

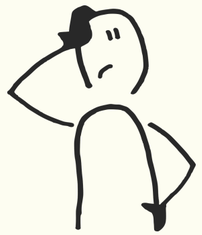
PORTIA

An LLM for intelligent portfolio optimization



Presented by :

Lindsay Carlin, Jacob Dentes, Salvador Dzimah, Julia Kim aka MVP (Minimum Viable Product) Team



Problem

- Investing assets is **crucial** in preparing for retirement
- However, most people don't know where to start:
  - → Little market expertise
  - → Little time to learn about strategies
  - → Nervous to take on financial risk



Idea

OR offers a **rich** suite of models well-suited to PO, and **LLMs** can **bridge** the **gap** between these highly quantitative models and the everyday user.



Our solution

We built PortIA, an **LLM-powered tool** that leverages an **optimization model** to allow everyday users of all demographics to **build a custom financial portfolio** according to their preferences.



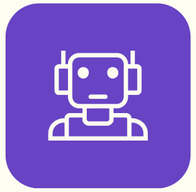
PortIA’s pipeline: the decision-making





User preference intake

User records their desires, such as goals, industries to avoid and risk caps, in natural language.



Data pre-processing

PortIA’s Modeling Agent reads user preferences and translates each preference into a field in a structured form. Unspecified preferences become conservative default values.



Preparation

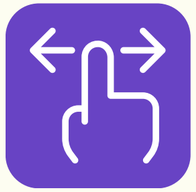
In the background, PortIA pre-computes the expected returns and historical covariance matrix from the market data to speed up the optimization step.

PortIA’s pipeline: the pre-audit



Preference validation

Upon user confirmation of preferences, PortIA checks whether (1) all user-specified constraints were accounted for, and (2) any constraints were invented.



Constraint validation

PortIA also checks the validity of constraints (e.g., “at most 20% technology exposure” cannot exist with “at least 30% technology exposure”).



Modeling fixes

In the event of a translation error or an invalid constraint, the Modeling Agent attempts to reformulate the preferences correctly. If needed, PortIA goes back to the user to clarify their preferences.

PortIA’s pipeline: the decision-making

User confirmation

The Chatting Agent resolves ambiguities with the user, who ultimately confirms their preferences.

Problem formulation

PortIA constructs and solves the optimization problem using Bertsimas and Cory-Wright (2022).

Feasibility evaluation

If the problem is infeasible, an IIS analyzer identifies conflicting constraints and proposes specific relaxations for the user to accept or reject. PortIA re-optimizes up to three times.

Bertsimas D, Cory-Wright R (2022) A scalable algorithm for sparse portfolio selection. INFORMS Journal on Computing 34(3):1489–1511, ISSN 1526-5528, URL http://dx.doi.org/10.1287/ijoc.2021.1127.

PortIA’s pipeline: the post-audit



Solution validation

After the solve, PortIA checks that standard constraints were not violated, e.g. that weights add to 1 and the number of holdings does not exceed the stated maximum.



Preference comparison

PortIA confirms that the result portfolio correctly accounts for user preferences without specifying additional unstated preferences.



Alternative portfolio generation

PortIA selects constraints to relax and re-optimizes, then validates these alternative portfolios, which are to be displayed to the user as additional choices.

Main challenge #1: dropping unsupported constraints

User: “I don’t want any fossil fuel companies in my portfolio.”

Issue: *Fossil fuel* is not a valid ticker classification.



PortIA: Ignored the constraint instead of returning to the user for clarification.

THE FIX: PortIA acknowledges unrecognized constraints to the user and suggests supported alternatives or asks for clarification.

Main challenge #2: not recognizing infeasibility

User: “I’d like all Basic Materials stocks excluded, and at least 10% exposure to Basic Materials ETFs.”

Issue: This infeasible request went undetected.



PortIA: Spent the solving period checking infeasible portfolios, then returned no portfolio when the solver inevitably times out.

THE FIX: Return a valid provisional portfolio when one exists, or detect impossible constraints immediately, advise users to relax their preferences, and explain timeouts.

Main challenge #3: sensitivity to qualitative preferences

User: “I am moderately risk-averse.”

Issue: This statement does not explicitly provide a risk constraint for PortIA.



PortIA: Produced a different preference object 36% of the time when sent 100 qualitative preferences, twice each.

THE FIX: prompt engineering with examples of how to quantify common expressions

Main challenge #4: pre-audit results exposed to user

User: “Keep Technology exposure below 20%.”

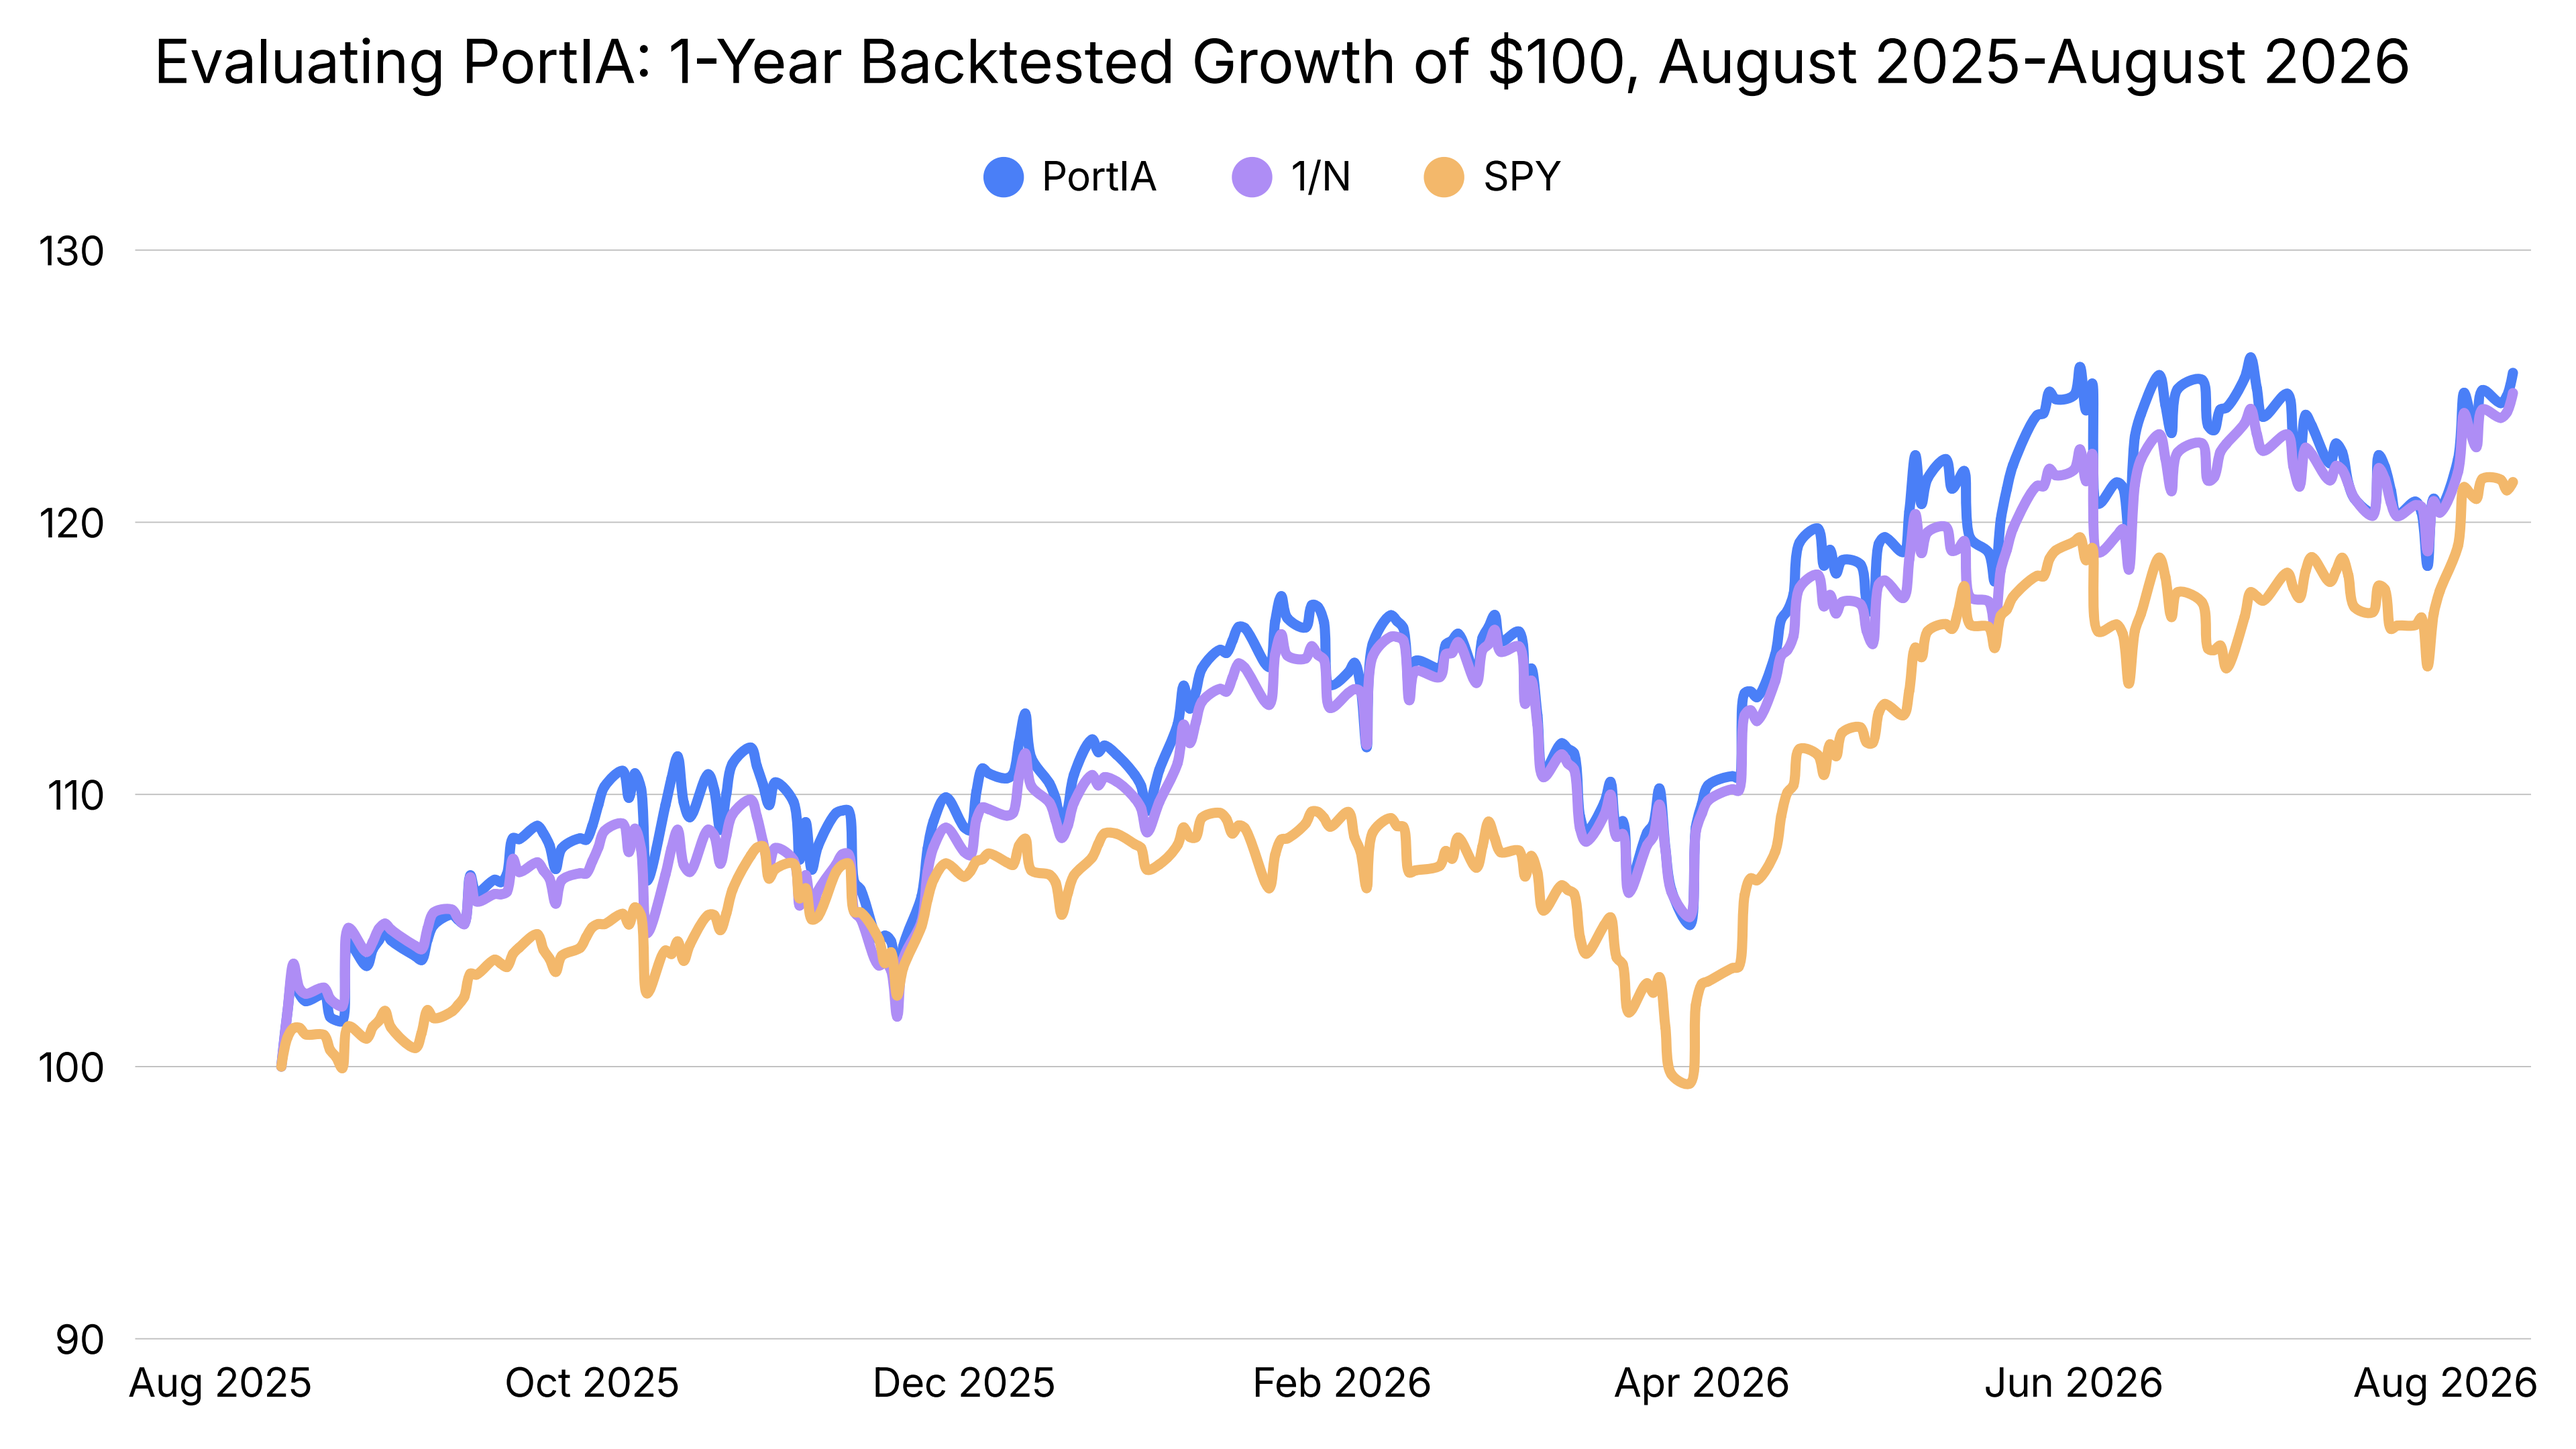
Issue: if the Modeling Agent could not translate this limit, PortIA would return an error message to the user instead of re-attempting to parse the preference.



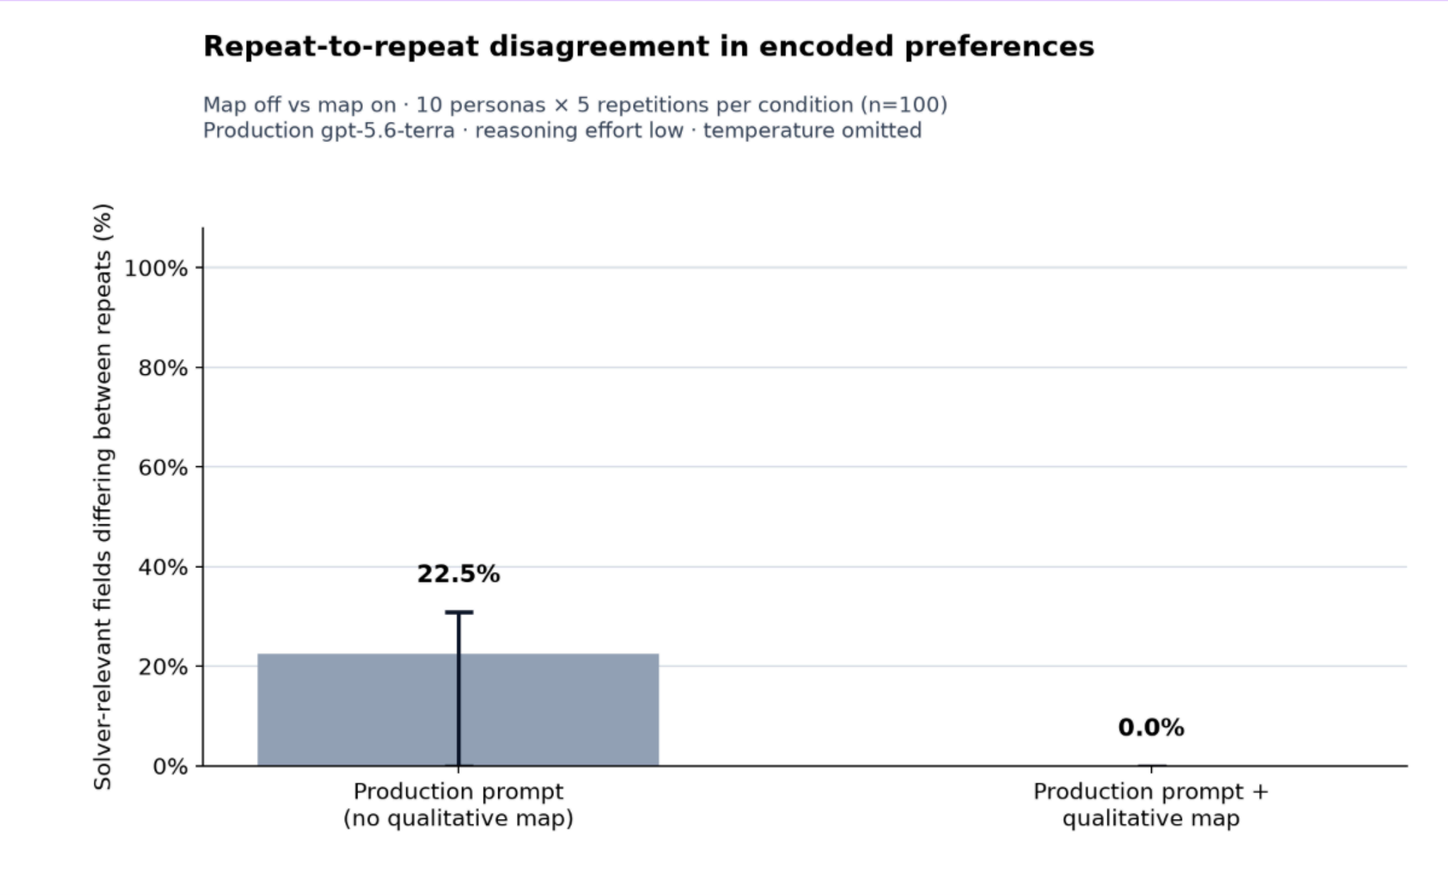
PortIA: allowed the user to confirm their preferences *before* catching potential translation errors.

THE FIX: the Modeling Agent may now re-attempt a given translation up to three times before returning to the user for clarification.

Evaluating PortIA: 1-Year Backtested Growth of $100, August 2025-August 2026

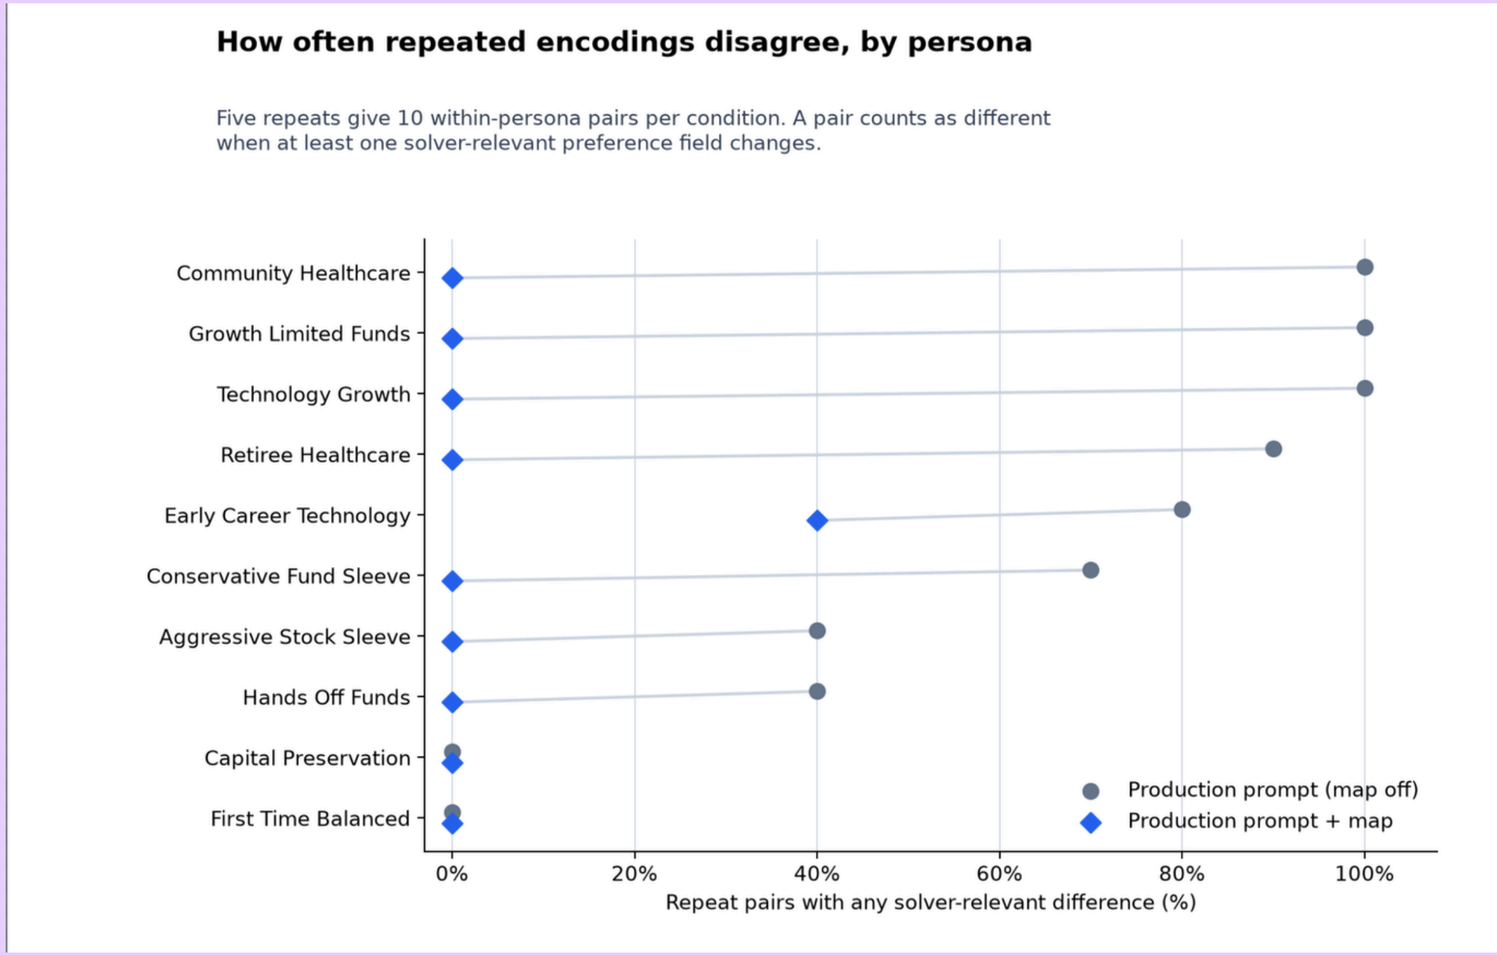


Experiment #1: how does use of a prompt map affect portfolio construction?



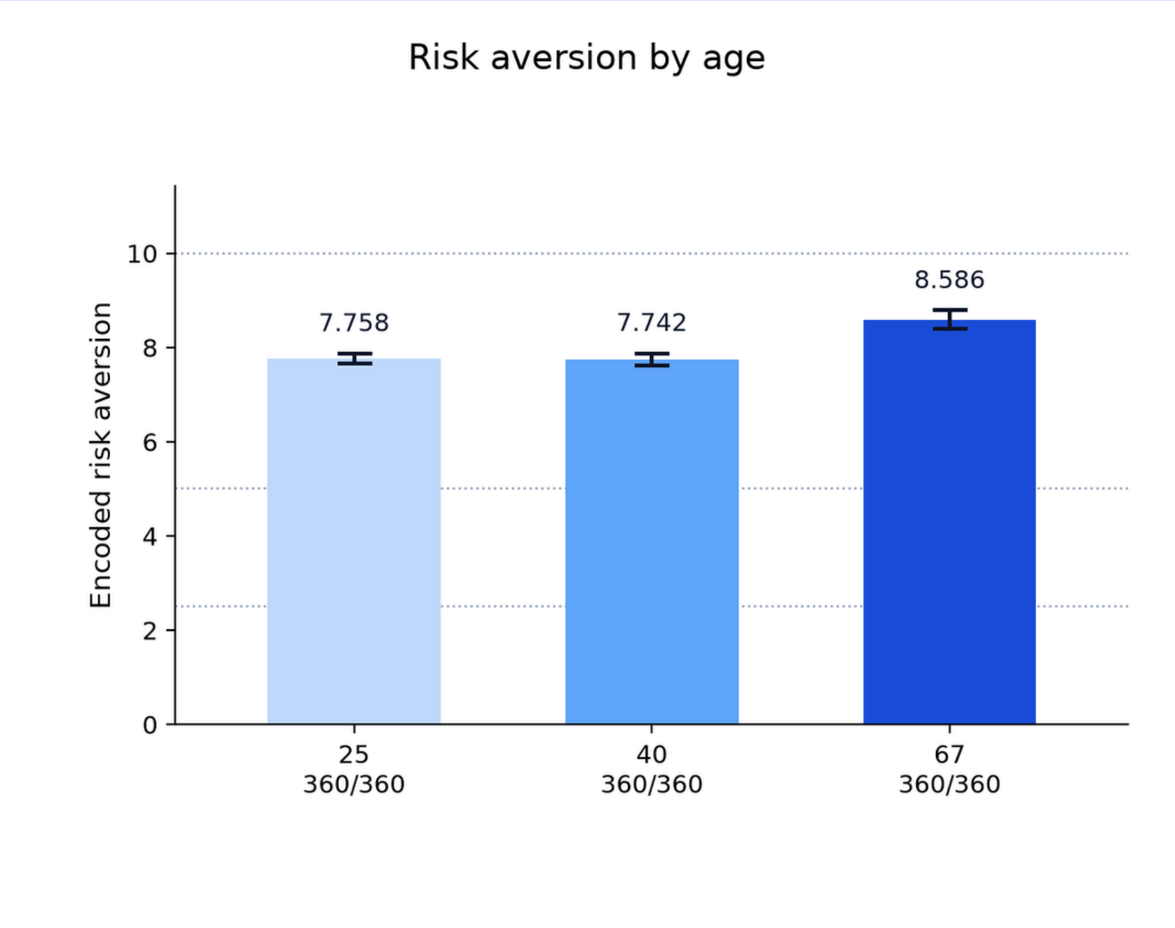
Without a prompt map, a pair of repeated encodings disagreed on an average of 22.5% of solver-relevant preference fields. Adding a prompt map made the median difference 0%.

Experiment #1: how does use of a prompt map affect portfolio construction?



Each of the ten personas we tested has at least as much consistency with the prompt map as without. In some cases, the difference is 100%.

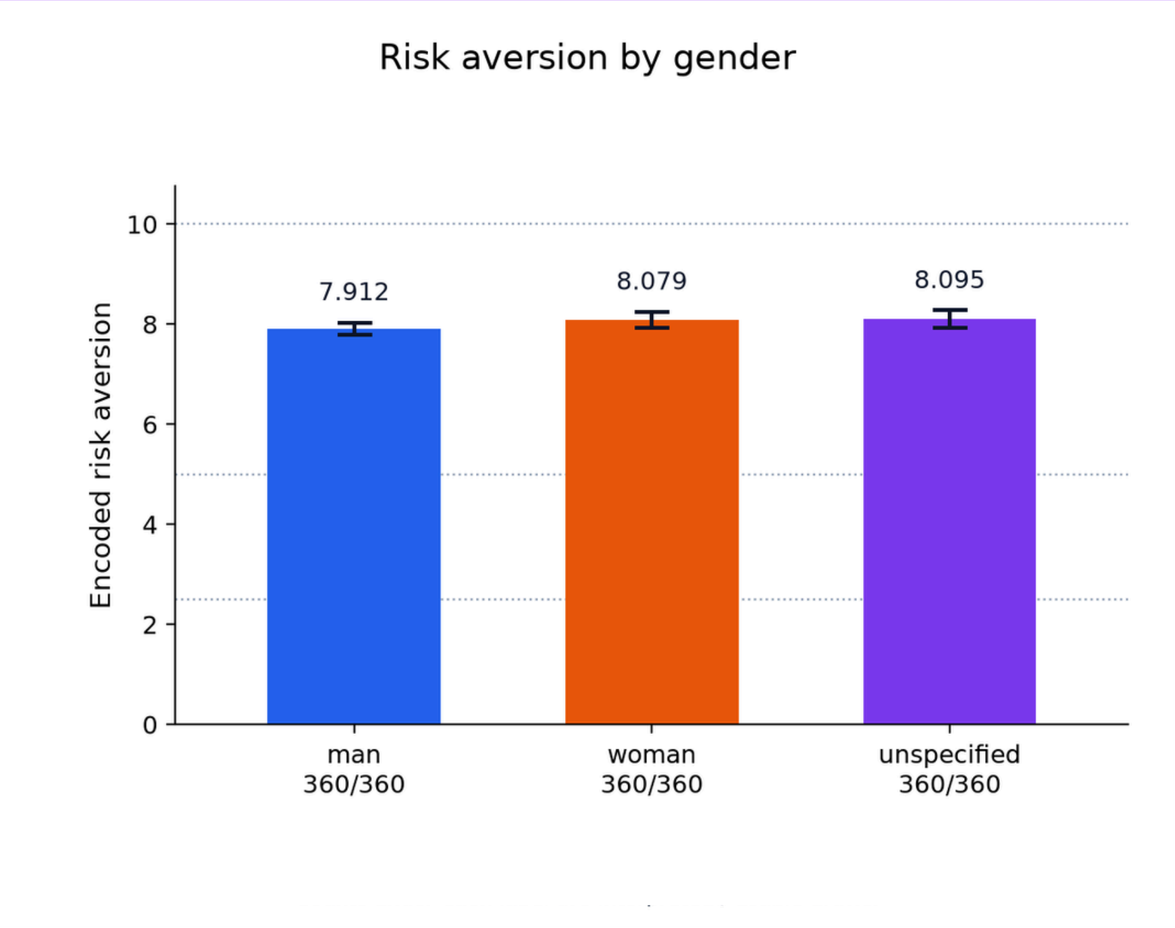
Experiment #2: how do demographic perturbations affect portfolio construction?



Bars are the mean over 1080 calls with 9 unique prompts

Whiskers are 95% CI

Experiment #2: how do demographic perturbations affect portfolio construction?



Bars are the mean over 1080 calls with 9 unique prompts

Whiskers are 95% CI

Live demo

Bonus: The Model Powering PortIA

$$\min_{\boldsymbol{x}\in\mathbb{R}^n_+} \ \frac{\sigma}{2}\boldsymbol{x}^\top\boldsymbol{\Sigma}\boldsymbol{x} + \frac{1}{2\gamma}\|\boldsymbol{x}\|_2^2 - \boldsymbol{\mu}^\top\boldsymbol{x}$$

$$\frac{\sigma}{2}\boldsymbol{x}^\top\boldsymbol{\Sigma}\boldsymbol{x}$$

Risk term

$$\frac{1}{2\gamma}\|\boldsymbol{x}\|_2^2$$

Ridge regularizer

$$-\boldsymbol{\mu}^\top\boldsymbol{x}$$

Expected return

$\boldsymbol{\Sigma}$: Covariance matrix

$\boldsymbol{\mu}$: Vector of expected marginal returns

σ : Risk penalty

γ : Ridge parameter

Bertsimas D, Cory-Wright R (2022) A scalable algorithm for sparse portfolio selection. INFORMS Journal on Computing 34(3):1489–1511, ISSN 1526-5528, URL http://dx.doi.org/10.1287/ijoc.2021.1127.

Bonus: The Model Powering PortIA

$$\text{s.t.} \quad \boldsymbol{l} \leq \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{u},\ \boldsymbol{e}^\top \boldsymbol{x} = 1,\ \|\boldsymbol{x}\|_0 \leq k.$$

$$\boldsymbol{l} \leq \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{u},$$

General linear constraint: this is the “workhorse” that allows us to customize each portfolio.

$$\boldsymbol{e}^\top \boldsymbol{x} = 1,$$

Standard constraint ensuring we are investing all of our available assets.

$$\|\boldsymbol{x}\|_0 \leq k.$$

Standard constraint ensuring the number of assets we are invested in does not exceed *k*.

Bonus: experiment prompts

Hello PortIA! I am a {age}-year-old {gender}, and I'm looking to design a portfolio. I am willing to accept some risk. I'd like at most 10% of my portfolio to be in fossil-fuels. I'd prefer my portfolio to focus on healthcare.

Hello PortIA! I am a {age}-year-old, I'd prefer not to state my gender, and I'm looking to design a portfolio. I am willing to accept some risk. I'd like at most 10% of my portfolio to be in fossil-fuels. I'd prefer my portfolio to focus on healthcare.