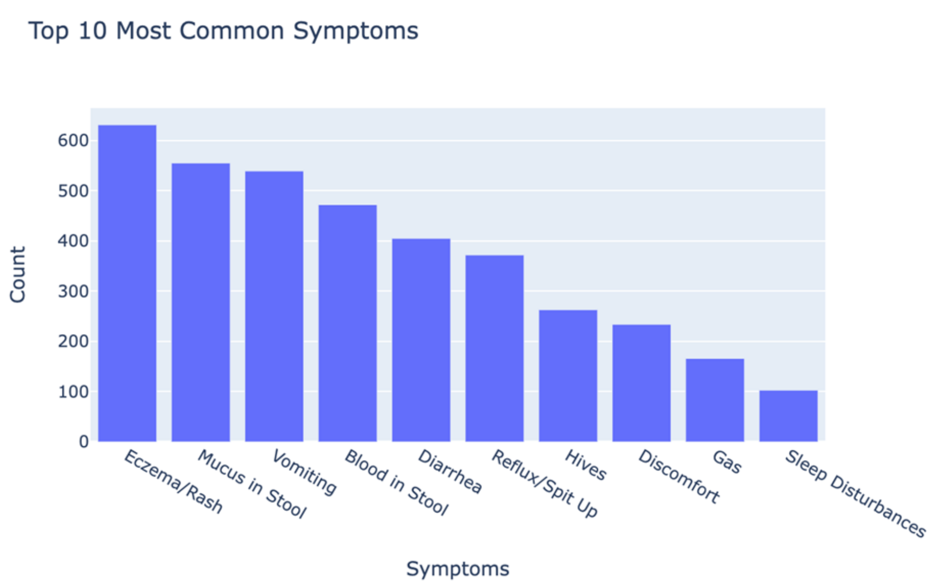

Figure 3. Top ten most common physical symptoms in food-reactive children after exposure to trigger food as reported by parents.

the parentally reported age of reintroduction, the reintroduced trigger. These results are shown in Figure 7

The highest success rates were reported for tree nuts and peanuts, which have reintroduction success rates of 80% and 76%, respectfully. The lowest success rates were reported for dairy (42.9%) and soy (40.9%). It is important to note that while dairy is the most common food trigger, it has one of the lowest success rates. This highlights the struggles most children and families face on their journey with food reactivity. We also performed additional exploratory analysis on the relationship between the age of reintroduction and probability of success. The results from a logistic regression model (with and without sliding window successrates) were statistically insignificant at the α=0.01 level.

Social Impacts of Food Reactivity

Our final aim of this report is to understand how social determinants hinder families from receiving proper care and support while navigating their child's food reactivity. It is especially important to examine health-related issues holistically as there are numerous environmental factors that contribute to the overall well-being of a child and their family. By understanding these factors, we can identify gaps in care and work towards creating more equitable and accessible solutions for those affected by food reactivity.

First, we investigate the prevalence of certain health-related disparities experienced by families in the data set. This data was extracted from a series of yes/no questions asking participants if they had experienced any of the following challenges while navigating their child's food reactivity.

Most families—above 80%—experience increased anxiety, stress, or trauma during their food reactivity journeys. In addition to external stressors, participants report difficulty integrating themselves into social environments and trouble finding adequate child care.

Infant food reactivity challenges also put a significant financial strain on many families. We used data on participants' expenses to understand how this issue impacts families on an economic level. By extrapolating a series of yes/no questions from the survey asking families to indicate which expenses increased while receiving care, we examined how different types of medical and dietary interventions affect overall household spending.

This analysis revealed specific areas where families experience financial burdens, such as higher costs for specialized foods, frequent medical visits, and additional childcare support. Understanding these economic impacts helps to highlight the broader implications of food reactivity beyond health alone, underscoring the need for more accessible and cost-effective solutions.

By analyzing correlations between health disparities (as shown in Figure 8) and expenses (as shown in Figure 9), we investigated the underlying causes of the high prevalence of increased stress, anxiety, and trauma reported by families.

To further analyze factors that affect whether families grapple with certain health disparities and expenses, we looked into whether the number of food triggers affects it. Families are more likely to experience health disparities and expenses if their child has more food triggers. These results are shown in

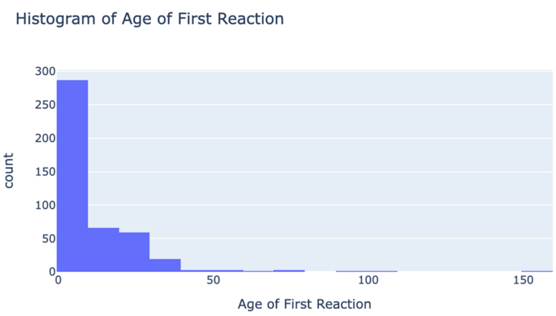

Figure 4. Histogram showing age of child at time of first reported reaction to trigger food.

SURVEY REDESIGN

Throughout our research, we were often limited by the structure and quality of our data. Due to a large amount of text-entry responses, it was at first hard to analyze and aggregate the data, which as mentioned earlier led to us utilizing LLMs. While recognizing the importance of thorough detailed data in medical research, we also faced the other side of the problem with the hardship of applying data science to data that lacked structure completely. Therefore, we propose a redesigned survey to avoid excessive data cleaning and make data more suitable for data science research for Free to Feed's future research.

While the original survey was done with 123FormBuilder, our redesigned proposal was implemented in Qualtrics, which has some additional functionality. There were two main improvements we made to the

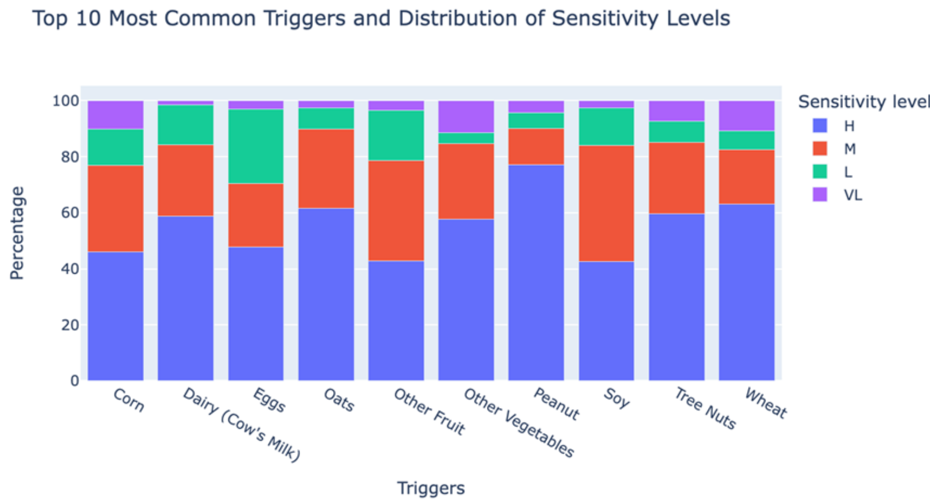

Figure 5. Reported reactivity or sensitivity levels of children to top ten most commonly reported trigger foods.

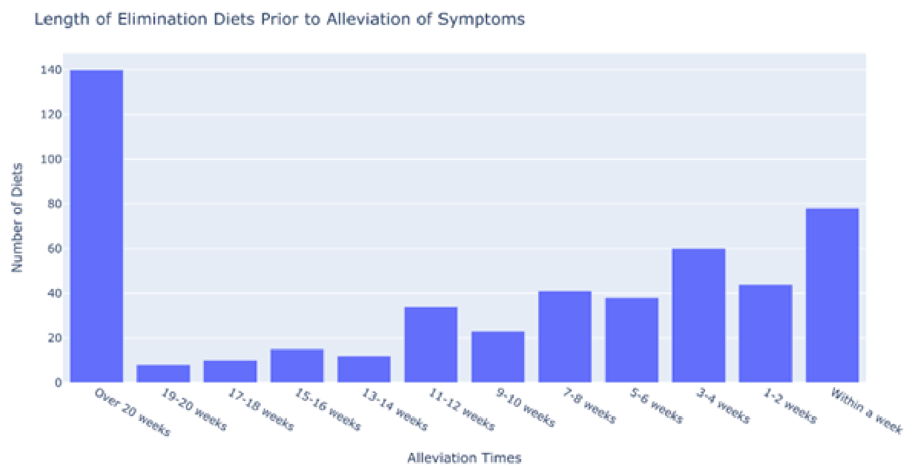

Figure 6. Distribution of durations of all elimination diets reported by survey participants.

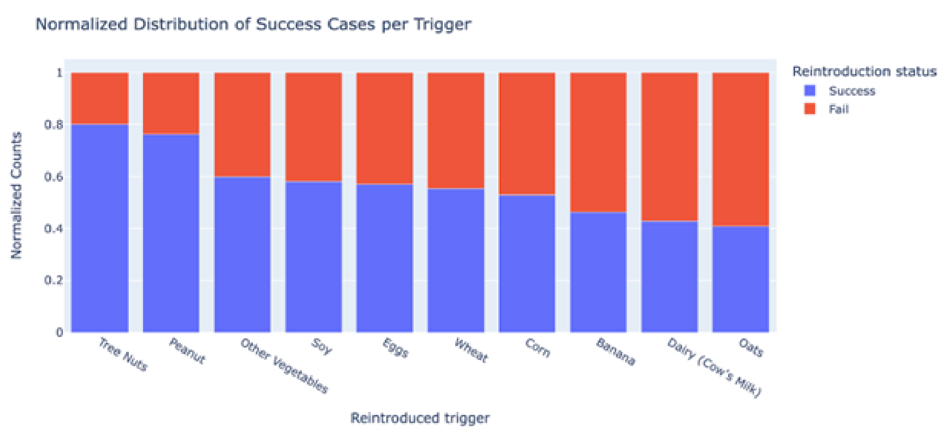

Figure 7. Reported success and failure rates for food reintroductions for the top ten most commonly reported triggers.

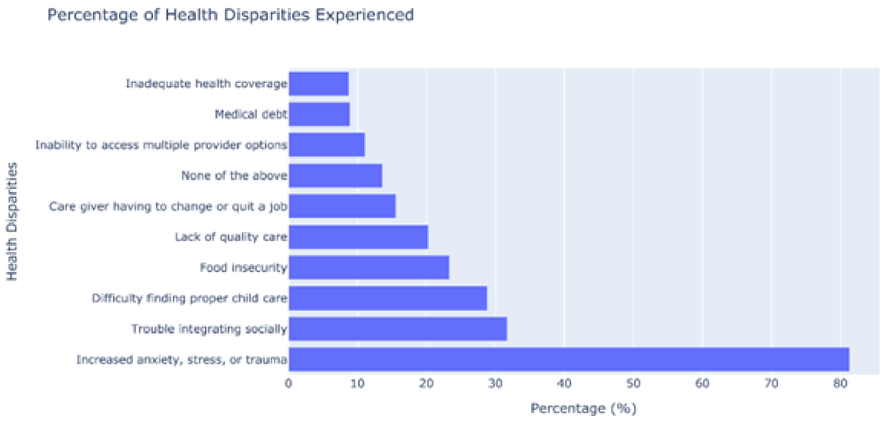

Figure 8. Percentages of health disparities experienced by children in the surveyed population.

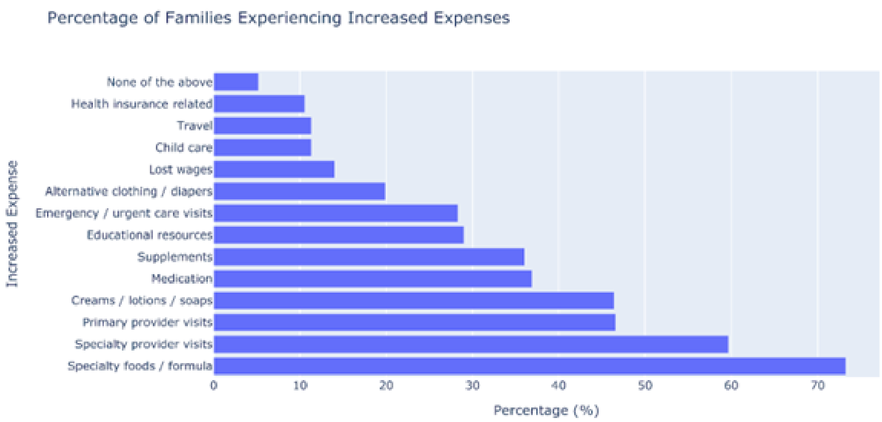

Figure 9. Percentages families reporting increased expenses for various care items and resources.

survey: integrating survey flow and standardizing most answer options to questions.

The need for survey flow came due to the structure of the data collected: while some questions were family-specific (e.g. health expenses questions), others were child-specific (e.g. what is the child reactive to). The original survey asked participants to fill out the form several times for each allergic child. While this was somewhat accommodating the data structure, it was inconvenient for the parents. Many parents accidentally reported different family-specific information when filling out the form several times. Moreover, it is generally more time-consuming and surveys are more likely to be left incomplete when respondents are asked to fill out the same information several times. Therefore, we integrated the survey flow.

In our Qualtrics survey, parents are once asked family-specific questions and how many allergic children they have. Then, they are taken on a loop where they answer the same questions for each reactive child (Figure 13). This way, parents fill out the form only once and do not repeat family-specific information.

Since one food trigger can be named in several different ways by various parents, a free-text response for trigger labeling allows for the introduction of large amounts of noise and error into the dataset. Free to Feed provided us with a standardized list of food triggers and physical symptoms to make available as answer choices.

We also standardized the structure of responses. For example, age variables are now collected through a combination of number-entry boxes and scale drop-

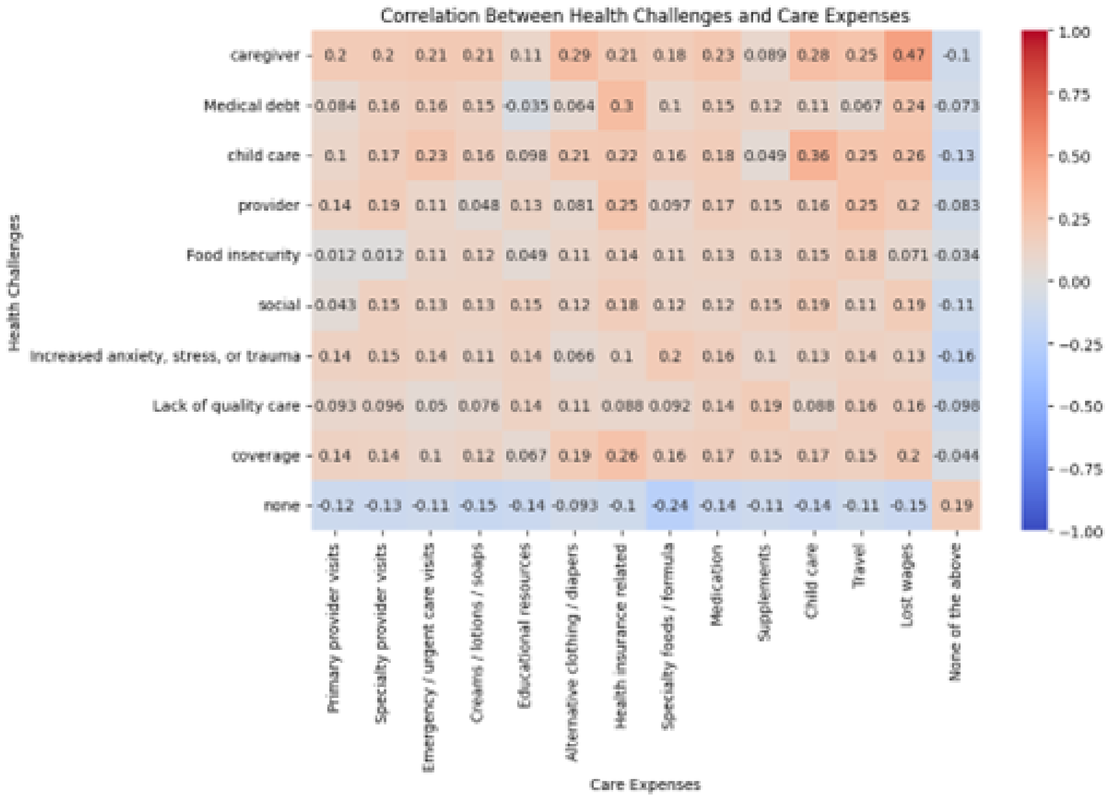

Figure 10. Correlation matrix between health-related disparities and increased expenses.

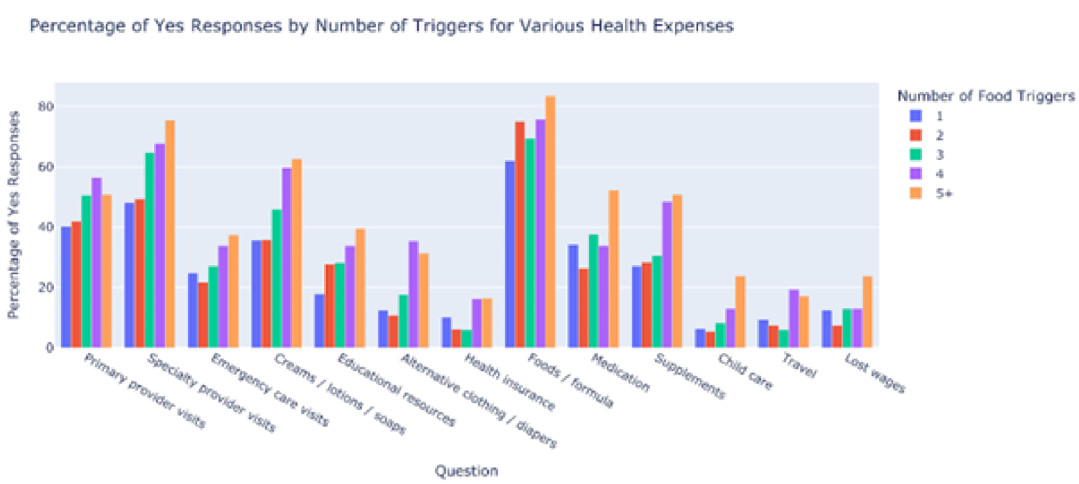

Figure 11. Bar plot of percentage of "Yes" responses to expenses questions grouped by number of food triggers.

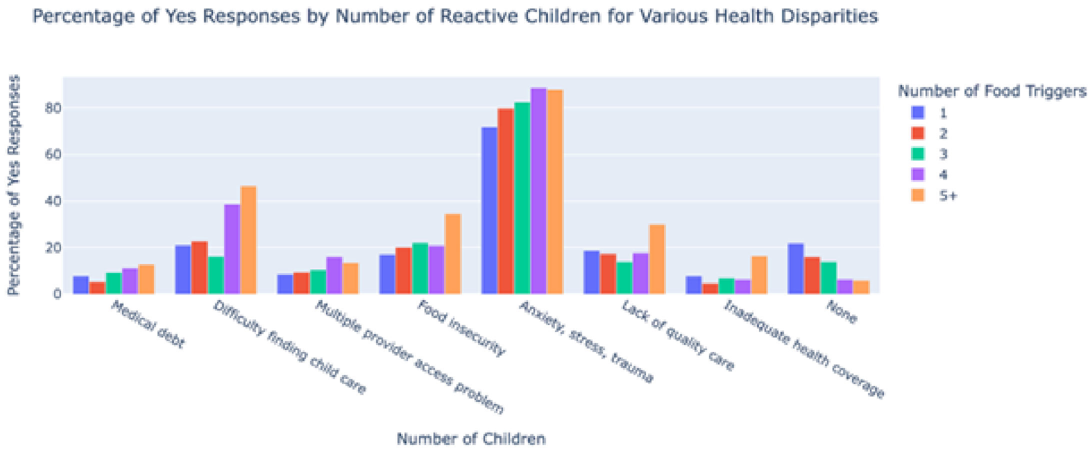

Figure 12. Bar plot of percentage of "Yes" responses to health disparities questions grouped by number of food triggers.

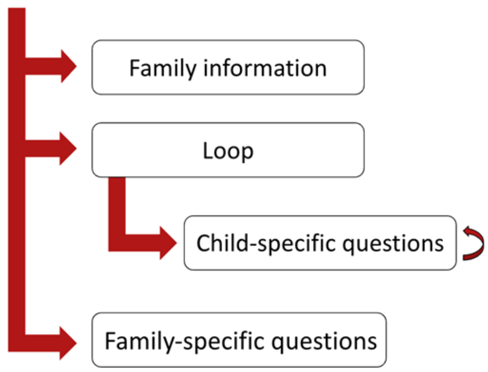

Figure 13. Diagram showing flow of redesigned survey.

downs (e.g. weeks, months, and years). Likewise, sensitivity levels are now selected using a drop-down menu. Figure 14 shows an example of a redesigned question. The redesigned survey and supporting documentation was provided to Free to Feed for future research.

CONCLUSION

This study on parentally-reported food reactivity in children highlights several critical insights into the prevalence, symptoms, and social implications of infant food allergies. Our analysis identified common symptoms, such as skin reactions, gastrointestinal distress, and respiratory issues, often triggered by common food items like dairy, nuts, and gluten. Furthermore, our examination of dietary modifications revealed that while some approaches, like elimination diets, showed promise as diagnostic tools, their effectiveness varied widely among individuals, emphasizing a need for personalized care.

Additionally, the social and economic ramifications of food reactivity underscore the far-reaching impact of this issue. Specifically, food allergies put a significant financial strain on families due to specialized treatments and clinical visits.

The implications of these findings are significant for families, healthcare professionals, and policymakers. Families must navigate the complexities of food reactivity with greater awareness and support, particularly in the context of dietary choices and social support. Healthcare providers should consider integrating more personalized dietary assessments and management plans to better address the needs of each child, and this individualized approach must begin in infancy.

For policymakers, the study highlights the need for more robust regulations and support systems to ensure that families have access to the necessary resources, including better insurance coverage for dietary interventions and child care.

Further research is essential to deepen our understanding of food reactivity in children. Longitudinal studies could provide valuable insights into the long-term outcomes of dietary interventions and the development of chronic conditions related to allergies. Exploring the genetic and environmental factors that contribute to food reactivity could pave the way for more targeted and effective treatments. Finally, future

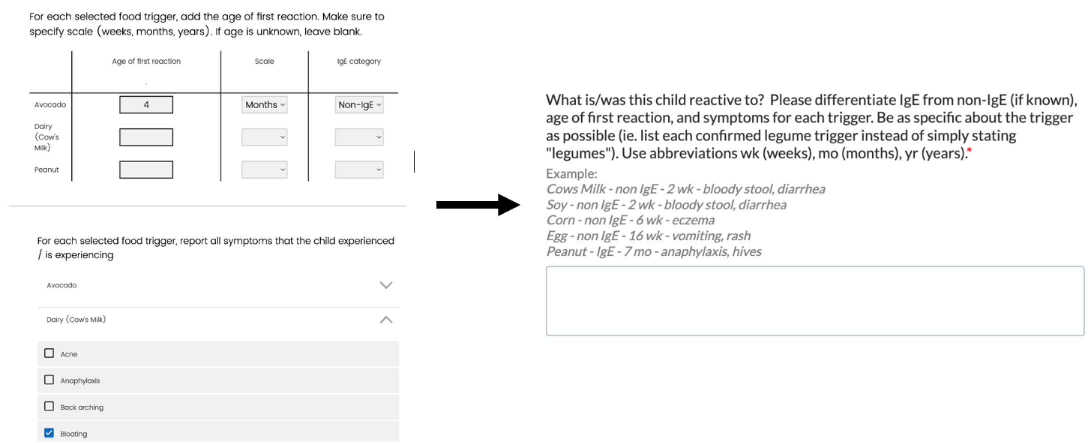

Figure 14. Example of redesigned survey question with original format on left and redesigned format on right.

research should continue to consider the broader social and economic impacts of food reactivity, with a focus on developing more support systems to guide families and help them navigate the health journeys of their children.

ACKNOWLEDGMENTS

This work was supported by the Stanford Data Science for Social Good summer program, which trains aspiring researchers to work on data science projects with social impact. Working closely with governments and nonprofits, participants take on real-world problems in education, health, energy, public safety, transportation, economic development, international development and more. Participants include a diverse and inclusive cohort of students who spend eight weeks of the summer working with Stanford researchers and technical mentors, learning insights from data that benefit society.

Free to Feed is a biotechnology company founded to help parents navigate food reactivity and provide the tools and knowledge necessary to support their infants with food allergies (https://www.freetofeed.com/). Dr. Trillitye Paullin the co-founder and CEO of Free to Feed. Dr. Paullin is also a Hoover Institution Veteran Fellow at Stanford.

We would like to extend a special thanks for Dr. Balasubramanian Narasimhan (Stanford University), Dr. Shilaan Alzahawi (Stanford University), Dr. Sarah Reyes, Cassandra Partridge, and the rest of the Stanford DSSG team.

REFERENCES

1. U.S. Census Bureau. Age and Sex Composition in the United States: 2022. Retrieved from https://www.census.gov/data/ tables/2022/demo/age-and-sex/2022-age-sex-composition.html on September 7, 2023.
2. Gupta RS, Warren CM, Smith BM, Blumenstock JA, Jiang J, Davis MM, Nadeau KC. The Public Health Impact of Parent-Reported Childhood Food Allergies in the United States. Pediatrics 2018; 142(6):e20181235.
3. OpenAI, et al. GPT-4 Technical Report. 2023; revised 2024. https://doi.org/10.48550/arXiv.2303.08774.

Amira Mahmedjan is an undergraduate at Stanford University, majoring in Computer Science, potentially on the Systems track. Amira's interests within Computer Science are broad: with experience in data structures and algorithms, machine learning, and app development, she is passionate about developing solutions tackling social problems and making software and technology an accessible tool for the world. At Stanford, she serves as a board member of Muslim Student Union and Central Asian Student Association, where she shares her identity and fosters a sense of belonging.

Sean Nesamoney is an undergraduate at Stanford University, majoring in Bioengineering with a minor in Creative Writing. Passionate about the intersection of computer science and medicine, Sean is particularly focused on how technology can advance precision care in psychiatry. At Stanford, he is actively involved in a cappella, The Stanford Daily, Students in Biodesign, and the Undergraduate Research Association. With a keen interest in biotechnology, Sean aims to make significant contributions in the subfield of early detection and diagnostic wearable technology.

Michael Howes is a PhD candidate in Statistics at Stanford University. He is an expert in Monte Carlo Significance Tests. Michael is a technical mentor Stanford Data Science for Social Good.

Annie K. Lamar received a BA in Classics and a BS in Computer Science from the University of Puget Sound in 2019. She also holds an MA in Education Data Science from the Stanford Graduate School of Education and a PhD in Classics from Stanford University. Her research is primarily focused on low-resource machine learning and computational linguistics. At Stanford, Annie was a Data Science Scholar and mentor for Data Science for Social Good. Annie is currently an Assistant Professor of Computational Classics at the University of California, Santa Barbara.