Measurement of Political Preferences

Presented at PolMeth Europe
May 14, 2026

Asya Magazinnik
Professor of Social Data Science, Hertie School

Measuring Political Preferences

Needs no long motivation: a central part of our work as political scientists.

- Democratic politics
    - Assessing representation & responsiveness
    - Predicting & explaining electoral outcomes
- Authoritarian politics/democratic backsliding: no less important
    - Autocrats have a notion of latent support $\rightarrow$ how much to repress, threat from opposition
    - Citizens' political participation depends on their perception of others
    - Backsliding: citizen support for political institutions

Plan for the Talk

A walk through my research on preference elicitation & measurement, with my brilliant collaborators:

Scott Abramson
University of Rochester

Lukas Stoetzer
University of Witten/Herdecke

Korhan Kocak
IE Madrid

Anton Strezhnev
University of Wisconsin-Madison

Structured by an argument about the broader literature:

- The problems methodologists have to solve with our designs
- How we approach the problem (as a causal vs. a descriptive exercise)
- An illustration of the divergence between intuitive causal and descriptive estimands
- An example of each type of work from my own research
- Virtues, challenges, and promising ways forward on each side

Defining the Problem

Our Five Basic Challenges

1. Preferences are latent. We can only observe:
   - Responses to questions *about* preferences
   - Choices
2. Preferences are multidimensional.
   - E.g. candidates for office, immigrants, tax plans...
   - They are interdependent (not separable)
3. We need meaningful estimands to summarize their complexity.
   - As social scientists, we seek answers to questions like:

     Do voters support or oppose redistribution?

     Do women or racial minorities face electoral bias?
   - This requires *aggregating* over choice tasks, respondents, other attributes.
4. Preferences are heterogeneous in both direction and intensity.
   - *Majorities* determine electoral winners
   - But intensity of preference drives turnout, feature relevance, donations, protest, lobbying, interest group mobilization...
5. Survey subjects impose limitations.
   5.1 Limited time/attention: satisficing, speeding, low quality data
   5.2 Social desirability, preference falsification, demand effects, cheap talk...

The Purpose of the Exercise

Are we engaging in a causal or a descriptive exercise?

The Descriptive Exercise

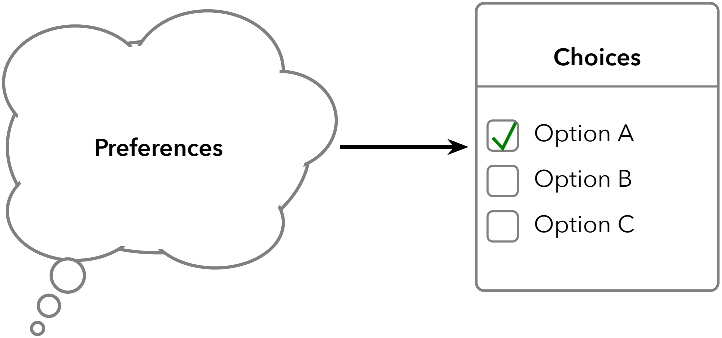

What Do Preferences Look Like?

Figure 1: A rank-ordering of possible alternatives.

What Do Preferences Look Like?

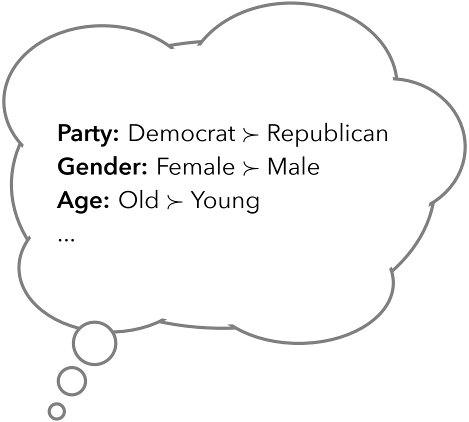

Figure 2: A preference ordering over features + importance ordering over attributes.

What Do Preferences Look Like?

$$
\begin{aligned}
U_{ij} = {} & \beta_{i,\text{party}} \mathbf{1}(\text{Party}_j = \text{Democrat}) \\
& + \beta_{i,\text{gender}} \mathbf{1}(\text{gender}_j = \text{female}) \\
& + \beta_{i,\text{age}} \mathbf{1}(\text{Age}_j = \text{old}) + \varepsilon_{ij}
\end{aligned}
$$

Figure 3: A utility function.

The Descriptive Exercise

The task is to recover the latent preferences from some observed choices.

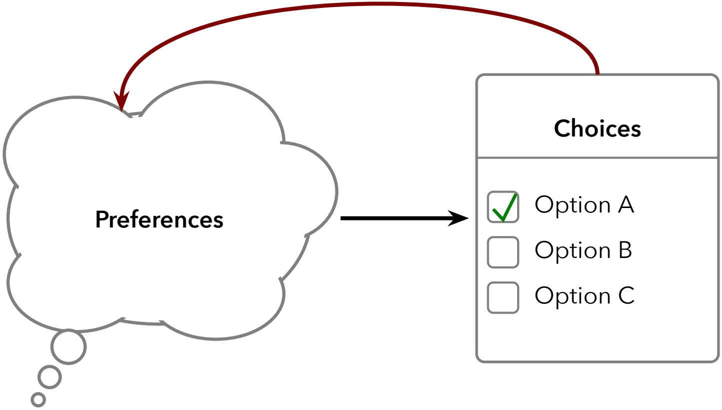

This is not what the applied literature is doing most of the time!

The Causal Exercise

Hainmueller, Hopkins, and Yamamoto (2014) introduces the average marginal component effect (AMCE) as a causal estimand:

> $AMCE(M, F) :=$ *the average change in the probability of choosing a profile with feature M vs. F, holding fixed the rest of this profile and the other profile.*

TABLE 2 Preferences over Candidate Profiles

Rank	V1	V2	V3	V4	V5
1	MR	MR	MR	FD	FD
2	FR	FR	FR	FR	FR
3	MD	MD	MD	MD	MD
4	FD	FD	FD	MR	MR

Note: This table presents preferences over profiles constructed from preferences over attributes.

FIGURE 1 Obtaining the AMCE

$\bar{Y}(MR, MD)$	$-$	$\bar{Y}(FR, MD)$	$=$	-2
$\bar{Y}(MR, FD)$	$-$	$\bar{Y}(FR, FD)$	$=$	0
$\bar{Y}(MR, MR)$	$-$	$\bar{Y}(FR, MR)$	$=$	$1/2$
$\bar{Y}(MR, FR)$	$-$	$\bar{Y}(FR, FR)$	$=$	$1/2$
$\bar{Y}(MD, MD)$	$-$	$\bar{Y}(FD, MD)$	$=$	$1/2$
$\bar{Y}(MD, FD)$	$-$	$\bar{Y}(FD, FD)$	$=$	$1/2$
$\bar{Y}(MD, MR)$	$-$	$\bar{Y}(FD, MR)$	$=$	0
$\bar{Y}(MD, FR)$	$-$	$\bar{Y}(FD, FR)$	$=$	-2
				-2
(# of profiles) $\times$ (# of voters) $\times$(# of profiles with a given gender)			$=$	40
AMCE			$=$	$-1/20$

Notes: Computing the AMCE for male over female candidates based on the aggregate preferences summarized in Table 3. $\bar{Y}(C_1, C_2)$ denotes the number of votes candidate C_1 obtains when running against candidate C_2.

The Causal vs. Descriptive Exercises

Abramson, Scott, Korhan Kocak, and Asya Magazinnik. 2022. "What Do We Learn About Voter Preferences from Conjoint Experiments?" *American Journal of Political Science* 66(4): 1008-1020.

These are fundamentally different questions.

- The descriptive question asks: *what are the parameters of the latent preference structure?*
- The causal question asks: *what is the effect on observed choices of a particular manipulation on the choice set?*

And unsurprisingly, they yield different answers.

- We show that an AMCE in favor of female over male *does not imply*:
  - A majority of voters prefer female to male candidates
  - Female candidates beat male candidates in the majority of possible head-to-head electoral contests
- It can, with some additional assumptions, recover an average ideal point

Related Literature

- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1): 1-30.
- Abramson, Scott, Korhan Kocak, and Asya Magazinnik. 2022. "What Do We Learn About Voter Preferences from Conjoint Experiments?" *American Journal of Political Science* 66(4): 1008-1020.
- Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2023. "Using Conjoint Experiments to Analyze Election Outcomes: The Essential Role of the Average Marginal Component Effect." *Political Analysis* 31(4): 500-518.
- Abramson, Scott, Korhan Kocak, Asya Magazinnik, and Anton Strezhnev. 2026. "Aggregation, Interpretation, and Estimation of Preferences in Conjoint Experiments." R&R, *Political Science Research and Methods*. Working paper.

Reflection

In Abramson, Kocak, and Magazinnik (2022), we wanted to build a bridge between causal estimands and descriptive parameters of interest (TIEM)

- My view today is that these are two branching paths
- Different questions, different advantages and disadvantages
- Both have exciting research frontiers!
- But we would be well-served as a field to be precise about which one we're doing (and why), and to keep the two distinct

Divergence Between Causal and Descriptive Estimands

The Audit Study

Research Question: Is there a bias against female job applicants in the labor market?

Research Design:

- Construct identical resumes except for the gender of the applicant and randomly send the female version to 50% of employers and the male version to 50% of employers
- **Causal Estimand:** Difference in proportion of sent resumes that get a callback between employers in the male condition and employers in the female condition

(see, e.g., Bertrand and Mullainathan 2003 and many others)

Stephen Nichols

nichols.stephenw@gmail.com

Education: Wesleyan University, B.A. in Computer Science, 2017 G.P.A. 3.6

Languages:	C++, C#, Java, Python, Ruby, JavaScript, TypeScript, Go
Font-End:	React, Unity, Node, Swift, PHP, HTML/CSS
Cloud Computing:	Google Cloud Platform, AWS
Infrastructure/Scalability:	Docker, Kubernetes, MongoDB, MySQL, jQuery

EXPERIENCE

Blue Slate Solutions Albany, NY
Software Engineer 2018-Present

- Built a platform to analyze 50K+ Skype call recordings daily using Spark
- Led a 5-person team to create an interactive KPI platform to build & share dashboards
- Developed a cloud app to review performance of IT services daily
- Trained new hires in agile methodology, documentation protocols, & Git procedures
- Built RESTful API endpoints and tests, using Java and SAP HANA SQL database

Astro Technology (acquired by Slack) Palo Alto, CA
Software Development Engineer 2017-2018

- Integrated customer metrics data flow and improved data visualization protocols to give product and marketing teams weekly consumer experience summaries
- Started "Bug Busters," led bi-weekly code review sessions, fixing 100s of bugs
- Deployed TensorFlow jobs to convert Kafka streams of email & voice messages to dynamic calendar events

eBay San Jose, CA
Software Engineer Intern June-August, 2016

- Created a scoring system for advertising unit to evaluate advertiser goals
- Collaborated with data scientists to improve insight reports for 1M+ clients

PROJECTS

Stephanie Nichols

nichols.stephenw@gmail.com

Education: Wesleyan University, B.A. in Computer Science, 2017 G.P.A. 3.6

Languages:	C++, C#, Java, Python, Ruby, JavaScript, TypeScript, Go
Font-End:	React, Unity, Node, Swift, PHP, HTML/CSS
Cloud Computing:	Google Cloud Platform, AWS
Infrastructure/Scalability:	Docker, Kubernetes, MongoDB, MySQL, jQuery

EXPERIENCE

Blue Slate Solutions Albany, NY
Software Engineer 2018-Present

- Built a platform to analyze 50K+ Skype call recordings daily using Spark
- Led a 5-person team to create an interactive KPI platform to build & share dashboards
- Developed a cloud app to review performance of IT services daily
- Trained new hires in agile methodology, documentation protocols, & Git procedures
- Built RESTful API endpoints and tests, using Java and SAP HANA SQL database

Astro Technology (acquired by Slack) Palo Alto, CA
Software Development Engineer 2017-2018

- Integrated customer metrics data flow and improved data visualization protocols to give product and marketing teams weekly consumer experience summaries
- Started "Bug Busters," led bi-weekly code review sessions, fixing 100s of bugs
- Deployed TensorFlow jobs to convert Kafka streams of email & voice messages to dynamic calendar events

eBay San Jose, CA
Software Engineer Intern June-August, 2016

- Created a scoring system for advertising unit to evaluate advertiser goals
- Collaborated with data scientists to improve insight reports for 1M+ clients

PROJECTS

The Audit Study

Suppose there are three employer types, each constituting 1/3 of the employer population.

Rank	Type 1	Type 2	Type 3
1	Male Candidate	Female Candidate	Outside Option
2	Female Candidate	Outside Option	Male Candidate
3	Outside Option	Male Candidate	Female Candidate

Table 1: Preference Orderings for Three Employer Types

What proportion of employers prefer male to female candidates?

$$\frac{1}{3} + \frac{1}{3} = \frac{2}{3}$$

And what is the causal estimate from our experiment?

$$\frac{1}{3} - \frac{2}{3} = -\frac{1}{3}$$

An Unusual Application of Causal Inference

Let's back up: why did we do causal inference here?

- Because that's what we always do...?
- To solve the fundamental problem of causal inference?

$$ITE_i = Y_i(T_i = 1) - Y_i(T_i = 0)$$

  It doesn't exist in *quite* the usual sense (though there are relevant respondent-imposed constraints)

- Because this causal quantity is inherently interesting?
- To measure a preference?

  Preferences are fixed and latent; the *choice set* is manipulated by the researcher:

$$ITE_i = Y_i(\text{Preferences}_i, T_i = 1) - Y_i(\text{Preferences}_i, T_i = 0)$$

Taking the Descriptive Road

Step 1: What Do Preferences Look Like?

Let candidates be fully described by three binary attributes:

$$\text{Gender} \in \{\text{M}, \text{F}\},\ \text{Party} \in \{\text{D}, \text{R}\},\ \text{and Age} \in \{\text{Y}, \text{O}\}$$

We assume only that our respondents are able to compare every possible candidate pairing and form a preference (completeness):

$$\begin{aligned}
&\text{FRY} \succ \text{FDY}\\
&\text{FRY} \succ \text{MDY}\\
&\text{FRY} \succ \text{FDO}\\
&\text{FRY} \succ \text{MDO}\\
&\text{FRY} \prec \text{MRY}\\
&\ldots
\end{aligned}$$

Step 2: How to Summarize These Preferences?

FRY $\succ$ FDY

FRY $\succ$ MDY

FRY $\succ$ FDO

FRY $\succ$ MDO

FRY $\prec$ MRY

...

We define the Feature Choice Probability (FCP) of M vs. F as: *the proportion of tasks in which M is chosen over F out of all possible tasks where the two are directly compared*.

A Nonparametric Estimand: The AFCP

Abramson, Scott, Korhan Kocak, Asya Magazinnik, and Anton Strezhnev. 2026. "Aggregation, Interpretation, and Estimation of Preferences in Conjoint Experiments." R&R, *Political Science Research and Methods*. Working paper.

In this paper, we:

- Define the AFCP estimand and an unbiased estimator
- Argue for the AFCP as the most intuitive, direct measure of preference (in a nonparametric world)
- Decompose other estimands in the literature as a function of AFCPs
  - The AMCE (HHY 2014), the ACP (Ganter 2023), and the marginal mean (Leeper, Hobolt, and Tilley 2020)
- Show that the AMCE uses both direct and indirect AFCPs, thus increasing statistical power while chipping away at interpretability as a preference

Using Direct and Indirect Comparisons

$$\text{AMCE}(\text{Ukraine}, \text{Syria}) =$$

$$f\begin{pmatrix} \text{AFCP}(\text{Ukraine}, \text{Syria}) & \text{AFCP}(\text{Ukraine}, \text{Libya}) & \text{AFCP}(\text{Syria}, \text{Libya}) \\ & \text{AFCP}(\text{Ukraine}, \text{Spain}) & \text{AFCP}(\text{Syria}, \text{Spain}) \\ & \text{AFCP}(\text{Ukraine}, \text{Greece}) & \text{AFCP}(\text{Syria}, \text{Greece}) \end{pmatrix}$$

Comparing the AMCE to the AFCP

Problem	AMCE	AFCP
1. Preferences are latent	"Model-free"	"Model-free"
2. Preferences are multidimensional	✓	✓
3. Meaningful estimand	😐	average preference
4. Heterogeneity in direction and intensity	×	×
5a. Survey subjects: sample limitations	Direct + indirect comparisons	Only direct comparisons
5b. Survey subjects: social desirability bias	Decently*	Decently*

*See Horiuchi, Markovich, and Yamamoto (2022); avenue for future research.

A Structural Approach

Abramson, Scott, Korhan Kocak, and Asya Magazinnik. "Ideology, Electability, and Valence." Working paper available upon request.

- An application to a Democratic primary election in Florida to challenge the Republican incumbent Senator, Rick Scott
- Sample: Florida likely Democratic voters
- We want to learn how party loyalists trade off:
  1. Their preferred ideology
  2. Personal attributes, e.g. gender, race, and political experience
  3. Electability (expected performance against the incumbent)

Model Setup

- Unidimensional policy space, voter i's ideal point is $x_i \in \mathbb{R}$
- Let $d_{ij} = |x_i - x_j|$ denote ideological distance between voter i and candidate j
- Candidate j also has personal attributes $v_j \in \mathbb{R}^\tau$ (e.g. gender, race, experience)
- Voter i assigns weights r_i to ideology and k_i to personal attributes
- Unobserved error term ε_{ij}
- Then, voter i's payoff from having candidate j in office is given by:

$$u_{ij} = -d_{ij} r_i + v_j k_i + \varepsilon_{ij}$$

Instrumental vs. Expressive Concerns

Instrumental: by voting for j, voter i increases the probability j wins

But voters also derive expressive utility from voting for a candidate they like:

- Warm glow,
- Signaling.

We can build in instrumental (r_i, k_i) and expressive (ρ_i, κ_i) concerns over ideology and personal attributes:

- Let q_{jc} denote probability primary candidate j wins general election against the other party's challenger, c
- Then voter i's payoff from voting for primary candidate j is given by:

$$u_{ij} = \underbrace{-d_{ij}(q_{jc} r_i + \rho_i)}_{\text{policy term, } j} + \underbrace{v_j(q_{jc} k_i + \kappa_i)}_{\text{personal attributes, } j} + \underbrace{(1 - q_{jc})(-d_{ic} r_i + v_c k_i)}_{\text{challenger, } c} + \varepsilon_{ij}$$

Estimating Equation

Assume $\varepsilon_{ij} \sim$ Type I Extreme Value.

Then, the log odds ratio for voter i's probability of voting for some primary candidate j over another primary candidate j' is:

$$\ln\left(\frac{p_{ij}}{p_{ij'}}\right) = -r_i\left(q_{jc}(d_{ij} - d_{ic}) - q_{j'c}(d_{ij'} - d_{ic})\right) - \rho_i(d_{ij} - d_{ij'}) + \\ b_i(q_{jc} - q_{j'c}) + k_i\left(q_{jc}(v_j - v_c) - q_{j'c}(v_{j'} - v_c)\right) + \kappa_i(v_j - v_{j'}).$$

Estimating Equation

Assume $\epsilon_{ij} \sim$ Type I Extreme Value.

Then, the log odds ratio for voter i's probability of voting for j over j' is:

$$\ln\left(\frac{p_{ij}}{p_{ij'}}\right) = -r_i\left(q_{jc}(d_{ij} - d_{ic}) - q_{j'c}(d_{ij'} - d_{ic})\right) - \rho_i(d_{ij} - d_{ij'}) + \\ b_i(q_{jc} - q_{j'c}) + k_i\left(q_{jc}(v_j - v_c) - q_{j'c}(v_{j'} - v_c)\right) + \kappa_i(v_j - v_{j'}).$$

Now, we can create an intentional randomized survey design where:

- We observe p, q, d, and v

Estimating Equation

Assume $\epsilon_{ij} \sim$ Type I Extreme Value.

Then, the log odds ratio for voter i's probability of voting for j over j' is:

$$\ln\left(\frac{p_{ij}}{p_{ij'}}\right) = -r_i\left(q_{jc}(d_{ij} - d_{ic}) - q_{j'c}(d_{ij'} - d_{ic})\right) - \rho_i(d_{ij} - d_{ij'}) + b_i(q_{jc} - q_{j'c}) + k_i\left(q_{jc}(v_j - v_c) - q_{j'c}(v_{j'} - v_c)\right) + \kappa_i(v_j - v_{j'}).$$

Now, we can create an intentional randomized survey design where:

- We observe p, q, d, and v
- and we estimate r, ρ, b, k, and κ (via OLS)

Quantities of Interest

Willingness-to-pay for personal attribute v_j in terms of ideology:

$$WTP_i^d(v_j) = \frac{q_{jc}k_i + \kappa_i}{q_{jc}r_i + \rho_i}.$$

Willingness-to-pay for personal attribute v_j in terms of electability:

$$WTP_i^q(v_j) = \frac{q_{jc}k_i + \kappa_i}{r_i(d_{ij} - d_{ic}) - (k_i + b_i)}.$$

(See Wiswall and Zafar 2018)

We estimate these quantities and calculate measures of uncertainty using block bootstrap.

Survey Design

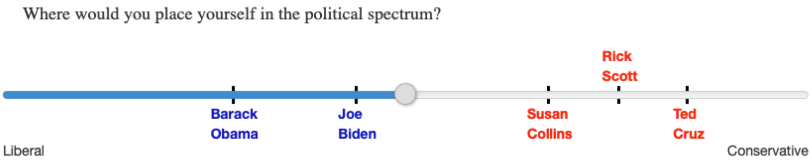

This gives us our x_i for every survey respondent.

We randomize q_{jc} (probability of beating c), and three personal attributes: gender, race, and political experience. Respondents give us p_{ij} for every randomized candidate.

For each candidate below, you are given estimates of their ideology on the political spectrum, the probability they will win in the general election, their gender, race, and their experience in politics. Please indicate your probability of voting for each of the candidates.

Scenario 1

	A	B	C	D	Rick Scott
Probability of beating Rick Scott	57%	32%	51%	35%	NA
Gender	Female	Male	Female	Male	Male
Race	White	White	White	Black	White
Years in politics	5	29	15	11	10

This is how your ideology compares with the candidates:

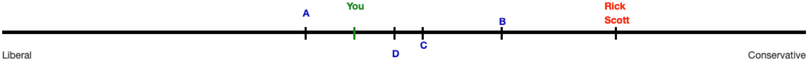

Question 1: What is the probability you would vote for each Democratic primary candidate?

Candidate A	0	25	100
Candidate B	0	25	100
Candidate C	0	25	100
Candidate D	0	25	100

‹ Previous › Next

Results: WTP in Ideology for Candidate Features

	Median	1st quartile	3rd quartile	Prop. who prefer
Female vs. male	0.02	-0.14	0.18	0.54
	[0.01, 0.02]	[-0.16, -0.12]	[0.16, 0.20]	[0.52, 0.56]
Black vs. white	0.01	-0.25	0.26	0.51
	[-0.00, 0.02]	[-0.29, -0.22]	[0.22, 0.30]	[0.49, 0.53]
Latinx vs. white	0.01	-0.22	0.22	0.52
	[-0.00, 0.02]	[-0.25, -0.18]	[0.20, 0.24]	[0.50, 0.54]
Asian vs. white	0.02	-0.29	0.32	0.48
	[0.00, 0.03]	[-0.33, -0.25]	[0.29, 0.36]	[0.46, 0.50]
Experience	-0.00	-0.01	0.01	0.50
	[-0.00, 0.00]	[-0.01, -0.01]	[0.01, 0.01]	[0.48, 0.51]

Challenges for Descriptive Estimands

1. Statistical power
   - The estimator for the AFCP has less data to work with than the AMCE estimator (unless you design your randomization scheme otherwise)
   - Similarly, highly parameterized structural models will have many parameters to estimate (especially if you want them at the individual level)

2. (For the structural route) model dependence
   - You have to believe your general model of the world
   - ...as well as its particular functional form

Future Directions: Structural Estimation + ML

- Acharya, Hainmueller, and Xu (2026)

  "We propose a structural approach that embeds a deep neural network within a random utility logit model, allowing preference parameters to vary as a fully flexible function of respondent characteristics."
  "We think of this approach as lean structure -- the upside of a structural model without its usual parametric downside."

- Already an active research area in economics and marketing

  Farrell, Max H., Tengyuan Liang, and Sanjog Misra. 2025. "Deep Learning for Individual Heterogeneity." *arXiv preprint*: https://arxiv.org/pdf/2010.14694.
  Wei, Yanhao and Zhenling Jiang. 2025. "Estimating Parameters of Structural Models Using Neural Networks." *Marketing Science* 44(1): 102-128.

- Matrix completion to fill in respondents' full preference orderings

Taking the Causal Road

Measuring Attribute Importance

Stoetzer, Lukas and Asya Magazinnik. 2026. "Measuring Attribute Relevance in Conjoint Analysis." Working paper available upon request.

Motivation:

- We showed in Abramson, Kocak, and Magazinnik (2022) that the AMCE conflates the direction and intensity of preferences
- In general, our usual conjoint estimands are poorly suited to measure the importance or relevance of an attribute for decisionmaking
- We build on the literature on causal attribution (Pearl 1999; Yamamoto 2012; Rosenbaum 2002) to define a causal estimand that specifically measures attribute importance

Our Proposed Estimand: The APNS

		Counterfactual outcome $Y_i(t_p)$	
		Select (1)	Not select (0)
$Y_i(t_q)$	Select (1)	S, not N	**N & S**
	Not select (0)	**N & S**	N, not S

The Probability of Necessary and Sufficient Causation (PNS) captures the joint probability that a respondent would select the profile with t_q but not with t_p, or vice versa (green cells).

$$PNS(t_q, t_p, t_{[-l]}, \mathbf{t}_{[-j]}) = P\left(Y_i(t_q, t_{[-l]}, \mathbf{t}_{[-j]}) = 1,\ Y_i(t_p, t_{[-l]}, \mathbf{t}_{[-j]}) = 0\right)$$

The APNS takes the expectation of the PNS over the joint distribution of other attribute levels and other profiles in the experiment.

Advantages of the APNS

- Directly interpretable (as a probability between 0 and 1)
- Can be applied to standard conjoint data
- Causal inference framework serves a substantive purpose: focus on causes of effects helps pin down relevance

Wrapping Up

My Approach to Preference Elicitation

1. What is the goal? A particular causal question? A descriptive task?
2. What is the assumed preference structure?
   - Nonparametric
   - More structure: e.g., single-peakedness, separability
   - A fully structural model
3. What is the estimand?
   - How does it aggregate preferences in asymptopia?
   - Estimation concerns
4. How am I handling the usual five challenges?
   - Weaknesses on some dimensions are inevitable. What is the use case?