
pyblp
Release 0.8.1

Jeff Gortmaker

Nov 02, 2019

TABLE OF CONTENTS

| | | |
|----------|--|-----------|
| I | User Documentation | 1 |
| 1 | Introduction | 3 |
| 1.1 | Citation | 3 |
| 1.2 | Installation | 3 |
| 1.3 | Other Languages | 4 |
| 1.4 | Features | 4 |
| 1.5 | Features Slated for Future Versions | 5 |
| 1.6 | Bugs and Requests | 5 |
| 2 | Notation | 7 |
| 2.1 | Indices | 7 |
| 2.2 | Dimensions | 8 |
| 2.3 | Sets | 8 |
| 2.4 | Matrices, Vectors, and Scalars | 8 |
| 3 | Background | 11 |
| 3.1 | The Model | 11 |
| 3.2 | Estimation | 12 |
| 3.3 | Fixed Effects | 15 |
| 3.4 | Micro Moments | 15 |
| 3.5 | Random Coefficients Nested Logit | 16 |
| 3.6 | Logit and Nested Logit | 16 |
| 3.7 | Equilibrium Prices | 17 |
| 4 | Tutorial | 19 |
| 4.1 | Logit and Nested Logit Tutorial | 20 |
| 4.2 | Random Coefficients Logit Tutorial with the Fake Cereal Data | 32 |
| 4.3 | Random Coefficients Logit Tutorial with the Automobile Data | 48 |
| 4.4 | Post-Estimation Tutorial | 55 |
| 4.5 | Problem Simulation Tutorial | 70 |
| 5 | API Documentation | 77 |
| 5.1 | Configuration Classes | 77 |
| 5.2 | Data Construction Functions | 90 |
| 5.3 | Problem Class | 109 |
| 5.4 | Micro Moment Classes | 119 |
| 5.5 | Problem Results Class | 120 |
| 5.6 | Boostrapped Problem Results Class | 141 |
| 5.7 | Optimal Instrument Results Class | 143 |
| 5.8 | Simulation Class | 147 |

| | | |
|------------|---------------------------------------|------------|
| 5.9 | Simulation Results Class | 156 |
| 5.10 | Structured Data Classes | 159 |
| 5.11 | Multiprocessing | 160 |
| 5.12 | Options and Example Data | 164 |
| 5.13 | Exceptions | 170 |
| 6 | References | 179 |
| 6.1 | Conlon and Gortmaker (2019) | 179 |
| 6.2 | Other References | 179 |
| 7 | Legal | 185 |
| II | Developer Documentation | 187 |
| 8 | Contributing | 189 |
| 9 | Testing | 191 |
| 9.1 | Testing Requirements | 191 |
| 9.2 | Running Tests | 191 |
| 9.3 | Test Organization | 191 |
| 10 | Version Notes | 193 |
| 10.1 | 0.8 | 193 |
| 10.2 | 0.7 | 193 |
| 10.3 | 0.6 | 193 |
| 10.4 | 0.5 | 194 |
| 10.5 | 0.4 | 194 |
| 10.6 | 0.3 | 194 |
| 10.7 | 0.2 | 195 |
| 10.8 | 0.1 | 195 |
| III | Indices | 197 |
| | Python Module Index | 199 |
| | Index | 201 |

Part I

User Documentation

INTRODUCTION

Note: This package is in beta. The API may change in future versions. Please use the [GitHub issue tracker](#) to report bugs or to request features.

The `pyblp` package is a Python 3 implementation of routines for estimating the demand for differentiated products with BLP-type random coefficients logit models. This package was created by [Jeff Gortmaker](#) in collaboration with [Chris Conlon](#).

Development of the package has been guided by the work of many researchers and practitioners. For a full list of references, including the original work of [Berry, Levinsohn, and Pakes \(1995\)](#), refer to the [references](#) section of the documentation.

1.1 Citation

If you use `pyblp` in your research, we ask that you also cite [Conlon and Gortmaker \(2019\)](#), which describes the advances implemented in the package.

1.2 Installation

The `pyblp` package has been tested on [Python](#) versions 3.6 and 3.7. The [SciPy instructions](#) for installing related packages is a good guide for how to install a scientific Python environment. A good choice is the [Anaconda Distribution](#), since, along with many other packages that are useful for scientific computing, it comes packaged with `pyblp`'s only required dependencies: [NumPy](#), [SciPy](#), [SymPy](#), and [Patsy](#).

However, `pyblp` may not work with old versions of its dependencies. You can update `pyblp`'s dependencies in [Anaconda](#) with:

```
conda update numpy scipy sympy patsy
```

You can install the current release of `pyblp` with `pip`:

```
pip install pyblp
```

You can upgrade to a newer release with the `--upgrade` flag:

```
pip install --upgrade pyblp
```

If you lack permissions, you can install `pyblp` in your user directory with the `--user` flag:

```
pip install --user pyblp
```

Alternatively, you can download a wheel or source archive from [PyPI](#). You can find the latest development code on [GitHub](#) and the latest development documentation [here](#).

1.3 Other Languages

Once installed, pyblp can be incorporated into projects written in many other languages with the help of various tools that enable interoperability with Python.

For example, the [reticulate](#) package makes interacting with pyblp in R straightforward:

```
library(reticulate)
pyblp <- import("pyblp")
```

Similarly, [PyCall](#) can be used to incorporate pyblp into a Julia workflow:

```
using PyCall
pyblp = pyimport("pyblp")
```

The [py command](#) serves a similar purpose in MATLAB:

```
py.pyblp
```

1.4 Features

- R-style formula interface
- Bertrand-Nash supply-side moments
- Multiple equation GMM
- Demographic interactions
- Micro moments that match product and agent characteristic covariances
- Fixed effect absorption
- Nonlinear functions of product characteristics
- Concentrating out linear parameters
- Parameter bounds and constraints
- Random coefficients nested logit (RCNL)
- Varying nesting parameters across groups
- Logit and nested logit benchmarks
- Classic BLP instruments
- Differentiation instruments
- Optimal instruments
- Tests of overidentifying and model restrictions
- Parametric bootstrapping post-estimation outputs

- Elasticities and diversion ratios
- Marginal costs and markups
- Profits and consumer surplus
- Merger simulation
- Custom counterfactual simulation
- Synthetic data construction
- SciPy or Artleys Knitro optimization
- Fixed point acceleration
- Monte Carlo, product rule, or sparse grid integration
- Custom optimization and iteration routines
- Robust and clustered errors
- Linear or log-linear marginal costs
- Partial ownership matrices
- Analytic gradients
- Finite difference Hessians
- Market-by-market parallelization
- Extended floating point precision
- Robust error handling

1.5 Features Slated for Future Versions

- Other micro moments
- Fast, “Robust,” and Approximately Correct (FRAC) estimation
- Analytic Hessians
- Mathematical Program with Equilibrium Constraints (MPEC)
- Generalized Empirical Likelihood (GEL)
- Discrete types
- Pure characteristics model
- Newton methods for computing equilibrium prices

1.6 Bugs and Requests

Please use the [GitHub issue tracker](#) to submit bugs or to request features.

NOTATION

The notation in pyblp is a customized amalgamation of the notation employed by *Berry, Levinsohn, and Pakes (1995)*, *Nevo (2000)*, *Morrow and Skerlos (2011)*, *Grigolon and Verboven (2014)*, and others.

2.1 Indices

| Index | Description |
|-------|-----------------------|
| j | Products |
| t | Markets |
| i | Agents or individuals |
| f | Firms |
| h | Nests |
| c | Clusters |
| m | Micro moments |

2.2 Dimensions

| Dimension | Description |
|-----------|--|
| T | Markets |
| N | Products across all markets |
| F | Firms across all markets |
| I | Agents across all markets |
| J_t | Products in market t |
| F_t | Firms in market t |
| I_t | Agents in market t |
| K_1 | Linear product characteristics |
| K_1^x | Exogenous linear product characteristics |
| K_1^p | Endogenous linear product characteristics |
| K_2 | Nonlinear product characteristics |
| K_3 | Cost product characteristics |
| D | Demographic variables |
| M_D | Demand-side instruments, which is the number of exluded demand-side instruments plus K_1^x |
| M_S | Supply-side instruments, which is the number of exluded supply-side instruments plus K_3 |
| M_M | Micro moments |
| M | Moments: the sum of M_D , M_S , and M_M |
| E_D | Absorbed dimensions of demand-side fixed effects |
| E_S | Absorbed dimensions of supply-side fixed effects |
| H | Nesting groups |
| C | Clusters |
| P | Parameters |

2.3 Sets

| Set | Description |
|--------------------|---|
| \mathcal{J}_{ft} | Products produced by firm f in market t |
| \mathcal{J}_{ht} | Products in nesting group h and market t |
| \mathcal{J}_{ct} | Products in cluster c and market t |
| \mathcal{T}_m | Markets over which micro moment m is averaged |

2.4 Matrices, Vectors, and Scalars

| Symbol | Dimensions | Description |
|----------|------------------|--|
| X_1 | $N \times K_1$ | Linear product characteristics |
| X_1^x | $N \times K_1^x$ | Exogenous linear product characteristics |
| X_1^p | $N \times K_1^p$ | Endogenous linear product characteristics |
| X_2 | $N \times K_2$ | Nonlinear product characteristics |
| X_3 | $N \times K_3$ | Cost product characteristics |
| ξ | $N \times 1$ | Unobserved demand-side product characteristics |
| ω | $N \times 1$ | Unobserved supply-side product characteristics |
| p | $N \times 1$ | Prices |

Continued on next page

Table 1 – continued from previous page

| Symbol | Dimensions | Description |
|--|------------------------|---|
| $s(s_{jt})$ | $N \times 1$ | Marketshares |
| $s(s_{ht})$ | $H \times 1$ | Group shares in a market t |
| $s(s_{jti})$ | $N \times I_t$ | Choice probabilities in a market t |
| c | $N \times 1$ | Marginal costs |
| \tilde{c} | $N \times 1$ | Linear or log-linear marginal costs, c or $\log c$ |
| η | $N \times 1$ | Markup term from the BLP-markup equation |
| ζ | $N \times 1$ | Markup term from the ζ -markup equation |
| O | $J_t \times J_t$ | Ownership matrix in market t |
| κ | $F_t \times F_t$ | Cooperation matrix in market t |
| Δ | $J_t \times J_t$ | Intra-firm matrix of (negative) demand derivatives in market t |
| Λ | $J_t \times J_t$ | Diagonal matrix used to decompose η and ζ in market t |
| Γ | $J_t \times J_t$ | Another matrix used to decompose η and ζ in market t |
| d | $I_t \times D$ | Observed agent characteristics called demographics in market t |
| ν | $I_t \times K_2$ | Unobserved agent characteristics called integration nodes in market t |
| w | $I_t \times 1$ | Integration weights in market t |
| δ | $N \times 1$ | Mean utility |
| μ | $J_t \times I_t$ | Agent-specific portion of utility in market t |
| ϵ | $N \times 1$ | Type I Extreme Value idiosyncratic preferences |
| $\bar{\epsilon}(\bar{\epsilon}_{jti})$ | $N \times 1$ | Type I Extreme Value term used to decompose ϵ |
| $\bar{\epsilon}(\bar{\epsilon}_{hti})$ | $N \times 1$ | Group-specific term used to decompose ϵ |
| U | $J_t \times I_t$ | Indirect utilities |
| $V(V_{jti})$ | $J_t \times I_t$ | Indirect utilities minus ϵ |
| $V(V_{hti})$ | $J_t \times I_t$ | Inclusive value of a nesting group |
| $\pi(\pi_{jt})$ | $N \times 1$ | Population-normalized gross expected profits |
| $\pi(\pi_{ft})$ | $F_t \times 1$ | Population-normalized gross expected profits of a firm in market t |
| β | $K_1 \times 1$ | Demand-side linear parameters |
| β^x | $K_1^x \times 1$ | Parameters in β on exogenous product characteristics |
| α | $K_1^p \times 1$ | Parameters in β on endogenous product characteristics |
| Σ | $K_2 \times K_2$ | Cholesky root of the covariance matrix for unobserved taste heterogeneity |
| Π | $K_2 \times D$ | Parameters that measures how agent tastes vary with demographics |
| ρ | $H \times 1$ | Parameters that measures within nesting group correlation |
| γ | $K_3 \times 1$ | Supply-side linear parameters |
| θ | $P \times 1$ | Parameters |
| Z_D | $N \times M_D$ | Excluded demand-side instruments and X_1 , except for X_1^p |
| Z_S | $N \times M_S$ | Excluded supply-side instruments and X_3 |
| W | $M \times M$ | Weighting matrix |
| S | $M \times M$ | Moment covariances |
| q | 1×1 | Objective value |
| g_D | $N \times M_D$ | Demand-side moments |
| g_S | $N \times M_S$ | Supply-side moments |
| g_M | $I \times M_M$ | Micro moments |
| $g(g_{jt})$ | $N \times (M_D + M_S)$ | Demand- and supply-side moments |
| $g(g_c)$ | $C \times (M_D + M_S)$ | Clustered demand- and supply-side moments |
| \bar{g}_D | $M_D \times 1$ | Averaged demand-side moments |
| \bar{g}_S | $M_S \times 1$ | Averaged supply-side moments |
| \bar{g}_M | $M_M \times 1$ | Averaged micro moments |
| \bar{g} | $M \times 1$ | Averaged moments |
| G | $M \times P$ | Jacobian of the averaged moments with respect to θ |
| ϵ | $J_t \times J_t$ | Elasticities of demand in market t |

Continued on next page

Table 1 – continued from previous page

| Symbol | Dimensions | Description |
|---------------|------------------|--|
| \mathcal{D} | $J_t \times J_t$ | Diversion ratios in market t |
| \mathcal{D} | $J_t \times J_t$ | Long-run diversion ratios in market t |
| \mathcal{M} | $N \times 1$ | Markups |
| \mathcal{E} | 1×1 | Aggregate elasticity of demand of a market |
| CS | 1×1 | Population-normalized consumer surplus of a market |
| HHI | 1×1 | Herfindahl-Hirschman Index of a market |

BACKGROUND

The following sections provide a very brief overview of the BLP model and how it is estimated. This goal is to concisely introduce the notation and terminology used throughout the rest of the documentation. For a more in-depth overview, refer to *Conlon and Gortmaker (2019)*.

3.1 The Model

There are $t = 1, 2, \dots, T$ markets, each with $j = 1, 2, \dots, J_t$ products produced by $f = 1, 2, \dots, F_t$ firms, for a total of N products across all markets. There are $i = 1, 2, \dots, I_t$ agents who choose among the J_t products and an outside good $j = 0$.

3.1.1 Demand

Observed demand-side product characteristics are contained in the $N \times K_1$ matrix of linear characteristics, X_1 , and the $N \times K_2$ matrix of nonlinear characteristics, X_2 , which is typically a subset of X_1 . Unobserved demand-side product characteristics, ξ , are a $N \times 1$ vector.

In market t , observed agent characteristics are a $I_t \times D$ matrix called demographics, d . Unobserved agent characteristics are a $I_t \times K_2$ matrix, ν .

The indirect utility of agent i from purchasing product j in market t is

$$U_{jti} = \underbrace{\delta_{jt} + \mu_{jti}}_{V_{jti}} + \epsilon_{jti}, \quad (3.1)$$

in which the mean utility is, in vector-matrix form,

$$\delta = \underbrace{X_1^p \alpha + X_1^x \beta^x}_{X_1 \beta} + \xi. \quad (3.2)$$

The $K_1 \times 1$ vector of demand-side linear parameters, β , is partitioned into two components: α is a $K_1^p \times 1$ vector of parameters on the $N \times K_1^p$ submatrix of endogenous characteristics, X_1^p , and β^x is a $K_1^x \times 1$ vector of parameters on the $N \times K_1^x$ submatrix of exogenous characteristics, X_1^x . Usually, $X_1^p = p$, prices, so α is simply a scalar.

The agent-specific portion of utility in a single market is

$$\mu = X_2(\Sigma \nu' + \Pi d'). \quad (3.3)$$

The model incorporates both observable (demographic) and unobservable taste heterogeneity through random coefficients. For the unobserved heterogeneity, we let ν denote independent draws from the standard normal distribution. These are scaled by a $K_2 \times K_2$ lower-triangular matrix Σ , which denotes the Cholesky root of the covariance matrix for unobserved taste heterogeneity. The $K_2 \times D$ matrix Π measures how agent tastes vary with demographics.

Random idiosyncratic preferences, ϵ_{jti} , are assumed to be Type I Extreme Value, so that conditional on the heterogeneous coefficients, marketshares follow the well known logit form. Aggregate marketshares are obtained by integrating over the distribution of individual heterogeneity. They are approximated with Monte Carlo integration or quadrature rules defined by the $I_t \times K_2$ matrix of integration nodes, ν , and a $I_t \times 1$ vector of integration weights, w :

$$s_{jt} \approx \sum_{i=1}^{I_t} w_{it} s_{jti}, \quad (3.4)$$

where the probability that agent i chooses product j in market t is

$$s_{jti} = \frac{\exp V_{jti}}{1 + \sum_{k=1}^{J_t} \exp V_{kti}}. \quad (3.5)$$

There is a one in the denominator because the utility of the outside good is normalized to $U_{0ti} = 0$.

3.1.2 Supply

Observed supply-side product characteristics are contained in the $N \times K_3$ matrix of cost characteristics, X_3 . Prices cannot be cost characteristics, but non-price product characteristics often overlap with the demand-side characteristics in X_1 and X_2 . Unobserved supply-side product characteristics, ω , are a $N \times 1$ vector.

Firm f chooses prices in market t to maximize the profits of its products $\mathcal{J}_{ft} \subset \{1, 2, \dots, J_t\}$:

$$\pi_{ft} = \sum_{j \in \mathcal{J}_{ft}} (p_{jt} - c_{jt}) s_{jt}. \quad (3.6)$$

In a single market, the corresponding multi-product differentiated Bertrand first order conditions are, in vector-matrix form,

$$p - c = \underbrace{\Delta^{-1}}_{\eta} s, \quad (3.7)$$

where the multi-product Bertrand markup η depends on Δ , a $J_t \times J_t$ matrix of intra-firm (negative) demand derivatives:

$$\Delta = -O \odot \frac{\partial s}{\partial p}. \quad (3.8)$$

Here, O denotes the market-level ownership matrix, where O_{jk} is typically 1 if the same firm produces products j and k , and 0 otherwise.

To include a supply side, we must specify a functional form for marginal costs:

$$\tilde{c} = f(c) = X_3 \gamma + \omega. \quad (3.9)$$

The most common choices are $f(c) = c$ and $f(c) = \log(c)$.

3.2 Estimation

A demand side is always estimated but including a supply side is optional. With only a demand side, there are three sets of parameters to be estimated: β (which may include α), Σ and Π . With a supply side, there is also γ . The linear parameters, β and γ , are typically concentrated out of the problem. The exception is α , which cannot be concentrated out when there is a supply side because it is needed to compute demand derivatives and hence marginal costs. Linear parameters that are not concentrated out along with unknown nonlinear parameters in Σ and Π are collectively denoted θ , a $P \times 1$ vector.

The GMM problem is

$$\min_{\theta} q(\theta) = \bar{g}(\theta)' W \bar{g}(\theta), \quad (3.10)$$

in which $q(\theta)$ is the GMM objective (n.b., in some of the BLP literature and earlier versions of this package, this expression was scaled by N^2 , but it is no longer scaled by this number), W is a $M \times M$ weighting matrix and \bar{g} is a $M \times 1$ vector of averaged demand- and supply-side moments:

$$\bar{g} = \begin{bmatrix} \bar{g}_D \\ \bar{g}_S \end{bmatrix} = \frac{1}{N} \begin{bmatrix} \sum_{j,t} Z'_{D,jt} \xi_{jt} \\ \sum_{j,t} Z'_{S,jt} \omega_{jt} \end{bmatrix} \quad (3.11)$$

where Z_D and Z_S are $N \times M_D$ and $N \times M_S$ matrices of demand- and supply-side instruments containing excluded instruments along with X_1^x and X_3 , respectively. When there are only demand- and supply-side moments, $M = M_D + M_S$.

The vector \bar{g} contains sample analogues of the demand- and supply-side moment conditions $E[g_{D,jt}] = E[g_{S,jt}] = 0$ where

$$\begin{bmatrix} g_{D,jt} & g_{S,jt} \end{bmatrix} = \begin{bmatrix} \xi_{jt} Z_{D,jt} & \omega_{jt} Z_{S,jt} \end{bmatrix}. \quad (3.12)$$

In each GMM stage, a nonlinear optimizer finds the $\hat{\theta}$ that minimizes the GMM objective value $q(\theta)$.

3.2.1 The Objective

Given a $\hat{\theta}$, the first step to computing the objective $q(\hat{\theta})$ is to compute $\delta(\hat{\theta})$ in each market with the following standard contraction:

$$\delta_{jt} \leftarrow \delta_{jt} + \log s_{jt} - \log s_{jt}(\delta, \hat{\theta}) \quad (3.13)$$

where s are the market's observed shares and $s(\delta, \hat{\theta})$ are calculated marketshares. Iteration terminates when the norm of the change in $\delta(\hat{\theta})$ is less than a small number.

With a supply side, marginal costs are then computed according to (3.7):

$$c_{jt}(\hat{\theta}) = p_{jt} - \eta_{jt}(\hat{\theta}). \quad (3.14)$$

Concentrated out linear parameters are recovered with linear IV-GMM:

$$\begin{bmatrix} \hat{\beta}^x \\ \hat{\gamma} \end{bmatrix} = (X' Z W Z' X)^{-1} X' Z W Z' Y(\hat{\theta}) \quad (3.15)$$

where

$$X = \begin{bmatrix} X_1^x & 0 \\ 0 & X_3 \end{bmatrix}, \quad Z = \begin{bmatrix} Z_D & 0 \\ 0 & Z_S \end{bmatrix}, \quad Y(\hat{\theta}) = \begin{bmatrix} \delta(\hat{\theta}) - X_1^p \hat{\alpha} & 0 \\ 0 & \tilde{c}(\hat{\theta}) \end{bmatrix}. \quad (3.16)$$

With only a demand side, α can be concentrated out, so $X = X_1$, $Z = Z_D$, and $Y = \delta(\hat{\theta})$ recover the full $\hat{\beta}$ in (3.15).

Finally, the unobserved product characteristics (structural errors),

$$\begin{bmatrix} \xi(\hat{\theta}) \\ \omega(\hat{\theta}) \end{bmatrix} = \begin{bmatrix} \delta(\hat{\theta}) - X_1 \hat{\beta} \\ \tilde{c}(\hat{\theta}) - X_3 \hat{\gamma} \end{bmatrix}, \quad (3.17)$$

are interacted with the instruments to form $\bar{g}(\hat{\theta})$ in (3.11), which give the GMM objective $q(\hat{\theta})$ in (3.10).

3.2.2 The Gradient

The gradient of the GMM objective in (3.10) is

$$\nabla q(\theta) = 2\bar{G}(\theta)'W\bar{g}(\theta) \quad (3.18)$$

where

$$\bar{G} = \begin{bmatrix} \bar{G}_D \\ \bar{G}_S \end{bmatrix} = \frac{1}{N} \begin{bmatrix} \sum_{j,t} Z'_{D,jt} \frac{\partial \xi_{jt}}{\partial \theta} \\ \sum_{j,t} Z'_{S,jt} \frac{\partial \omega_{jt}}{\partial \theta} \end{bmatrix}. \quad (3.19)$$

Writing δ as an implicit function of s in (3.4) gives the demand-side Jacobian:

$$\frac{\partial \xi}{\partial \theta} = \frac{\partial \delta}{\partial \theta} = - \left(\frac{\partial s}{\partial \delta} \right)^{-1} \frac{\partial s}{\partial \theta}. \quad (3.20)$$

The supply-side Jacobian is derived from the definition of \tilde{c} in (3.9):

$$\frac{\partial \omega}{\partial \theta} = \frac{\partial \tilde{c}}{\partial \theta_p} = - \frac{\partial \tilde{c}}{\partial c} \frac{\partial \eta}{\partial \theta}. \quad (3.21)$$

The second term in this expression is derived from the definition of η in (3.7):

$$\frac{\partial \eta}{\partial \theta} = -\Delta^{-1} \left(\frac{\partial \Delta}{\partial \theta} \eta + \frac{\partial \Delta}{\partial \xi} \eta \frac{\partial \xi}{\partial \theta} \right). \quad (3.22)$$

3.2.3 Weighting Matrices

Conventionally, the 2SLS weighting matrix is used in the first stage:

$$W = \begin{bmatrix} (Z'_D Z_D / N)^{-1} & 0 \\ 0 & (Z'_S Z_S / N)^{-1} \end{bmatrix}. \quad (3.23)$$

With two-step GMM, W is updated before the second stage according to

$$W = S^{-1}. \quad (3.24)$$

For heteroscedasticity robust weighting matrices,

$$S = \frac{1}{N} \sum_{j,t} g_{jt} g'_{jt}. \quad (3.25)$$

For clustered weighting matrices, which account for arbitrary correlation within $c = 1, 2, \dots, C$ clusters,

$$S = \frac{1}{N} \sum_{c=1}^C g_c g'_c, \quad (3.26)$$

where, letting the set $\mathcal{J}_c \subset \{1, 2, \dots, N\}$ denote products in cluster c ,

$$g_c = \sum_{t=1}^T \sum_{j \in \mathcal{J}_{ct}} g_{jt}. \quad (3.27)$$

For unadjusted weighting matrices,

$$S = \frac{1}{N} \begin{bmatrix} \sigma_\xi^2 Z'_D Z_D & \sigma_{\xi\omega} Z'_D Z_S \\ \sigma_{\xi\omega} Z'_S Z_D & \sigma_\omega^2 Z'_S Z_S \end{bmatrix} \quad (3.28)$$

where

$$\text{Var}(\xi, \omega) = \begin{bmatrix} \sigma_\xi^2 & \sigma_{\xi\omega} \\ \sigma_{\xi\omega} & \sigma_\omega^2 \end{bmatrix}. \quad (3.29)$$

3.2.4 Standard Errors

The covariance matrix of the estimated parameters is

$$\text{Var}(\hat{\theta}) = (\bar{G}'W\bar{G})^{-1}\bar{G}'WSW\bar{G}(\bar{G}'W\bar{G})^{-1}. \quad (3.30)$$

Standard errors are the square root of the diagonal of this matrix divided by N .

If the weighting matrix was chosen such that $W = S^{-1}$, this simplifies to

$$\text{Var}(\hat{\theta}) = (\bar{G}'W\bar{G})^{-1}. \quad (3.31)$$

Standard errors extracted from this simpler expression are called unadjusted.

3.3 Fixed Effects

The unobserved product characteristics can be partitioned into

$$\begin{bmatrix} \xi_{jt} \\ \omega_{jt} \end{bmatrix} = \begin{bmatrix} \xi_{k_1} + \xi_{k_2} + \cdots + \xi_{k_{E_D}} + \Delta\xi_{jt} \\ \omega_{\ell_1} + \omega_{\ell_2} + \cdots + \omega_{\ell_{E_S}} + \Delta\omega_{jt} \end{bmatrix} \quad (3.32)$$

where k_1, k_2, \dots, k_{E_D} and $\ell_1, \ell_2, \dots, \ell_{E_S}$ index unobserved characteristics that are fixed across E_D and E_S dimensions. For example, with $E_D = 1$ dimension of product fixed effects, $\xi_{jt} = \xi_j + \Delta\xi_{jt}$.

Small numbers of fixed effects can be estimated with dummy variables in X_1, X_3, Z_D , and Z_S . However, this approach does not scale with high dimensional fixed effects because it requires constructing and inverting an infeasibly large matrix in (3.15).

Instead, fixed effects are typically absorbed into X, Z , and $Y(\hat{\theta})$ in (3.15). With one fixed effect, these matrices are simply de-meaned within each level of the fixed effect. Both X and Z can be de-meaned just once, but $Y(\hat{\theta})$ must be de-meaned for each new $\hat{\theta}$.

This procedure is equivalent to replacing each column of the matrices with residuals from a regression of the column on the fixed effect. The Frish-Waugh-Lovell (FWL) theorem of *Frisch and Waugh (1933)* and *Lovell (1963)* guarantees that using these residualized matrices gives the same results as including fixed effects as dummy variables. When $E_D > 1$ or $E_S > 1$, the matrices are residualized with more involved algorithms.

Once fixed effects have been absorbed, estimation is as described above with the structural errors $\Delta\xi$ and $\Delta\omega$.

3.4 Micro Moments

In the spirit of *Imbens and Lancaster (1994)*, *Petrin (2002)*, and *Berry, Levinsohn, and Pakes (2004)*, more detailed micro data on individual agent decisions can be used to supplement the standard demand- and supply-side moments \bar{g}_D and \bar{g}_S in (3.11) with an additional $m = 1, 2, \dots, M_M$ averaged micro moments, \bar{g}_M , for a total of $M = M_D + M_S + M_M$ averaged moments:

$$\bar{g} = \begin{bmatrix} \bar{g}_D \\ \bar{g}_S \\ \bar{g}_M \end{bmatrix}. \quad (3.33)$$

Each micro moment m is approximated in a set $\mathcal{T}_m \subset \{1, 2, \dots, T\}$ of markets in which its micro data are relevant and then averaged across these markets:

$$\bar{g}_{M,m} \approx \frac{1}{|\mathcal{T}_m|} \sum_{t \in \mathcal{T}_m} \sum_{i=1}^{I_t} w_{it} g_{M,mti}. \quad (3.34)$$

The vector \bar{g}_M contains sample analogues of micro moment conditions $E[g_{M,mti}] = 0$ where $g_{M,mti}$ is typically a function of choice probabilities, data in market t , and a statistic computed from survey data that the moment aims to match.

Mico moments are computed for each $\hat{\theta}$ and contribute to the GMM objective $q(\hat{\theta})$ in (3.10). Their derivatives with respect to θ are added as rows to \bar{G} in (3.19), and blocks are added to both W and S in (3.23) and (3.24). The covariance between standard moments and micro moments is assumed to be zero, so these matrices will be block-diagonal. The covariance between micro moments m and n in S is set to zero if $\mathcal{I}_{mn} = \mathcal{I}_m \cap \mathcal{I}_n = \emptyset$ and otherwise is approximated by

$$\text{Cov}(\bar{g}_{M,m}, \bar{g}_{M,n}) \approx \frac{1}{|\mathcal{I}_{mn}|} \sum_{t \in \mathcal{I}_{mn}} \sum_{i=1}^{I_t} w_{it} (g_{M,mti} - \bar{g}_{M,mt}) (g_{M,nti} - \bar{g}_{M,nt}) \quad (3.35)$$

where $\bar{g}_{M,mt} = \sum_i w_{it} g_{M,mti}$.

3.5 Random Coefficients Nested Logit

Incorporating parameters that measure within nesting group correlation gives the random coefficients nested logit (RCNL) model of *Brenkers and Verboven (2006)* and *Grigolon and Verboven (2014)*. There are $h = 1, 2, \dots, H$ nesting groups and each product j is assigned to a group $h(j)$. The set $\mathcal{J}_{ht} \subset \{1, 2, \dots, J_t\}$ denotes the products in group h and market t .

In the RCNL model, idiosyncratic preferences are partitioned into

$$\epsilon_{jti} = \bar{\epsilon}_{h(j)ti} + (1 - \rho_{h(j)}) \bar{\epsilon}_{jti} \quad (3.36)$$

where $\bar{\epsilon}_{jti}$ is Type I Extreme Value and $\bar{\epsilon}_{h(j)ti}$ is distributed such that ϵ_{jti} is still Type I Extreme Value.

The nesting parameters, ρ , can either be a $H \times 1$ vector or a scalar so that for all groups $\rho_h = \rho$. Letting $\rho \rightarrow 0$ gives the standard BLP model and $\rho \rightarrow 1$ gives division by zero errors. With $\rho_h \in (0, 1)$, the expression for choice probabilities in (3.5) becomes more complicated:

$$s_{jti} = \frac{\exp[V_{jti}/(1 - \rho_{h(j)})]}{\exp[V_{h(j)ti}/(1 - \rho_{h(j)})]} \cdot \frac{\exp V_{h(j)ti}}{1 + \sum_{h=1}^H \exp V_{hti}} \quad (3.37)$$

where

$$V_{hti} = (1 - \rho_h) \log \sum_{k \in \mathcal{J}_{ht}} \exp[V_{kti}/(1 - \rho_h)]. \quad (3.38)$$

The contraction for $\delta(\hat{\theta})$ in (3.13) is also slightly different:

$$\delta_{jt} \leftarrow \delta_{jt} + (1 - \rho_{h(j)}) [\log s_{jt} - \log s_{jt}(\delta, \hat{\theta})]. \quad (3.39)$$

Otherwise, estimation is as described above with ρ included in θ .

3.6 Logit and Nested Logit

Letting $\Sigma = 0$ gives the simpler logit (or nested logit) model where there is a closed-form solution for δ . In the logit model,

$$\delta_{jt} = \log s_{jt} - \log s_{0t}, \quad (3.40)$$

and a lack of nonlinear parameters means that nonlinear optimization is not needed.

In the nested logit model, ρ must be optimized over, but there is still a closed-form solution for δ :

$$\delta_{jt} = \log s_{jt} - \log s_{0t} - \rho_{h(j)} [\log s_{jt} - \log s_{h(j)t}]. \quad (3.41)$$

where

$$s_{ht} = \sum_{j \in \mathcal{J}_{ht}} s_{jt}. \quad (3.42)$$

In both models, a supply side can still be estimated jointly with demand. Estimation is as described above with a representative agent in each market: $I_t = 1$ and $w_1 = 1$.

3.7 Equilibrium Prices

Counterfactual evaluation, synthetic data simulation, and optimal instrument generation often involve solving for prices implied by the Bertrand first order conditions in (3.7). Solving this system with Newton's method is slow and iterating over $p \leftarrow c + \eta(p)$ may not converge because it is not a contraction.

Instead, *Morrow and Skerlos (2011)* reformulate the solution to (3.7):

$$p - c = \underbrace{\Lambda^{-1}(O \odot \Gamma)'(p - c) - \Lambda^{-1}}_{\zeta} \quad (3.43)$$

where Λ is a diagonal $J_t \times J_t$ matrix approximated by

$$\Lambda_{jj} \approx \sum_{i=1}^{I_t} w_{it} s_{jti} \frac{\partial U_{jti}}{\partial p_{jt}} \quad (3.44)$$

and Γ is a dense $J_t \times J_t$ matrix approximated by

$$\Gamma_{jk} \approx \sum_{i=1}^{I_t} w_{it} s_{jti} s_{kti} \frac{\partial U_{jti}}{\partial p_{jt}}. \quad (3.45)$$

Equilibrium prices are computed by iterating over the ζ -markup equation in (3.43),

$$p \leftarrow c + \zeta(p), \quad (3.46)$$

which, unlike (3.7), is a contraction. Iteration terminates when the norm of firms' first order conditions, $\|\Lambda(p)(p - c - \zeta(p))\|$, is less than a small number.

TUTORIAL

This section uses a series of [Jupyter Notebooks](#) to explain how `pyblp` can be used to solve example problems, compute post-estimation outputs, and simulate problems. Other examples can be found in the [API Documentation](#).

The [online version](#) of the following section may be easier to read.

4.1 Logit and Nested Logit Tutorial

```
import pyblp
import numpy as np
import pandas as pd

pyblp.options.digits = 2
pyblp.options.verbose = False
pyblp.__version__

'0.8.1'
```

In this tutorial, we'll use data from *Nevo (2000)* to solve the paper's fake cereal problem. Locations of CSV files that contain the data are in the `data` module.

We will compare two simple models, the plain (IIA) logit model and the nested logit (GEV) model using the fake cereal dataset of *Nevo (2000)*.

4.1.1 Theory of Plain Logit

Let's start with the plain logit model under independence of irrelevant alternatives (IIA). In this model (indirect) utility is given by

$$U_{jti} = \alpha p_{jt} + x_{jt}\beta^x + \xi_{jt} + \epsilon_{jti}, \quad (4.1)$$

where ϵ_{jti} is distributed IID with the Type I Extreme Value (Gumbel) distribution. It is common to normalize the mean utility of the outside good to zero so that $U_{0ti} = \epsilon_{0ti}$. This gives us aggregate marketshares

$$s_{jt} = \frac{\exp(\alpha p_{jt} + x_{jt}\beta^x + \xi_{jt})}{1 + \sum_k \exp(\alpha p_{kt} + x_{kt}\beta^x + \xi_{kt})}. \quad (4.2)$$

If we take logs we get

$$\log s_{jt} = \alpha p_{jt} + x_{jt}\beta^x + \xi_{jt} - \log \sum_k \exp(\alpha p_{kt} + x_{kt}\beta^x + \xi_{kt}) \quad (4.3)$$

and

$$\log s_{0t} = -\log \sum_k \exp(\alpha p_{kt} + x_{kt}\beta^x + \xi_{kt}). \quad (4.4)$$

By differencing the above we get a linear estimating equation:

$$\log s_{jt} - \log s_{0t} = \alpha p_{jt} + x_{jt} \beta^x + \xi_{jt}. \quad (4.5)$$

Because the left hand side is data, we can estimate this model using linear IV GMM.

4.1.2 Application of Plain Logit

A Logit *Problem* can be created by simply excluding the formulation for the nonlinear parameters, X_2 , along with any agent information. In other words, it requires only specifying the *linear component* of demand.

We'll set up and solve a simple version of the fake data cereal problem from *Nevo (2000)*. Since we won't include any nonlinear characteristics or parameters, we don't have to worry about configuring an optimization routine.

There are some important reserved variable names:

- `market_ids` are the unique market identifiers which we subscript with t .
- `shares` specifies the marketshares which need to be between zero and one, and within a market ID, $\sum_j s_{jt} \leq 1$.
- `prices` are prices p_{jt} . These have some special properties and are *always* treated as endogenous.
- `demand_instruments0`, `demand_instruments1`, and so on are numbered demand instruments. These represent only the *excluded* instruments. The exogenous regressors in X_1 will be automatically added to the set of instruments.

We begin with two steps:

1. Load the product data which at a minimum consists of `market_ids`, `shares`, `prices`, and at least a single column of demand instruments, `demand_instruments0`.
2. Define a *Formulation* for the X_1 (linear) demand model.
 - This and all other formulas are similar to R and `patsy` formulas.
 - It includes a constant by default. To exclude the constant, specify either a 0 or a -1.
 - To efficiently include fixed effects, use the `absorb` option and specify which categorical variables you would like to absorb.
 - Some model reduction may happen automatically. The constant will be excluded if you include fixed effects and some precautions are taken against collinearity. However, you will have to make sure that differently-named variables are not collinear.
3. Combine the *Formulation* and product data to construct a *Problem*.
4. Use `Problem.solve` to estimate parameters.

Loading the Data

The `product_data` argument of *Problem* should be a structured array-like object with fields that store data. Product data can be a structured NumPy array, a pandas DataFrame, or other similar objects.

```
product_data = pd.read_csv(pyblp.data.NEVO_PRODUCTS_LOCATION)
product_data.head()
```

| | market_ids | city_ids | quarter | product_ids | firm_ids | brand_ids | shares | \ |
|---|------------|----------|---------|-------------|----------|-----------|----------|---|
| 0 | C01Q1 | 1 | 1 | F1B04 | 1 | 4 | 0.012417 | |
| 1 | C01Q1 | 1 | 1 | F1B06 | 1 | 6 | 0.007809 | |
| 2 | C01Q1 | 1 | 1 | F1B07 | 1 | 7 | 0.012995 | |
| 3 | C01Q1 | 1 | 1 | F1B09 | 1 | 9 | 0.005770 | |
| 4 | C01Q1 | 1 | 1 | F1B11 | 1 | 11 | 0.017934 | |

| | prices | sugar | mushy | ... | demand_instruments10 | demand_instruments11 | \ |
|---|----------|-------|-------|-----|----------------------|----------------------|---|
| 0 | 0.072088 | 2 | 1 | ... | 2.116358 | -0.154708 | |
| 1 | 0.114178 | 18 | 1 | ... | -7.374091 | -0.576412 | |
| 2 | 0.132391 | 4 | 1 | ... | 2.187872 | -0.207346 | |
| 3 | 0.130344 | 3 | 0 | ... | 2.704576 | 0.040748 | |
| 4 | 0.154823 | 12 | 0 | ... | 1.261242 | 0.034836 | |

| | demand_instruments12 | demand_instruments13 | demand_instruments14 | \ |
|---|----------------------|----------------------|----------------------|---|
| 0 | -0.005796 | 0.014538 | 0.126244 | |
| 1 | 0.012991 | 0.076143 | 0.029736 | |
| 2 | 0.003509 | 0.091781 | 0.163773 | |
| 3 | -0.003724 | 0.094732 | 0.135274 | |
| 4 | -0.000568 | 0.102451 | 0.130640 | |

| | demand_instruments15 | demand_instruments16 | demand_instruments17 | \ |
|---|----------------------|----------------------|----------------------|---|
| 0 | 0.067345 | 0.068423 | 0.034800 | |
| 1 | 0.087867 | 0.110501 | 0.087784 | |
| 2 | 0.111881 | 0.108226 | 0.086439 | |
| 3 | 0.088090 | 0.101767 | 0.101777 | |
| 4 | 0.084818 | 0.101075 | 0.125169 | |

| | demand_instruments18 | demand_instruments19 |
|---|----------------------|----------------------|
| 0 | 0.126346 | 0.035484 |
| 1 | 0.049872 | 0.072579 |
| 2 | 0.122347 | 0.101842 |
| 3 | 0.110741 | 0.104332 |
| 4 | 0.133464 | 0.121111 |

(continues on next page)

(continued from previous page)

```
[5 rows x 30 columns]
```

The product data contains `market_ids`, `product_ids`, `firm_ids`, `shares`, `prices`, a number of other IDs and product characteristics, and some pre-computed excluded `demand_instruments0`, `demand_instruments1`, and so on. The `product_ids` will be incorporated as fixed effects.

For more information about the instruments and the example data as a whole, refer to the [data](#) module.

Setting Up the Problem

We can combine the *Formulation* and `product_data` to construct a *Problem*. We pass the *Formulation* first and the `product_data` second. We can also display the properties of the problem by typing its name.

```
logit_formulation = pyblp.Formulation('prices', absorb='C(product_ids)')
logit_formulation

prices + Absorb[C(product_ids)]
```

```
problem = pyblp.Problem(logit_formulation, product_data)
problem
```

Dimensions:

```
=====
  T   N   F   K1  MD  ED
  ---
94  2256  5   1   20  1
=====
```

Formulations:

```
=====
      Column Indices:          0
-----
X1: Linear Characteristics  prices
=====
```

Two sets of properties are displayed:

1. Dimensions of the data.
2. Formulations of the problem.

The dimensions describe the shapes of matrices as laid out in *Notation*. They include:

- T is the number of markets.
- N is the length of the dataset (the number of products across all markets).
- F is the number of firms, which we won't use in this example.
- K_1 is the dimension of the linear demand parameters.
- M_D is the dimension of the instrument variables (excluded instruments and exogenous regressors).
- E_D is the number of fixed effect dimensions (one-dimensional fixed effects, two-dimensional fixed effects, etc.).

There is only a single *Formulation* for this model.

- X_1 is the linear component of utility for demand and depends only on prices (after the fixed effects are removed).

Solving the Problem

The `Problem.solve` method always returns a `ProblemResults` class, which can be used to compute post-estimation outputs. See the *post estimation* tutorial for more information.

```
logit_results = problem.solve()
logit_results

Problem Results Summary:
=====
GMM   Objective   Weighting Matrix
Step   Value       Condition Number
-----
  2    +8.3E-02      +5.7E+07
=====

Cumulative Statistics:
=====
Computation   Objective
  Time       Evaluations
-----
  00:00:00         2
=====

Beta Estimates (Robust SEs in Parentheses):
=====
prices
```

(continues on next page)

```
-----
-3.0E+01
(+1.0E+00)
=====
```

4.1.3 Theory of Nested Logit

We can extend the logit model to allow for correlation within a group h so that

$$U_{jti} = \alpha p_{jt} + x_{jt} \beta^x + \xi_{jt} + \bar{\epsilon}_{h(j)ti} + (1 - \rho) \bar{\epsilon}_{jti}. \quad (4.6)$$

Now, we require that $\epsilon_{jti} = \bar{\epsilon}_{h(j)ti} + (1 - \rho) \bar{\epsilon}_{jti}$ is distributed IID with the Type I Extreme Value (Gumbel) distribution. As $\rho \rightarrow 1$, all consumers stay within their group. As $\rho \rightarrow 0$, this collapses to the IIA logit. Note that if we wanted, we could allow ρ to differ between groups with the notation $\rho_{h(j)}$.

This gives us aggregate marketshares as the product of two logits, the within group logit and the across group logit:

$$s_{jt} = \frac{\exp[V_{jt}/(1 - \rho)]}{\exp[V_{h(j)t}/(1 - \rho)]} \cdot \frac{\exp V_{h(j)t}}{1 + \sum_h \exp V_{ht}}, \quad (4.7)$$

where $V_{jt} = \alpha p_{jt} + x_{jt} \beta^x + \xi_{jt}$.

After some work we again obtain the linear estimating equation:

$$\log s_{jt} - \log s_{0t} = \alpha p_{jt} + x_{jt} \beta^x + \rho \log s_{j|h(j)t} + \xi_{jt}, \quad (4.8)$$

where $s_{j|h(j)t} = s_{jt}/s_{h(j)t}$ and $s_{h(j)t}$ is the share of group h in market t . See *Berry (1994)* or *Cardell (1997)* for more information.

Again, the left hand side is data, though the $\ln s_{j|h(j)t}$ is clearly endogenous which means we must instrument for it. Rather than include $\ln s_{j|h(j)t}$ along with the linear components of utility, X_1 , whenever `nesting_ids` are included in `product_data`, ρ is treated as a nonlinear X_2 parameter. This means that the linear component is given instead by

$$\log s_{jt} - \log s_{0t} - \rho \log s_{j|h(j)t} = \alpha p_{jt} + x_{jt} \beta^x + \xi_{jt}. \quad (4.9)$$

This is done for two reasons:

1. It forces the user to treat ρ as an endogenous parameter.
2. It extends much more easily to the RCNL model of *Brenkers and Verboven (2006)*.

A common choice for an additional instrument is the number of products per nest.

4.1.4 Application of Nested Logit

By including `nesting_ids` (another reserved name) as a field in `product_data`, we tell the package to estimate a nested logit model, and we don't need to change any of the formulas. We show how to construct the category groupings in two different ways:

1. We put all products in a single nest (only the outside good in the other nest).
2. We put products into two nests (either mushy or non-mushy).

We also construct an additional instrument based on the number of products per nest. Typically this is useful as a source of exogenous variation in the within group share $\ln s_{j|h(j)t}$. However, in this example because the number of products per nest does not vary across markets, if we include product fixed effects, this instrument is irrelevant.

We'll define a function that constructs the additional instrument and solves the nested logit problem. We'll exclude product ID fixed effects, which are collinear with `mushy`, and we'll choose $\rho = 0.7$ as the initial value at which the optimization routine will start.

```
def solve_nl(df):
    groups = df.groupby(['market_ids', 'nesting_ids'])
    df['demand_instruments20'] = groups['shares'].transform(np.size)
    nl_formulation = pyblp.Formulation('0 + prices')
    problem = pyblp.Problem(nl_formulation, df)
    return problem.solve(rho=0.7)
```

First, we'll solve the problem when there's a single nest for all products, with the outside good in its own nest.

```
df1 = product_data.copy()
df1['nesting_ids'] = 1
nl_results1 = solve_nl(df1)
nl_results1
```

Problem Results Summary:

```
=====
GMM   Objective   Projected   Reduced   Weighting Matrix   Covariance Matrix
Step   Value         Gradient   Norm      Hessian           Condition Number   Condition Number
-----
  2    +9.0E-02     +7.6E-12   +4.8E+00   +2.0E+09           +3.0E+04
=====
```

Cumulative Statistics:

```
=====
Computation   Optimization   Objective
  Time         Iterations     Evaluations
-----
```

(continues on next page)

```

00:00:04      3      8
=====
Rho Estimates (Robust SEs in Parentheses):
=====
All Groups
-----
+9.8E-01
(+1.4E-02)
=====
Beta Estimates (Robust SEs in Parentheses):
=====
prices
-----
-1.2E+00
(+4.0E-01)
=====

```

When we inspect the *Problem*, the only changes from the plain logit model is the additional instrument that contributes to M_D and the inclusion of H , the number of nesting categories.

```

nl_results1.problem
Dimensions:
=====
  T      N      F      K1      MD      H
  ---      ---      ---      ---      ---      ---
 94     2256     5      1      21      1
=====
Formulations:
=====
      Column Indices:      0
-----
X1: Linear Characteristics prices
=====

```

Next, we'll solve the problem when there are two nests for mushy and non-mushy.

```
df2 = product_data.copy()
df2['nesting_ids'] = df2['mushy']
nl_results2 = solve_n1(df2)
nl_results2
```

Problem Results Summary:

```
=====
GMM   Objective   Projected   Reduced   Weighting Matrix   Covariance Matrix
Step  Value       Gradient Norm Hessian   Condition Number   Condition Number
-----
  2    +3.1E-01    +2.4E-12   +2.5E+00   +5.1E+08           +2.0E+04
=====
```

Cumulative Statistics:

```
=====
Computation   Optimization   Objective
  Time         Iterations     Evaluations
-----
00:00:04         3             8
=====
```

Rho Estimates (Robust SEs in Parentheses):

```
=====
All Groups
-----
+8.9E-01
(+1.9E-02)
=====
```

Beta Estimates (Robust SEs in Parentheses):

```
=====
prices
-----
-7.8E+00
(+4.8E-01)
=====
```

For both cases we find that $\hat{\rho} > 0.8$.

Finally, we'll also look at the adjusted parameter on prices, $\alpha/(1 - \rho)$.


```
nl_results1.beta[0] / (1 - nl_results1.rho)
array([[ -67.39338889]])
```

```
nl_results2.beta[0] / (1 - nl_results2.rho)
array([[ -72.27074638]])
```

Treating Within Group Shares as Exogenous

The package is designed to prevent the user from treating the within group share, $\log s_{j|h(j)t}$, as an exogenous variable. For example, if we were to compute a `group_share` variable and use the algebraic functionality of *Formulation* by including the expression `log(shares / group_share)` in our formula for X_1 , the package would raise an error because the package knows that `shares` should not be included in this formulation.

To demonstrate why this is a bad idea, we'll override this feature by calculating $\log s_{j|h(j)t}$ and including it as an additional variable in X_1 . To do so, we'll first re-define our function for setting up and solving the nested logit problem.

```
def solve_nl2(df):
    groups = df.groupby(['market_ids', 'nesting_ids'])
    df['group_share'] = groups['shares'].transform(np.sum)
    df['within_share'] = df['shares'] / df['group_share']
    df['demand_instruments20'] = groups['shares'].transform(np.size)
    nl2_formulation = pyblp.Formulation('0 + prices + log(within_share)')
    problem = pyblp.Problem(nl2_formulation, df.drop(columns=['nesting_ids']))
    return problem.solve()
```

Again, we'll solve the problem when there's a single nest for all products, with the outside good in its own nest.

```
nl2_results1 = solve_nl2(df1)
nl2_results1
```

Problem Results Summary:

```
=====
GMM   Objective  Weighting Matrix  Covariance Matrix
Step   Value      Condition Number  Condition Number
-----
  2    +9.0E-02      +2.1E+09          +1.1E+04
=====
```

Cumulative Statistics:

(continues on next page)

(continued from previous page)

```

=====
Computation   Objective
   Time       Evaluations
-----
00:00:00         2
=====

```

Beta Estimates (Robust SEs in Parentheses):

```

=====
prices      log(within_share)
-----
-1.0E+00    +9.9E-01
(+2.4E-01)  (+7.9E-03)
=====

```

And again, we'll solve the problem when there are two nests for mushy and non-mushy.

```

nl2_results2 = solve_nl2(df2)
nl2_results2

```

Problem Results Summary:

```

=====
GMM   Objective   Weighting Matrix   Covariance Matrix
Step   Value         Condition Number   Condition Number
-----
2      +3.1E-01       +5.5E+08           +7.7E+03
=====

```

Cumulative Statistics:

```

=====
Computation   Objective
   Time       Evaluations
-----
00:00:00         2
=====

```

Beta Estimates (Robust SEs in Parentheses):

```

=====
prices      log(within_share)
-----
-6.8E+00    +9.3E-01

```

(continues on next page)

(continued from previous page)

```
(+2.9E-01)      (+1.1E-02)
=====
```

One can observe that we obtain parameter estimates which are quite different than above.

```
n12_results1.beta[0] / (1 - n12_results1.beta[1])
```

```
array([-86.37368447])
```

```
n12_results2.beta[0] / (1 - n12_results2.beta[1])
```

```
array([-100.14496892])
```

The [online version](#) of the following section may be easier to read.

4.2 Random Coefficients Logit Tutorial with the Fake Cereal Data

```
import pyblp
import numpy as np
import pandas as pd

pyblp.options.digits = 2
pyblp.options.verbose = False
pyblp.__version__

'0.8.1'
```

In this tutorial, we'll use data from *Nevo (2000)* to solve the paper's fake cereal problem. Locations of CSV files that contain the data are in the `data` module.

4.2.1 Theory of Random Coefficients Logit

The random coefficients model extends the plain logit model by allowing for correlated tastes for different product characteristics. In this model (indirect) utility is given by

$$u_{jti} = \alpha_i p_{jt} + x_{jt} \beta_i^x + \xi_{jt} + \epsilon_{jti} \quad (4.10)$$

The main addition is that $\beta_i = (\alpha_i, \beta_i^x)$ have individual specific subscripts i .

Conditional on β_i , the individual marketshares follow the same logit form as before. But now we must integrate over heterogeneous individuals to get the aggregate marketshares:

$$s_{jt}(\alpha, \beta, \theta) = \int \frac{\exp(\alpha_i p_{jt} + x_{jt} \beta_i^x + \xi_{jt})}{1 + \sum_k \exp(\alpha_i p_{jt} + x_{kt} \beta_i^x + \xi_{kt})} f(\alpha_i, \beta_i | \theta). \quad (4.11)$$

In general, this integral needs to be calculated numerically. This also means that we can't directly linearize the model. It is common to re-parametrize the model to separate the aspects of mean utility that all individuals agree on, $\delta_{jt} = \alpha p_{jt} + x_{jt} \beta^x + \xi_{jt}$, from the individual specific heterogeneity, $\mu_{jti}(\theta)$. This gives us

$$s_{jt}(\delta_{jt}, \theta) = \int \frac{\exp(\delta_{jt} + \mu_{jti})}{1 + \sum_k \exp(\delta_{kt} + \mu_{kti})} f(\mu_{it} | \theta). \quad (4.12)$$

Given a guess of θ we can solve the system of nonlinear equations for the vector δ which equates observed and predicted marketshares $s_{jt} = s_{jt}(\delta, \theta)$. Now we can perform a linear IV GMM regression of the form

$$\delta_{jt}(\theta) = \alpha p_{jt} + x_{jt} \beta^x + \xi_{jt}. \quad (4.13)$$

The moments are constructed by interacting the predicted residuals $\hat{\xi}_{jt}(\theta)$ with instruments z_{jt} to form

$$\bar{g}(\theta) = \frac{1}{N} \sum_{j,t} z'_{jt} \hat{\xi}_{jt}(\theta). \quad (4.14)$$

4.2.2 Application of Random Coefficients Logit with the Fake Cereal Data

To include random coefficients we need to add a *Formulation* for the nonlinear characteristics X_2 .

Just like in the logit case we have the same reserved field names in `product_data`:

- `market_ids` are the unique market identifiers which we subscript t .
- `shares` specifies the marketshares which need to be between zero and one, and within a market ID, $\sum_j s_{jt} < 1$.
- `prices` are prices p_{jt} . These have some special properties and are *always* treated as endogenous.
- `demand_instruments0`, `demand_instruments1`, and so on are numbered demand instruments. These represent only the *excluded* instruments. The exogenous regressors in X_1 (of which X_2 is typically a subset) will be automatically added to the set of instruments.

We proceed with the following steps:

1. Load the `product_data` which at a minimum consists of `market_ids`, `shares`, `prices`, and at least a single column of demand instruments, `demand_instruments0`.
2. Define a *Formulation* for the X_1 (linear) demand model.
 - This and all other formulas are similar to R and `patsy` formulas.
 - It includes a constant by default. To exclude the constant, specify either a 0 or a `-1`.
 - To efficiently include fixed effects, use the `absorb` option and specify which categorical variables you would like to absorb.
 - Some model reduction may happen automatically. The constant will be excluded if you include fixed effects and some precautions are taken against collinearity. However, you will have to make sure that differently-named variables are not collinear.
3. Define a *Formulation* for the X_2 (nonlinear) demand model.
 - Include only the variables over which we want random coefficients.
 - Do not absorb or include fixed effects.

- It will include a random coefficient on the constant (to capture inside good vs. outside good preference) unless you specify not to with a 0 or a -1.
4. Define an *Integration* configuration to solve the market share integral from several available options:
 - Monte Carlo integration (pseudo-random draws).
 - Product rule quadrature.
 - Sparse grid quadrature.
 5. Combine *Formulation* classes, `product_data`, and the *Integration* configuration to construct a *Problem*.
 6. Use the *Problem.solve* method to estimate parameters.
 - It is required to specify an initial guess of the nonlinear parameters. This serves two primary purposes: speeding up estimation and indicating to the solver through initial values of zero which parameters are restricted to be always zero.

Specification of Random Taste Parameters

It is common to assume that $f(\beta_i | \theta)$ follows a multivariate normal distribution, and to break it up into three parts:

1. A mean $K_1 \times 1$ taste which all individuals agree on, β .
2. A $K_2 \times K_2$ covariance matrix, Σ .
3. Any $K_2 \times D$ interactions, Π , with observed $D \times 1$ demographic data, d_i .

Together this gives us that

$$\beta_i \sim N(\beta + \Pi d_i, \Sigma). \quad (4.15)$$

Problem.solve takes an initial guess Σ_0 of Σ . It guarantees that $\hat{\Sigma}$ (the estimated parameters) will have the same sparsity structure as Σ_0 . So any zero element of Σ is restricted to be zero in the solution $\hat{\Sigma}$. For example, a popular restriction is that Σ is diagonal, this can be achieved by passing a diagonal matrix as Σ_0 .

As is common with multivariate normal distributions, Σ is not estimated directly. Rather, its matrix square (Cholesky) root L is estimated where $LL' = \Sigma$.

Loading the Data

The `product_data` argument of *Problem* should be a structured array-like object with fields that store data. Product data can be a structured NumPy array, a pandas DataFrame, or other similar objects.

```
product_data = pd.read_csv(pyblp.data.NEVO_PRODUCTS_LOCATION)
product_data.head()
```

```

market_ids  city_ids  quarter  product_ids  firm_ids  brand_ids  shares  \
0  C01Q1      1      1      F1B04      1          4  0.012417
1  C01Q1      1      1      F1B06      1          6  0.007809
2  C01Q1      1      1      F1B07      1          7  0.012995
3  C01Q1      1      1      F1B09      1          9  0.005770
4  C01Q1      1      1      F1B11      1         11  0.017934

prices  sugar  mushy  ...  demand_instruments10  demand_instruments11  \
0  0.072088    2    1  ...          2.116358          -0.154708
1  0.114178   18    1  ...         -7.374091          -0.576412
2  0.132391    4    1  ...          2.187872          -0.207346
3  0.130344    3    0  ...          2.704576           0.040748
4  0.154823   12    0  ...          1.261242           0.034836

demand_instruments12  demand_instruments13  demand_instruments14  \
0          -0.005796          0.014538          0.126244
1           0.012991          0.076143          0.029736
2           0.003509          0.091781          0.163773
3          -0.003724          0.094732          0.135274
4          -0.000568          0.102451          0.130640

demand_instruments15  demand_instruments16  demand_instruments17  \
0           0.067345          0.068423          0.034800
1           0.087867          0.110501          0.087784
2           0.111881          0.108226          0.086439
3           0.088090          0.101767          0.101777
4           0.084818          0.101075          0.125169

demand_instruments18  demand_instruments19
0           0.126346          0.035484
1           0.049872          0.072579
2           0.122347          0.101842
3           0.110741          0.104332
4           0.133464          0.121111

[5 rows x 30 columns]

```

The product data contains `market_ids`, `product_ids`, `firm_ids`, `shares`, `prices`, a number of other firm IDs and product characteristics, and some pre-computed excluded `demand_instruments0`, `demand_instruments1`, and so on. The `product_ids` will be incorporated as fixed effects.

For more information about the instruments and the example data as a whole, refer to the [data](#) module.

Setting Up and Solving the Problem Without Demographics

Formulations, product data, and an integration configuration are collectively used to initialize a *Problem*. Once initialized, *Problem.solve* runs the estimation routine. The arguments to *Problem.solve* configure how estimation is performed. For example, *optimization* and *iteration* arguments configure the optimization and iteration routines that are used by the outer and inner loops of estimation.

We'll specify *Formulation* configurations for X_1 , the linear characteristics, and X_2 , the nonlinear characteristics.

- The formulation for X_1 consists of `prices` and fixed effects constructed from `product_ids`, which we will absorb using `absorb` argument of *Formulation*.
- If we were interested in reporting estimates for each fixed effect, we could replace the formulation for X_1 with `Formulation('prices + C(product_ids)')`.
- Because `sugar`, `mushy`, and the constant are collinear with `product_ids`, we can include them in X_2 but not in X_1 .

```
X1_formulation = pyblp.Formulation('0 + prices', absorb='C(product_ids)')
X2_formulation = pyblp.Formulation('1 + prices + sugar + mushy')
product_formulations = (X1_formulation, X2_formulation)
product_formulations
```

```
(prices + Absorb[C(product_ids)], 1 + prices + sugar + mushy)
```

We also need to specify an *Integration* configuration. We consider two alternatives:

1. Monte Carlo draws: we simulate 50 individuals from a random normal distribution.
2. Product rules: we construct nodes and weights according to a product rule that exactly integrates polynomials of degree $5 \times 2 - 1 = 9$ or less.

```
mc_integration = pyblp.Integration('monte_carlo', size=50, seed=0)
mc_integration
```

```
Configured to construct nodes and weights with Monte Carlo simulation.
```

```
pr_integration = pyblp.Integration('product', size=5)
pr_integration
```

```
Configured to construct nodes and weights according to the level-5 Gauss-Hermite product rule.
```

```
mc_problem = pyblp.Problem(product_formulations, product_data, integration=mc_integration)
mc_problem
```

```
Dimensions:
```

```
=====
```

(continues on next page)

(continued from previous page)

```

T      N      F      I      K1      K2      MD      ED
-----
94    2256    5     4700    1       4       20      1
=====

Formulations:
=====
      Column Indices:      0      1      2      3
-----
X1: Linear Characteristics      prices
X2: Nonlinear Characteristics    1      prices  sugar  mushy
=====

```

```

pr_problem = pyblp.Problem(product_formulations, product_data, integration=pr_integration)
pr_problem

```

```

Dimensions:
=====
T      N      F      I      K1      K2      MD      ED
-----
94    2256    5     58750    1       4       20      1
=====

Formulations:
=====
      Column Indices:      0      1      2      3
-----
X1: Linear Characteristics      prices
X2: Nonlinear Characteristics    1      prices  sugar  mushy
=====

```

As an illustration of how to configure the optimization routine, we'll use a simpler, non-default *Optimization* configuration that doesn't support parameter bounds.

```

bfgs = pyblp.Optimization('bfgs', {'gtol': 1e-10})
bfgs

```

```

Configured to optimize using the BFGS algorithm implemented in SciPy with analytic gradients and options {gtol: +1.0E-
↪10}.

```

We estimate three versions of the model:

1. An unrestricted covariance matrix for random tastes using Monte Carlo integration.

2. An unrestricted covariance matrix for random tastes using the product rule.
3. A restricted diagonal matrix for random tastes using Monte Carlo Integration.

Notice that the only thing that changes when we estimate the restricted covariance is our initial guess of Σ_0 . The upper diagonal in this initial guess is ignored because we are optimizing over the lower-triangular Cholesky root of Σ .

```
results1 = mc_problem.solve(sigma=np.ones((4, 4)), optimization=bfgs)
results1
```

Problem Results Summary:

```
=====
GMM   Objective  Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step  Value       Norm         Min Eigenvalue Max Eigenvalue  Condition Number  Condition Number
-----
  2    +6.1E-02   +3.8E-10     +2.3E-05     +2.7E+00       +5.6E+07          +1.8E+06
=====
```

Cumulative Statistics:

```
=====
Computation  Optimization  Objective  Fixed Point  Contraction
  Time       Iterations    Evaluations Iterations  Evaluations
-----
 00:05:12    180          350        215888     675540
=====
```

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

```
=====
Sigma:      1          prices      sugar      mushy
-----
 1          +1.6E-01
           (+5.2E+00)

prices     +3.7E+00   -8.9E+00
           (+4.2E+01)  (+1.8E+01)

sugar      -7.6E-02   +1.5E-01   +6.4E-02
           (+1.6E-01)  (+2.3E-01) (+1.4E-01)

mushy      -1.9E-01   -6.5E-02   -7.8E-02   +4.8E-01
           (+2.1E+00)  (+1.4E+00) (+1.8E+00) (+2.9E+00)
=====
```

(continues on next page)

(continued from previous page)

Beta Estimates (Robust SEs in Parentheses):

```
=====
prices
-----
-3.1E+01
(+5.6E+00)
=====
```

```
results2 = pr_problem.solve(sigma=np.ones((4, 4)), optimization=bfgs)
results2
```

Problem Results Summary:

```
=====
GMM   Objective   Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step  Value        Norm          Min Eigenvalue Max Eigenvalue  Condition Number  Condition Number
-----
  2    +7.0E-02    +5.6E-12      +7.0E-06      +2.3E+00        +5.3E+07          +1.4E+37
=====
```

Cumulative Statistics:

```
=====
Computation  Optimization  Objective  Fixed Point  Contraction
Time         Iterations    Evaluations Iterations  Evaluations
-----
00:06:59    182           321        182808      572217
=====
```

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

```
=====
Sigma:      1      prices      sugar      mushy
-----
  1      -7.4E-01
          (+2.0E+00)

prices     +1.3E+01    +1.0E-10
          (+6.1E+00)  (+8.2E-10)

sugar     -1.1E-01    -9.9E-12    -1.4E-11
          (+9.5E-02)  (+6.2E-08)  (+1.1E-07)
```

(continues on next page)

(continued from previous page)

```

mushy    +1.5E-01    +1.7E-10    +3.1E-10    -5.4E-10
          (+6.8E-01) (+1.1E-10) (+8.4E-10) (+8.6E-09)
=====

```

Beta Estimates (Robust SEs in Parentheses):

```

=====
prices
-----
-3.1E+01
(+3.9E+00)
=====

```

```

results3 = mc_problem.solve(sigma=np.eye(4), optimization=bfgs)
results3

```

Problem Results Summary:

```

=====
GMM   Objective  Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step  Value       Norm         Min Eigenvalue Max Eigenvalue  Condition Number  Condition Number
-----
  2   +7.9E-02   +3.5E-10     +4.8E-04     +2.7E+00       +5.8E+07          +4.8E+04
=====

```

Cumulative Statistics:

```

=====
Computation  Optimization  Objective  Fixed Point  Contraction
  Time       Iterations    Evaluations Iterations  Evaluations
-----
00:00:57         61          147         66429      209785
=====

```

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

```

=====
Sigma:      1          prices      sugar      mushy
-----
  1         +5.2E-02
          (+1.1E+00)

prices     +0.0E+00    -4.3E-01
          (+8.0E+00)

```

(continues on next page)

(continued from previous page)

```
sugar    +0.0E+00    +0.0E+00    +3.6E-02
          (+5.8E-02)

mushy    +0.0E+00    +0.0E+00    +0.0E+00    +5.0E-01
          (+1.7E+00)
```

```
=====
Beta Estimates (Robust SEs in Parentheses):
=====
```

```
prices
-----
-3.0E+01
(+1.4E+00)
=====
```

We see that all three models give similar estimates of the price coefficient $\hat{\alpha} \approx -30$. Note a few of the estimated terms on the diagonal of Σ are negative. Since the diagonal consists of standard deviations, negative values are unrealistic. When using another optimization routine that supports bounds (like the default L-BFGS-B routine), these diagonal elements are by default bounded from below by zero.

Adding Demographics to the Problem

To add demographic data we need to make two changes:

1. We need to load `agent_data`, which for this cereal problem contains pre-computed Monte Carlo draws and demographics.
2. We need to add an `agent_formulation` to the model.

The `agent_data` has several reserved column names.

- `market_ids` are the index that link the `agent_data` to the `market_ids` in `product_data`.
- `weights` are the weights w_{it} attached to each agent. In each market, these should sum to one so that $\sum_i w_{it} = 1$. It is often the case the $w_{it} = 1/I_t$ where I_t is the number of agents in market t , so that each agent gets equal weight. Other possibilities include quadrature nodes and weights.
- `nodes0`, `nodes1`, and so on are the nodes at which the unobserved agent tastes μ_{jti} are evaluated. The nodes should be labeled from $0, \dots, K_2 - 1$ where K_2 is the number of random coefficients.
- Other fields are the realizations of the demographics d themselves.

```
agent_data = pd.read_csv(pyblp.data.NEVO_AGENTS_LOCATION)
agent_data.head()
```

| | market_ids | city_ids | quarter | weights | nodes0 | nodes1 | nodes2 | \ |
|---|------------|----------|---------|---------|-----------|-----------|-----------|---|
| 0 | C01Q1 | 1 | 1 | 0.05 | 0.434101 | -1.500838 | -1.151079 | |
| 1 | C01Q1 | 1 | 1 | 0.05 | -0.726649 | 0.133182 | -0.500750 | |
| 2 | C01Q1 | 1 | 1 | 0.05 | -0.623061 | -0.138241 | 0.797441 | |
| 3 | C01Q1 | 1 | 1 | 0.05 | -0.041317 | 1.257136 | -0.683054 | |
| 4 | C01Q1 | 1 | 1 | 0.05 | -0.466691 | 0.226968 | 1.044424 | |

| | nodes3 | income | income_squared | age | child |
|---|-----------|-----------|----------------|-----------|-----------|
| 0 | 0.161017 | 0.495123 | 8.331304 | -0.230109 | -0.230851 |
| 1 | 0.129732 | 0.378762 | 6.121865 | -2.532694 | 0.769149 |
| 2 | -0.795549 | 0.105015 | 1.030803 | -0.006965 | -0.230851 |
| 3 | 0.259044 | -1.485481 | -25.583605 | -0.827946 | 0.769149 |
| 4 | 0.092019 | -0.316597 | -6.517009 | -0.230109 | -0.230851 |

The agent formulation tells us which columns of demographic information to interact with X_2 .

```
agent_formulation = pyblp.Formulation('0 + income + income_squared + age + child')
agent_formulation
income + income_squared + age + child
```

This tells us to include demographic realizations for income, income_squared, age, and the presence of children, child, but to ignore other possible demographics when interacting demographics with X_2 . We should also generally exclude the constant from the demographic formula.

Now we configure the *Problem* to include the agent_formulation and agent_data, which follow the product_formulations and product_data.

When we display the class, it lists the demographic interactions in the table of formulations and reports $D = 4$, the dimension of the demographic draws.

```
nevo_problem = pyblp.Problem(
    product_formulations,
    product_data,
    agent_formulation,
    agent_data
)
nevo_problem
```

Dimensions:

| T | N | F | I | K1 | K2 | D | MD | ED |
|----|------|---|------|----|----|---|----|----|
| 94 | 2256 | 5 | 1880 | 1 | 4 | 4 | 20 | 1 |

(continues on next page)

(continued from previous page)

Formulations:

```

=====
      Column Indices:      0          1          2          3
-----
X1: Linear Characteristics  prices
X2: Nonlinear Characteristics  1          prices          sugar  mushy
      d: Demographics          income  income_squared  age  child
=====

```

The initialized problem can be solved with *Problem.solve*. We'll use the same starting values as *Nevo (2000)*. By passing a diagonal matrix as starting values for Σ , we're choosing to ignore covariance terms. Similarly, zeros in the starting values for Π mean that those parameters will be fixed at zero.

To replicate common estimates, we'll use the non-default BFGS optimization routine, and we'll configure *Problem.solve* to use 1-step GMM instead of 2-step GMM.

```

initial_sigma = np.diag([0.3302, 2.4526, 0.0163, 0.2441])
initial_pi = np.array([
    [ 5.4819, 0, 0.2037, 0 ],
    [15.8935, -1.2000, 0, 2.6342],
    [-0.2506, 0, 0.0511, 0 ],
    [ 1.2650, 0, -0.8091, 0  ]
])
nevo_results = nevo_problem.solve(
    initial_sigma,
    initial_pi,
    optimization='bfgs',
    method='1s'
)
nevo_results

```

Problem Results Summary:

```

=====
GMM   Objective  Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step   Value      Norm      Min Eigenvalue  Max Eigenvalue  Condition Number  Condition Number
-----
  1    +2.0E-03   +5.4E-11   +1.3E-08      +7.3E+00      +6.9E+07          +8.4E+08
=====

```

Cumulative Statistics:

```

=====

```

(continues on next page)

| Computation Time | Optimization Iterations | Objective Evaluations | Fixed Point Iterations | Contraction Evaluations |
|------------------|-------------------------|-----------------------|------------------------|-------------------------|
| 00:01:27 | 140 | 145 | 109732 | 341137 |

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

| Sigma: | | | | | Pi: | | | | |
|--------|------------------------|------------------------|------------------------|------------------------|--------|------------------------|------------------------|------------------------|------------------------|
| | 1 | prices | sugar | mushy | | income | income_squared | age | child |
| 1 | +5.6E-01 (+1.6E-01) | | | | 1 | +2.3E+00 (+1.2E+00) | +0.0E+00 | +1.3E+00 (+6.3E-01) | +0.0E+00 |
| prices | +0.0E+00 | +3.3E+00 (+1.3E+00) | | | prices | +5.9E+02 (+2.7E+02) | -3.0E+01 (+1.4E+01) | +0.0E+00 | +1.1E+01 (+4.1E+00) |
| sugar | +0.0E+00 | +0.0E+00 | -5.8E-03 (+1.4E-02) | | sugar | -3.8E-01 (+1.2E-01) | +0.0E+00 | +5.2E-02 (+2.6E-02) | +0.0E+00 |
| mushy | +0.0E+00 | +0.0E+00 | +0.0E+00 | +9.3E-02 (+1.9E-01) | mushy | +7.5E-01 (+8.0E-01) | +0.0E+00 | -1.4E+00 (+6.7E-01) | +0.0E+00 |

Beta Estimates (Robust SEs in Parentheses):

```
=====
prices
-----
-6.3E+01
(+1.5E+01)
=====
```

Results are similar to those in the original paper with a (non-scaled) objective value of $q(\hat{\theta}) = 8.96E-7 = 4.65/N^2$ and a price coefficient of $\hat{\alpha} = -62.7$.

Restricting Parameters

Because the interactions between price, income, and income_squared are potentially collinear, we might worry that $\hat{\Pi}_{21} = 588$ and $\hat{\Pi}_{22} = -30.2$ are pushing the price coefficient in opposite directions. Both are large in magnitude but statistically insignificant. One way of dealing with this is to restrict $\Pi_{22} = 0$.

There are two ways we can do this:

1. Change the initial Π_0 values to make this term zero.
2. Change the agent formula to drop `income_squared`.

First, we'll change the initial Π_0 values.

```
restricted_pi = initial_pi.copy()
restricted_pi[1, 1] = 0
nevo_problem.solve(
    initial_sigma,
    restricted_pi,
    optimization=bfgs,
    method='1s'
)
```

Problem Results Summary:

```
=====
GMM   Objective  Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step  Value       Norm         Min Eigenvalue Max Eigenvalue  Condition Number  Condition Number
-----
1     +6.8E-03    +1.8E-09     +2.1E-05     +7.3E+00       +6.9E+07         +5.7E+05
=====
```

Cumulative Statistics:

```
=====
Computation  Optimization  Objective  Fixed Point  Contraction
  Time       Iterations    Evaluations Iterations    Evaluations
-----
00:01:48      112          208        155160       482159
=====
```

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

```
=====
Sigma:      1      prices      sugar      mushy      |  Pi:      income      income_squared      age      child
-----
1          +3.7E-01      |  1          +3.1E+00      +0.0E+00      +1.2E+00      +0.0E+00
          (+1.2E-01)      |
prices     +0.0E+00      +1.8E+00      |  prices     +4.2E+00      +0.0E+00      +0.0E+00      +1.2E+01
          (+9.2E-01)      |          (+4.6E+00)      (+5.2E+00)
sugar      +0.0E+00      +0.0E+00      -4.4E-03      |  sugar      -1.9E-01      +0.0E+00      +2.8E-02      +0.0E+00
          (+1.2E-02)      |          (+3.5E-02)      (+3.2E-02)
=====
```

(continues on next page)

(continued from previous page)

```

mushy      +0.0E+00      +0.0E+00      +0.0E+00      +8.6E-02      |      mushy      +1.5E+00      +0.0E+00      -1.5E+00      +0.0E+00
              (+1.9E-01)      |              (+6.5E-01)              (+1.1E+00)
=====
Beta Estimates (Robust SEs in Parentheses):
=====
prices
-----
-3.2E+01
(+2.3E+00)
=====

```

Now we'll drop both `income_squared` and the corresponding column in Π_0 .

```

restricted_formulation = pyblp.Formulation('0 + income + age + child')
deleted_pi = np.delete(initial_pi, 1, axis=1)
restricted_problem = pyblp.Problem(
    product_formulations,
    product_data,
    restricted_formulation,
    agent_data
)
restricted_problem.solve(
    initial_sigma,
    deleted_pi,
    optimization='bfgs',
    method='ls'
)

```

Problem Results Summary:

```

=====
GMM   Objective  Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step   Value      Norm      Min Eigenvalue  Max Eigenvalue  Condition Number  Condition Number
-----
1     +6.8E-03  +2.0E-09  +2.1E-05      +7.3E+00      +6.9E+07      +5.7E+05
=====

```

Cumulative Statistics:

```

=====
Computation  Optimization  Objective  Fixed Point  Contraction

```

(continues on next page)

(continued from previous page)

| Time | Iterations | Evaluations | Iterations | Evaluations |
|----------|------------|-------------|------------|-------------|
| 00:01:28 | 111 | 157 | 117591 | 365555 |

=====

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

=====

| Sigma: | 1 | prices | sugar | mushy | Pi: | income | age | child |
|--------|------------------------|------------------------|------------------------|------------------------|--------|------------------------|------------------------|------------------------|
| 1 | +3.7E-01 (+1.2E-01) | | | | 1 | +3.1E+00 (+1.1E+00) | +1.2E+00 (+1.0E+00) | +0.0E+00 |
| prices | +0.0E+00 | +1.8E+00 (+9.2E-01) | | | prices | +4.2E+00 (+4.6E+00) | +0.0E+00 | +1.2E+01 (+5.2E+00) |
| sugar | +0.0E+00 | +0.0E+00 | -4.4E-03 (+1.2E-02) | | sugar | -1.9E-01 (+3.5E-02) | +2.8E-02 (+3.2E-02) | +0.0E+00 |
| mushy | +0.0E+00 | +0.0E+00 | +0.0E+00 | +8.6E-02 (+1.9E-01) | mushy | +1.5E+00 (+6.5E-01) | -1.5E+00 (+1.1E+00) | +0.0E+00 |

=====

Beta Estimates (Robust SEs in Parentheses):

=====

| prices |
|------------------------|
| -3.2E+01 (+2.3E+00) |

=====

The parameter estimates and standard errors are identical for both approaches. Based on the number of fixed point iterations, there is some evidence that the solver took a slightly different path for each problem, but both restricted problems arrived at identical answers. The $\hat{\Pi}_{12}$ interaction term is still insignificant.

The [online version](#) of the following section may be easier to read.

4.3 Random Coefficients Logit Tutorial with the Automobile Data

```
import pyblp
import numpy as np
import pandas as pd

pyblp.options.digits = 2
pyblp.options.verbose = False
pyblp.__version__

'0.8.1'
```

In this tutorial, we'll use data from *Berry, Levinsohn, and Pakes (1995)* to solve the paper's automobile problem.

4.3.1 Application of Random Coefficients Logit with the Automobile Data

This tutorial is similar to the *fake cereal tutorial*, but exhibits some other features of pyblp:

- Incorporating a supply side into demand estimation.
- Allowing for simple price-income demographic effects.
- Calculating clustered standard errors.

Loading the Data

We'll use `pandas` to load two sets of data:

1. `product_data`, which contains prices, shares, and other product characteristics.
2. `agent_data`, which contains draws from the distribution of heterogeneity.

```
product_data = pd.read_csv(pyblp.data.BLP_PRODUCTS_LOCATION)
product_data.head()
```

```

market_ids clustering_ids car_ids firm_ids region shares prices \
0 1971 AMGREM71 129 15 US 0.001051 4.935802
1 1971 AMHORN71 130 15 US 0.000670 5.516049
2 1971 AMJAVL71 132 15 US 0.000341 7.108642
3 1971 AMMATA71 134 15 US 0.000522 6.839506
4 1971 AMAMBS71 136 15 US 0.000442 8.928395

hpwt air mpd ... supply_instruments11 supply_instruments12 \
0 0.528997 0 1.888146 ... 2.024221 -0.251338
1 0.494324 0 1.935989 ... 2.026706 -0.205305
2 0.467613 0 1.716799 ... 1.882594 -0.417634
3 0.426540 0 1.687871 ... 1.846079 -0.445630
4 0.452489 0 1.504286 ... 1.760927 -0.610203

supply_instruments13 supply_instruments14 supply_instruments15 \
0 0.741272 0.820281 3.691881
1 0.675468 1.004707 3.628727
2 0.946640 0.429942 4.001801
3 0.982922 0.392283 4.061161
4 1.201745 -0.085373 4.340853

supply_instruments16 supply_instruments17 supply_instruments18 \
0 0.566070 -2.328128 0.385197
1 0.605683 -2.268975 0.363954
2 0.253461 -2.581003 -0.113763
3 0.183331 -2.626074 -0.253065
4 -0.043107 -2.785686 -0.556135

supply_instruments19 supply_instruments20
0 0.648526 1.052606
1 0.643205 0.906429
2 0.770387 1.511333
3 0.802037 1.583208
4 0.931611 2.177230

[5 rows x 47 columns]

```

The `product_data` contains market IDs, product IDs, firm IDs, shares, prices, a number of product characteristics, and some pre-computed excluded instruments constructed with principal component analysis. The product IDs are called `clustering_ids` because they will be used to compute clustered standard errors. For more information about the instruments and the example data as a whole, refer to the [data](#) module.

The `agent_data` contains market IDs, integration weights w_{it} , integration nodes ν_{it} , and demographics d_{it} . Here we use $I_t = 200$ equally weighted Monte

Carlo draws per market.

In non-example problems, it is usually a better idea to use many more draws, or a more sophisticated *Integration* configuration such as sparse grid quadrature.

```
agent_data = pd.read_csv(pyblp.data.BLP_AGENTS_LOCATION)
agent_data.head()
```

| | market_ids | weights | nodes0 | nodes1 | nodes2 | nodes3 | nodes4 | \ |
|---|------------|---------|----------|-----------|-----------|-----------|-----------|---|
| 0 | 1971 | 0.005 | 1.764052 | 0.997845 | 0.945508 | -1.962653 | 0.330046 | |
| 1 | 1971 | 0.005 | 0.400157 | 0.260081 | 0.422924 | -0.834032 | -0.000480 | |
| 2 | 1971 | 0.005 | 0.978738 | 0.925066 | -1.175680 | 1.993837 | 0.818116 | |
| 3 | 1971 | 0.005 | 2.240893 | 1.476076 | -0.204851 | -0.666660 | 0.428214 | |
| 4 | 1971 | 0.005 | 1.867558 | -1.879252 | 0.956495 | 0.976290 | -2.503947 | |

| | income |
|---|------------|
| 0 | 4.172843 |
| 1 | 112.503416 |
| 2 | 2.221875 |
| 3 | 0.814499 |
| 4 | 13.038725 |

Setting up the Problem

Unlike the fake cereal problem, we won't absorb any fixed effects in the automobile problem, so the linear part of demand X_1 has more components. We also need to specify a formula for the random coefficients X_2 , including a random coefficient on the constant, which captures correlation among all inside goods.

The primary new addition to the model relative to the fake cereal problem is that we add a supply side formula for product characteristics that contribute to marginal costs, X_3 . The *patsy*-style formulas support functions of regressors such as the `log` function used below.

We stack the three product formulations in order: X_1 , X_2 , and X_3 .

```
product_formulations = (
    pyblp.Formulation('1 + hpwt + air + mpd + space'),
    pyblp.Formulation('1 + prices + hpwt + air + mpd + space'),
    pyblp.Formulation('1 + log(hpwt) + air + log(mpg) + log(space) + trend')
)
product_formulations

(1 + hpwt + air + mpd + space,
 1 + prices + hpwt + air + mpd + space,
 1 + log(hpwt) + air + log(mpg) + log(space) + trend)
```

The original specification for the automobile problem includes the term $\log(y_i - p_j)$, in which y is income and p are prices. Instead of including this term, which gives rise to a host of numerical problems, we'll follow *Berry, Levinsohn, and Pakes (1999)* and use its first-order linear approximation, p_j/y_i .

The agent formulation for demographics, d , includes a column of $1/y_i$ values, which we'll interact with p_j . To do this, we will treat draws of y_i as demographic variables.

```
agent_formulation = pyblp.Formulation('0 + I(1 / income)')
agent_formulation
I(1 / income)
```

As in the cereal example, the *Problem* can be constructed by combining the `product_formulations`, `product_data`, `agent_formulation`, and `agent_data`. We'll also choose the functional form of marginal costs c_{jt} . A linear marginal cost specification is the default setting, so we'll need to use the `costs_type` argument of *Problem* to employ the log-linear specification used by *Berry, Levinsohn, and Pakes (1995)*.

```
problem = pyblp.Problem(product_formulations, product_data, agent_formulation, agent_data, costs_type='log')
problem
```

Dimensions:

```
=====
T      N      F      I      K1      K2      K3      D      MD      MS
-----
20    2217    26    4000     5       6       6       1     18     27
=====
```

Formulations:

```
=====
Column Indices:      0      1      2      3      4      5
-----
X1: Linear Characteristics      1      hpwt      air      mpd      space
X2: Nonlinear Characteristics      1      prices      hpwt      air      mpd      space
X3: Log Cost Characteristics      1      log(hpwt)      air      log(mpg)      log(space)      trend
d: Demographics      1/income
=====
```

The problem outputs a table of dimensions:

- T denotes the number of markets.
- N is the length of the dataset (the number of products across all markets).
- F denotes the number of firms.
- $I = \sum_t I_t$ is the total number of agents across all markets (200 draws per market times 20 markets).

- K_1 is the number of linear demand characteristics.
- K_2 is the number of nonlinear demand characteristics.
- K_3 is the number of linear supply characteristics.
- D is the number of demographic variables.
- M_D is the number of demand instruments, including exogenous regressors.
- M_S is the number of supply instruments, including exogenous regressors.

The formulations table describes all four formulas for linear characteristics, nonlinear characteristics, cost characteristics, and demographics.

Solving the Problem

The only remaining decisions are:

- Choosing Σ and Π starting values, Σ_0 and Π_0 .
- Potentially choosing bounds for Σ and Π .

The decisions we will use are:

- Use published estimates as our starting values in Σ_0 .
- Interact the inverse of income, $1/y_i$, only with prices, and use the published estimate on $\log(y_i - p_j)$ as our starting value for α in Π_0 .
- Bound Σ_0 to be positive since it is a diagonal matrix where the diagonal consists of standard deviations.
- Constrain the p_j/y_i coefficient to be negative. Specifically, we'll use a bound that's slightly smaller than zero because when the parameter is exactly zero, there are matrix inversion problems with estimating marginal costs.

When using a routine that supports bounds, *Problem* chooses some default bounds. These bounds are often very wide, so it's usually a good idea to set your own more restrictive bounds.

```
initial_sigma = np.diag([3.612, 0, 4.628, 1.818, 1.050, 2.056])
initial_pi = np.c_[[0, -43.501, 0, 0, 0, 0]]
sigma_bounds = (
    np.zeros_like(initial_sigma),
    np.diag([100, 0, 100, 100, 100, 100])
)
pi_bounds = (
    np.c_[[0, -100, 0, 0, 0, 0]],
    np.c_[[0, -0.1, 0, 0, 0, 0]]
)
```


Note that there are only 5 nonzeros on the diagonal of Σ , which means that we only need 5 columns of integration nodes to integrate over these 5 dimensions of unobserved heterogeneity. Indeed, `agent_data` contains exactly 5 columns of nodes. If we were to ignore the $\log(y_i - p_j)$ term (by not configuring `II`) and include a term on prices in Σ instead, we would have needed 6 columns of integration nodes in our `agent_data`.

A downside of the log-linear marginal costs specification is that nonpositive estimated marginal costs can create problems for the optimization routine when computing $\log c(\hat{\theta})$. We'll use the `costs_bounds` argument to bound marginal costs from below by a small number.

Finally, as in the original paper, we'll use `W_type` and `se_type` to cluster by product IDs, which were specified as `clustering_ids` in `product_data`.

```

results = problem.solve(
    initial_sigma,
    initial_pi,
    sigma_bounds=sigma_bounds,
    pi_bounds=pi_bounds,
    costs_bounds=(0.001, None),
    W_type='clustered',
    se_type='clustered'
)
results

```

Problem Results Summary:

| GMM Step | Objective Value | Projected Gradient Norm | Reduced Hessian Min Eigenvalue | Reduced Hessian Max Eigenvalue | Clipped Costs | Weighting Matrix Condition Number | Covariance Matrix Condition Number |
|----------|-----------------|-------------------------|--------------------------------|--------------------------------|---------------|-----------------------------------|------------------------------------|
| 2 | +1.5E-01 | +3.2E-09 | +8.5E-05 | +7.8E-02 | 2 | +9.7E+05 | +9.1E+07 |

Cumulative Statistics:

| Computation Time | Optimization Iterations | Objective Evaluations | Fixed Point Iterations | Contraction Evaluations |
|------------------|-------------------------|-----------------------|------------------------|-------------------------|
| 00:08:55 | 60 | 71 | 20751 | 63650 |

Nonlinear Coefficient Estimates (Robust SEs Adjusted for 999 Clusters in Parentheses):

| Sigma: | 1 | prices | hpwt | air | mpd | space | Pi: | 1/income |
|--------|------------------------|--------|------|-----|-----|-------|-----|----------|
| 1 | +6.0E+00 (+8.0E-01) | | | | | | 1 | +0.0E+00 |

(continues on next page)

(continued from previous page)

| | | | | | | | | | |
|--------|----------|----------|------------|------------|------------|------------|--|--------|------------|
| prices | +0.0E+00 | +0.0E+00 | | | | | | prices | -5.6E+01 |
| | | | | | | | | | (+4.9E+00) |
| hpwt | +0.0E+00 | +0.0E+00 | +3.2E+00 | | | | | hpwt | +0.0E+00 |
| | | | (+1.2E+00) | | | | | | |
| air | +0.0E+00 | +0.0E+00 | +0.0E+00 | +8.1E-01 | | | | air | +0.0E+00 |
| | | | | (+6.5E-01) | | | | | |
| mpd | +0.0E+00 | +0.0E+00 | +0.0E+00 | +0.0E+00 | +3.0E+00 | | | mpd | +0.0E+00 |
| | | | | | (+3.1E-01) | | | | |
| space | +0.0E+00 | +0.0E+00 | +0.0E+00 | +0.0E+00 | +0.0E+00 | +3.7E-01 | | space | +0.0E+00 |
| | | | | | | (+2.3E-01) | | | |

=====
Beta Estimates (Robust SEs Adjusted for 999 Clusters in Parentheses):
=====

| 1 | hpwt | air | mpd | space |
|------------|------------|------------|------------|------------|
| -7.1E+00 | +8.2E+00 | +2.0E+00 | -1.9E+00 | +5.4E+00 |
| (+9.1E-01) | (+1.2E+00) | (+2.6E-01) | (+3.6E-01) | (+4.7E-01) |

=====
Gamma Estimates (Robust SEs Adjusted for 999 Clusters in Parentheses):
=====

| 1 | log (hpwt) | air | log (mpg) | log (space) | trend |
|------------|------------|------------|------------|-------------|------------|
| +2.4E+00 | +6.0E-01 | +5.4E-01 | -2.0E-01 | +5.6E-02 | +1.2E-02 |
| (+9.6E-02) | (+5.9E-02) | (+4.5E-02) | (+6.1E-02) | (+9.6E-02) | (+1.6E-03) |

There are some discrepancies between our results and the original paper. Our instruments are different, we use different agent data, and the first-order linear approximation to the $\ln(y_i - p_j)$ term is different as well.

The [online version](#) of the following section may be easier to read.

4.4 Post-Estimation Tutorial

```
%matplotlib inline

import pyblp
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt

pyblp.options.digits = 2
pyblp.options.verbose = False
pyblp.__version__

'0.8.1'
```

This tutorial covers several features of `pyblp` which are available after estimation including:

1. Calculating elasticities and diversion ratios.
2. Calculating marginal costs and markups.
3. Computing the effects of mergers: prices, shares, and HHI.
4. Using a parametric bootstrap to estimate standard errors.
5. Estimating optimal instruments.

4.4.1 Problem Results

As in the *fake cereal tutorial*, we'll first solve the fake cereal problem from *Nevo (2000)*. We load the fake data and estimate the model as in the previous tutorial. We output the setup of the model to confirm we have correctly configured the *Problem*

```
product_data = pd.read_csv(pyblp.data.NEVO_PRODUCTS_LOCATION)
agent_data = pd.read_csv(pyblp.data.NEVO_AGENTS_LOCATION)
product_formulations = (
    pyblp.Formulation('0 + prices', absorb='C(product_ids)'),
```

(continues on next page)

(continued from previous page)

```

pyblp.Formulation('1 + prices + sugar + mushy')
)
agent_formulation = pyblp.Formulation('0 + income + income_squared + age + child')
problem = pyblp.Problem(product_formulations, product_data, agent_formulation, agent_data)
problem

```

Dimensions:

```

=====
T      N      F      I      K1      K2      D      MD      ED
-----
94    2256     5    1880     1       4       4      20     1
=====

```

Formulations:

```

=====
Column Indices:      0          1          2          3
-----
X1: Linear Characteristics  prices
X2: Nonlinear Characteristics  1          prices      sugar  mushy
   d: Demographics          income  income_squared  age  child
=====

```

We'll solve the problem in the same way as before. The `Problem.solve` method returns a `ProblemResults` class, which displays basic estimation results. The results that are displayed are simply formatted information extracted from various class attributes such as `ProblemResults.sigma` and `ProblemResults.sigma_se`.

```

initial_sigma = np.diag([0.3302, 2.4526, 0.0163, 0.2441])
initial_pi = [
    [ 5.4819,  0,      0.2037,  0      ],
    [15.8935, -1.2000,  0,      2.6342],
    [-0.2506,  0,      0.0511,  0      ],
    [ 1.2650,  0,      -0.8091,  0      ]
]
bfgs = pyblp.Optimization('bfgs')
results = problem.solve(
    initial_sigma,
    initial_pi,
    optimization=bfgs,
    method='ls'
)
results

```

Problem Results Summary:

| GMM Step | Objective Value | Gradient Norm | Hessian Min Eigenvalue | Hessian Max Eigenvalue | Weighting Matrix Condition Number | Covariance Matrix Condition Number |
|----------|-----------------|---------------|------------------------|------------------------|-----------------------------------|------------------------------------|
| 1 | +2.0E-03 | +3.9E-06 | +1.1E-08 | +7.3E+00 | +6.9E+07 | +8.4E+08 |

Cumulative Statistics:

| Computation Time | Optimization Iterations | Objective Evaluations | Fixed Point Iterations | Contraction Evaluations |
|------------------|-------------------------|-----------------------|------------------------|-------------------------|
| 00:02:21 | 132 | 137 | 102575 | 318945 |

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

| Sigma: | | | | | Pi: | | | | |
|--------|------------------------|------------------------|------------------------|------------------------|--------|------------------------|------------------------|------------------------|------------------------|
| | 1 | prices | sugar | mushy | | income | income_squared | age | child |
| 1 | +5.6E-01 (+1.6E-01) | | | | 1 | +2.3E+00 (+1.2E+00) | +0.0E+00 | +1.3E+00 (+6.3E-01) | +0.0E+00 |
| prices | +0.0E+00 | +3.3E+00 (+1.3E+00) | | | prices | +5.9E+02 (+2.7E+02) | -3.0E+01 (+1.4E+01) | +0.0E+00 | +1.1E+01 (+4.1E+00) |
| sugar | +0.0E+00 | +0.0E+00 | -5.8E-03 (+1.4E-02) | | sugar | -3.8E-01 (+1.2E-01) | +0.0E+00 | +5.2E-02 (+2.6E-02) | +0.0E+00 |
| mushy | +0.0E+00 | +0.0E+00 | +0.0E+00 | +9.3E-02 (+1.9E-01) | mushy | +7.5E-01 (+8.0E-01) | +0.0E+00 | -1.4E+00 (+6.7E-01) | +0.0E+00 |

Beta Estimates (Robust SEs in Parentheses):

```
=====
prices
-----
-6.3E+01
(+1.5E+01)
=====
```

Additional post-estimation outputs can be computed with *ProblemResults* methods.

4.4.2 Elasticities and Diversion Ratios

We can estimate elasticities, ε , and diversion ratios, \mathcal{D} , with `ProblemResults.compute_elasticities` and `ProblemResults.compute_diversion_ratios`.

As a reminder, elasticities in each market are

$$\varepsilon_{jk} = \frac{x_k}{s_j} \frac{\partial s_j}{\partial x_k}. \quad (4.16)$$

Diversion ratios are


$$\mathcal{D}_{jk} = -\frac{\partial s_k}{\partial x_j} / \frac{\partial s_j}{\partial x_j}. \quad (4.17)$$

Following *Conlon and Mortimer (2018)*, we report the diversion to the outside good D_{j0} on the diagonal instead of $D_{jj} = -1$.


```
elasticities = results.compute_elasticities()
diversions = results.compute_diversion_ratios()
```

Post-estimation outputs are computed for each market and stacked. We'll use `matplotlib` functions to display the matrices associated with a single market.

```
single_market = product_data['market_ids'] == 'C01Q1'
plt.colorbar(plt.matshow(elasticities[single_market]));
```



```
plt.colorbar(plt.matshow(diversions[single_market]));
```



The diagonal of the first image consists of own elasticities and the diagonal of the second image consists of diversion ratios to the outside good. As one might expect, own price elasticities are large and negative while cross-price elasticities are positive but much smaller.

Elasticities and diversion ratios can be computed with respect to variables other than prices with the name argument of *ProblemResults.compute_elasticities* and *ProblemResults.compute_diversion_ratios*. Additionally, *ProblemResults.compute_long_run_diversion_ratios* can be used to understand substitution when products are eliminated from the choice set.

The convenience methods *ProblemResults.extract_diagonals* and *ProblemResults.extract_diagonal_means* can be used to extract information about own elasticities of demand from elasticity matrices.

```
means = results.extract_diagonal_means(elasticities)
```

An alternative to summarizing full elasticity matrices is to use *ProblemResults.compute_aggregate_elasticities* to estimate aggregate elasticities of demand, E , in each market, which reflect the change in total sales under a proportional sales tax of some factor.

```
aggregates = results.compute_aggregate_elasticities(factor=0.1)
```

Since demand for an entire product category is generally less elastic than the average elasticity of individual products, mean own elasticities are generally larger in magnitude than aggregate elasticities.


```
plt.hist(
    [means.flatten(), aggregates.flatten()],
```

(continues on next page)

```

color=['red', 'blue'],
bins=50
);
plt.legend(['Mean Own Elasticities', 'Aggregate Elasticities']);

```




4.4.3 Marginal Costs and Markups

To compute marginal costs, c , the `product_data` passed to `Problem` must have had a `firm_ids` field. Since we included firm IDs when configuring the problem, we can use `ProblemResults.compute_costs`.

```


costs = results.compute_costs()
plt.hist(costs, bins=50);
plt.legend(["Marginal Costs"]);

```

Other methods that compute supply-side outputs often compute marginal costs themselves. For example, `ProblemResults.compute_markup`s will compute marginal costs when estimating markups, μ , but computation can be sped up if we just use our pre-computed values.

```
markups = results.compute_markup(costs=costs)
plt.hist(markups, bins=50);
plt.legend(["Markups"]);
```



4.4.4 Mergers

Before computing post-merger outputs, we'll supplement our pre-merger markups with some other outputs. We'll compute Herfindahl-Hirschman Indices, HHI, with `ProblemResults.compute_hhi`; population-normalized gross expected profits, π , with `ProblemResults.compute_profits`; and population-normalized consumer surpluses, CS, with `ProblemResults.compute_consumer_surpluses`.

```
hhi = results.compute_hhi()
profits = results.compute_profits(costs=costs)
cs = results.compute_consumer_surpluses()
```

To compute post-merger outputs, we'll create a new set of firm IDs that represent a merger of firms 2 and 1.

```
product_data['merger_ids'] = product_data['firm_ids'].replace(2, 1)
```

We can use `ProblemResults.compute_approximate_prices` or `ProblemResults.compute_prices` to estimate post-merger prices. The first method, which is discussed, for example, in *Nevo (1997)*, assumes that shares and their price derivatives are unaffected by the merger. The second method does not make these assumptions and iterates over the ζ -markup equation from *Morrow and Skerlos (2011)* to solve the full system of J_t equations and J_t unknowns in each market t . We'll use the latter, since it is fast enough for this example problem.

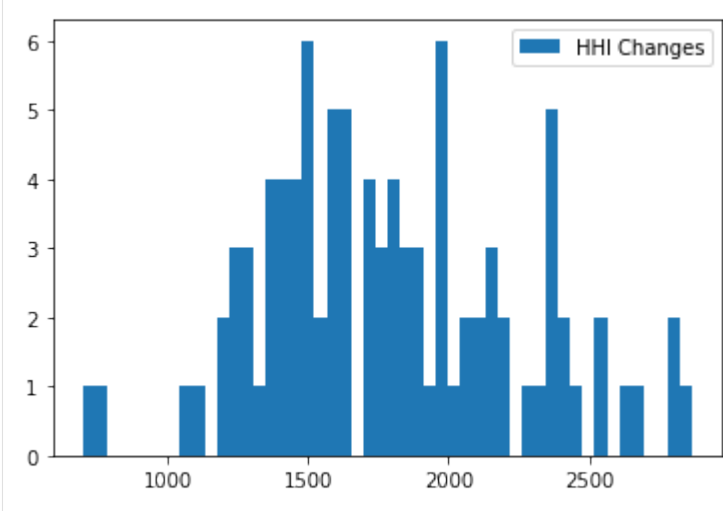
```
changed_prices = results.compute_prices(
    firm_ids=product_data['merger_ids'],
    costs=costs
)
```

We'll compute post-merger shares with `ProblemResults.compute_shares`.

```
changed_shares = results.compute_shares(changed_prices)
```


Post-merger prices and shares are used to compute other post-merger outputs. For example, HHI increases.

```
changed_hhi = results.compute_hhi(
    firm_ids=product_data['merger_ids'],
    shares=changed_shares
)
plt.hist(changed_hhi - hhi, bins=50);
plt.legend(["HHI Changes"]);
```

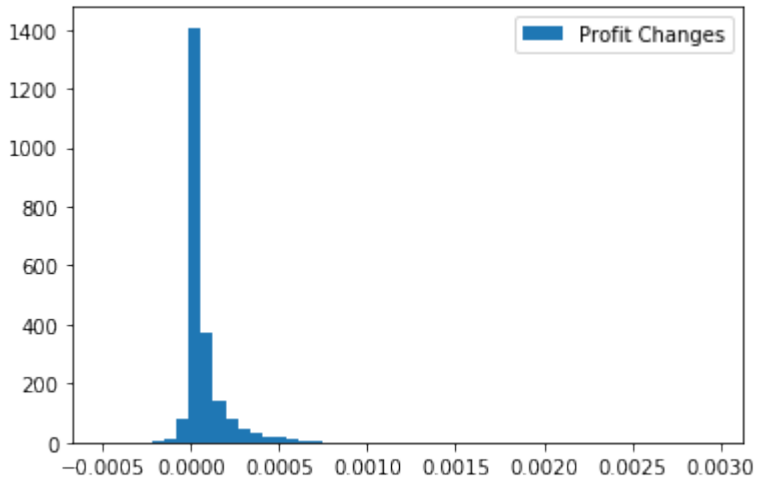


Markups, \mathcal{M} , and profits, π , generally increase as well.

```
changed_markups = results.compute_markups(changed_prices, costs)
plt.hist(changed_markups - markups, bins=50);
plt.legend(["Markup Changes"]);
```




```
changed_profits = results.compute_profits(changed_prices, changed_shares, costs)
plt.hist(changed_profits - profits, bins=50);
plt.legend(["Profit Changes"]);
```



On the other hand, consumer surpluses, CS, generally decrease.

```
changed_cs = results.compute_consumer_surpluses(changed_prices)
plt.hist(changed_cs - cs, bins=50);
plt.legend(["Consumer Surplus Changes"]);
```



4.4.5 Bootstrapping Results

Post-estimation outputs can be informative, but they don't mean much without a sense sample-to-sample variability. One way to estimate confidence intervals for post-estimation outputs is with a standard bootstrap procedure:

1. Construct a large number of bootstrap samples by sampling with replacement from the original product data.
2. Initialize and solve a *Problem* for each bootstrap sample.
3. Compute the desired post-estimation output for each bootstrapped *ProblemResults* and from the resulting empirical distribution, construct bootstrap confidence intervals.

Although appealing because of its simplicity, the computational resources required for this procedure are often prohibitively expensive. Furthermore, human oversight of the optimization routine is often required to determine whether the routine ran into any problems and if it successfully converged. Human oversight of estimation for each bootstrapped problem is usually not feasible.

A more reasonable alternative is a parametric bootstrap procedure:

1. Construct a large number of draws from the estimated joint distribution of parameters.
2. Compute the implied mean utility, δ , and shares, s , for each draw. If a supply side was estimated, also computed the implied marginal costs, c , and prices, p .
3. Compute the desired post-estimation output under each of these parametric bootstrap samples. Again, from the resulting empirical distribution, construct bootstrap confidence intervals.

Compared to the standard bootstrap procedure, the parametric bootstrap requires far fewer computational resources, and is simple enough to not require human oversight of each bootstrap iteration. The primary complication to this procedure is that when supply is estimated, equilibrium prices and shares need to be computed for each parametric bootstrap sample by iterating over the ζ -markup equation from *Morrow and Skerlos (2011)*. Although nontrivial, this fixed point iteration problem is much less demanding than the full optimization routine required to solve the BLP problem from the start.

An empirical distribution of results computed according to this parametric bootstrap procedure can be created with the `ProblemResults.bootstrap` method, which returns a `BootstrappedResults` class that can be used just like `ProblemResults` to compute various post-estimation outputs. The difference is that `BootstrappedResults` methods return arrays with an extra first dimension, along which bootstrapped results are stacked.

We'll construct 90% parametric bootstrap confidence intervals for estimated mean own elasticities in each market of the fake cereal problem. Usually, bootstrapped confidence intervals should be based on thousands of draws, but we'll only use a few for the sake of speed in this example.

```
bootstrapped_results = results.bootstrap(draws=100, seed=0)
bootstrapped_results
```

```
Bootstrapped Results Summary:
=====
Computation  Bootstrap
   Time      Draws
-----
00:00:20    100
=====
```

```
bounds = np.percentile(
    bootstrapped_results.extract_diagonal_means(
        bootstrapped_results.compute_elasticities()
    ),
    q=[10, 90],
    axis=0
)
table = pd.DataFrame(index=problem.unique_market_ids, data={
    'Lower Bound': bounds[0].flatten(),
    'Mean Own Elasticity': aggregates.flatten(),
    'Upper Bound': bounds[1].flatten()
})
table.round(2).head()
```

| | Lower Bound | Mean Own Elasticity | Upper Bound |
|-------|-------------|---------------------|-------------|
| C01Q1 | -14.06 | -0.81 | 10.66 |
| C01Q2 | -8.97 | -0.69 | 7.54 |
| C03Q1 | -11.64 | -0.55 | 7.98 |
| C03Q2 | -14.46 | -0.60 | 6.47 |
| C04Q1 | -12.00 | -0.68 | 5.61 |

4.4.6 Optimal Instruments

Given a consistent estimate of θ , we may want to compute the optimal instruments of *Chamberlain (1987)* and use them to re-solve the problem. Optimal instruments have been shown, for example, by *Reynaert and Verboven (2014)*, to reduce bias, improve efficiency, and enhance stability of BLP estimates.

The `ProblemResults.compute_optimal_instruments` method computes the expected Jacobians that comprise the optimal instruments by integrating over the density of ξ (and ω if a supply side was estimated). By default, the method approximates this integral by averaging over the Jacobian realizations computed under draws from the asymptotic normal distribution of the error terms. Since this process is computationally expensive and often doesn't make much of a difference, we'll use `method='approximate'` in this example to simply evaluate the Jacobians at the expected value of ξ , zero.

```
instrument_results = results.compute_optimal_instruments(method='approximate')
instrument_results
```

Optimal Instrument Results Summary:

```
=====
Computation Error Term
  Time      Draws
-----
00:00:01      1
=====
```

We can use the `OptimalInstrumentResults.to_problem` method to re-create the fake cereal problem with the estimated optimal excluded instruments.

```
updated_problem = instrument_results.to_problem()
updated_problem
```

Dimensions:

```
=====
  T    N    F    I    K1    K2    D    MD    ED
-----
94   2256    5  1880    1     4     4    14    1
=====
```

Formulations:

```
=====
      Column Indices:      0      1      2      3
-----
X1: Linear Characteristics  prices
X2: Nonlinear Characteristics  1      prices  sugar  mushy
      d: Demographics      income income_squared  age  child
=====
```

We can solve this updated problem just like the original one. We'll start at our consistent estimate of θ .

```

updated_results = updated_problem.solve(
    results.sigma,
    results.pi,
    optimization=pyblp.Optimization('bfgs'),
    method='1s'
)
updated_results

```

Problem Results Summary:

```

=====
GMM   Objective  Gradient      Hessian      Hessian      Weighting Matrix  Covariance Matrix
Step  Value       Norm         Min Eigenvalue Max Eigenvalue Condition Number  Condition Number
-----
1     +2.5E-03   +7.1E-06    +1.2E-16     +6.8E+00     +7.8E+07         +3.3E+20
=====

```

Cumulative Statistics:

```

=====
Computation  Optimization  Objective  Fixed Point  Contraction
Time         Iterations    Evaluations Iterations   Evaluations
-----
00:02:18    114          119        111996      346574
=====

```

Nonlinear Coefficient Estimates (Robust SEs in Parentheses):

```

=====
Sigma:      1          prices      sugar      mushy      |  Pi:      income      income_squared      age      child
-----
1          +2.2E-01      |  1          +6.2E+00      +0.0E+00      +1.0E-01      +0.0E+00
          (+8.1E-02)      |
prices     +0.0E+00      +3.1E+00      |  prices     +4.5E+01      -2.8E+00      +0.0E+00      +3.3E+00
          (+7.0E-01)      |          (+9.3E+01)      (+4.9E+00)      (+2.4E+00)
sugar      +0.0E+00      +0.0E+00      -5.8E-03      |  sugar      -2.8E-01      +0.0E+00      +3.7E-02      +0.0E+00
          (NAN)          |          (+3.6E-02)      (+1.6E-02)
mushy      +0.0E+00      +0.0E+00      +0.0E+00      +9.5E-01      |  mushy      +5.1E-01      +0.0E+00      -1.8E-01      +0.0E+00
          (+3.0E-01)      |          (+2.7E-01)      (+2.3E-01)
=====

```

(continues on next page)

(continued from previous page)

```
Beta Estimates (Robust SEs in Parentheses):
```

```
=====
prices
-----
-2.9E+01
(+4.5E+00)
=====
```

The [online version](#) of the following section may be easier to read.

4.5 Problem Simulation Tutorial

```
import pyblp
import numpy as np
import pandas as pd

pyblp.options.digits = 2
pyblp.options.verbose = False
pyblp.__version__

'0.8.1'
```

Before configuring and solving a problem with real data, it may be a good idea to perform Monte Carlo analysis on simulated data to verify that it is possible to accurately estimate model parameters. For example, before configuring and solving the example problems in the prior tutorials, it may have been a good idea to simulate data according to the assumed models of supply and demand. During such Monte Carlo analysis, the data would only be used to determine sample sizes and perhaps to choose reasonable true parameters.

Simulations are configured with the *Simulation* class, which requires many of the same inputs as *Problem*. The two main differences are:

1. Variables in formulations that cannot be loaded from `product_data` or `agent_data` will be drawn from independent uniform distributions.
2. True parameters and the distribution of unobserved product characteristics are specified.

First, we'll use *build_id_data* to build market and firm IDs for a model in which there are $T = 50$ markets, and in each market t , a total of $J_t = 20$ products produced by $F = 10$ firms.

```
id_data = pyblp.build_id_data(T=50, J=20, F=10)
```

Next, we'll create an *Integration* configuration to build agent data according to a Gauss-Hermite product rule that exactly integrates polynomials of degree $2 \times 9 - 1 = 17$ or less.

```
integration = pyblp.Integration('product', 9)
integration
```

```
Configured to construct nodes and weights according to the level-9 Gauss-Hermite product rule.
```

We'll then pass these data to *Simulation*. We'll use *Formulation* configurations to create an X_1 that consists of a constant, prices, and an exogenous characteristic; an X_2 that consists only of the same exogenous characteristic; and an X_3 that consists of the common exogenous characteristic and a cost-shifter.

```

simulation = pyblp.Simulation(
    product_formulations=(
        pyblp.Formulation('1 + prices + x'),
        pyblp.Formulation('0 + x'),
        pyblp.Formulation('0 + x + z')
    ),
    beta=[1, -2, 2],
    sigma=1,
    gamma=[1, 4],
    product_data=id_data,
    integration=integration,
    seed=0
)
simulation

```

Dimensions:

```

=====
T      N      F      I      K1      K2      K3
-----
50    1000    10     450     3        1        2
=====

```

Formulations:

```

=====
          Column Indices:          0      1      2
-----
X1: Linear Characteristics      1  prices  x
X2: Nonlinear Characteristics   x
X3: Linear Cost Characteristics x      z
=====

```

Nonlinear Coefficient True Values:

```

=====
Sigma:      x
-----
x          +1.0E+00
=====

```

Beta True Values:

```

=====
1          prices      x

```

(continues on next page)

(continued from previous page)

```

-----
+1.0E+00  -2.0E+00  +2.0E+00
=====

Gamma True Values:
=====
      x          z
-----
+1.0E+00  +4.0E+00
=====

```

When *Simulation* is initialized, it constructs *Simulation.agent_data* and simulates *Simulation.product_data*.

The *Simulation* can be further configured with other arguments that determine how unobserved product characteristics are simulated and how marginal costs are specified.

At this stage, simulated variables are not consistent with true parameters, so we still need to solve the simulation with *Simulation.replace_endogenous*. This method replaced simulated prices and marketshares with values that are consistent with the true parameters. Just like *ProblemResults.compute_prices*, to do so it iterates over the ζ -markup equation from *Morrow and Skerlos (2011)*.

```

simulation_results = simulation.replace_endogenous()
simulation_results

```

```

Simulation Results Summary:
=====
Computation  Fixed Point  Contraction
   Time      Iterations  Evaluations
-----
 00:00:00      721          721
=====

```

Now, we can try to recover the true parameters by creating and solving a *Problem*.

The convenience method *SimulationResults.to_problem* constructs some basic “sums of characteristics” BLP instruments that are functions of all exogenous numerical variables in the problem. In this example, excluded demand-side instruments are the cost-shifter *z* and traditional BLP instruments constructed from *x*. Excluded supply-side instruments are traditional BLP instruments constructed from *x* and *z*.

```

problem = simulation_results.to_problem()
problem

```

```

Dimensions:
=====

```

(continues on next page)

(continued from previous page)

| T | N | F | I | K1 | K2 | K3 | MD | MS |
|----|------|----|-----|----|----|----|----|----|
| 50 | 1000 | 10 | 450 | 3 | 1 | 2 | 5 | 6 |

Formulations:

| Column Indices: | 0 | 1 | 2 |
|---------------------------------|---|--------|---|
| X1: Linear Characteristics | 1 | prices | x |
| X2: Nonlinear Characteristics | x | | |
| X3: Linear Cost Characteristics | x | z | |

We'll choose starting values that are half the true parameters so that the optimization routine has to do some work. Note that since we're jointly estimating the supply side, we need to provide an initial value for the linear coefficient on prices because this parameter cannot be concentrated out of the problem (unlike linear coefficients on exogenous characteristics).

```
results = problem.solve(
    sigma=0.5 * simulation.sigma,
    pi=0.5 * simulation.pi,
    beta=[None, 0.5 * simulation.beta[1], None]
)
results
```

Problem Results Summary:

| GMM Step | Objective Value | Projected Gradient Norm | Reduced Hessian Min Eigenvalue | Reduced Hessian Max Eigenvalue | Weighting Matrix Condition Number | Covariance Matrix Condition Number |
|----------|-----------------|-------------------------|--------------------------------|--------------------------------|-----------------------------------|------------------------------------|
| 2 | +6.4E-03 | +1.9E-10 | +7.2E-03 | +3.8E+00 | +3.7E+04 | +1.5E+04 |

Cumulative Statistics:

| Computation Time | Optimization Iterations | Objective Evaluations | Fixed Point Iterations | Contraction Evaluations |
|------------------|-------------------------|-----------------------|------------------------|-------------------------|
| 00:00:23 | 25 | 30 | 7762 | 24616 |

(continues on next page)

(continued from previous page)

```
Nonlinear Coefficient Estimates (Robust SEs in Parentheses):
```

```
=====
Sigma:      x
-----
x          +7.8E-01
          (+5.2E-01)
=====
```

```
Beta Estimates (Robust SEs in Parentheses):
```

```
=====
      1          prices          x
-----
+9.6E-01   -2.0E+00   +2.1E+00
(+9.3E-02) (+2.4E-02) (+1.4E-01)
=====
```

```
Gamma Estimates (Robust SEs in Parentheses):
```

```
=====
      x          z
-----
+9.8E-01   +4.0E+00
(+8.7E-02) (+8.5E-02)
=====
```

The parameters seem to have been estimated reasonably well.

```
np.c_[simulation.beta, results.beta]
```

```
array([[ 1.          ,  0.96223514],
       [-2.          , -2.00792431],
       [ 2.          ,  2.10032015]])
```

```
np.c_[simulation.gamma, results.gamma]
```

```
array([[1.          ,  0.97820624],
       [4.          ,  4.03121577]])
```

```
np.c_[simulation.sigma, results.sigma]
```

```
array([[1.          , 0.78358853]])
```

In addition to checking that the configuration for a model based on actual data makes sense, the *Simulation* class can also be a helpful tool for better understanding under what general conditions BLP models can be accurately estimated. Simulations are also used extensively in pyblp's test suite.

API DOCUMENTATION

The majority of the package consists of classes, which compartmentalize different aspects of the BLP model. There are some convenience functions as well.

5.1 Configuration Classes

Various components of the package require configurations for how to approximate integrals, solve fixed point problems, and solve optimization problems. Such configurations are specified with the following classes.

| | |
|---|---|
| <i>Formulation</i> (formula[, absorb, absorb_method]) | Configuration for designing matrices and absorbing fixed effects. |
| <i>Integration</i> (specification, size[, seed]) | Configuration for building integration nodes and weights. |
| <i>Iteration</i> (method[, method_options, ...]) | Configuration for solving fixed point problems. |
| <i>Optimization</i> (method[, method_options, ...]) | Configuration for solving optimization problems. |

5.1.1 pyblp.Formulation

class `pyblp.Formulation` (*formula*, *absorb=None*, *absorb_method=None*)

Configuration for designing matrices and absorbing fixed effects.

Internally, the `patsy` package is used to convert data and R-style formulas into matrices. All of the standard `binary operators` can be used to design complex matrices of factor interactions:

- `+` - Set union of terms.
- `-` - Set difference of terms.
- `*` - Short-hand. The formula `a * b` is the same as `a + b + a:b`.
- `/` - Short-hand. The formula `a / b` is the same as `a + a:b`.
- `:` - Interactions between two sets of terms.
- `**` - Interactions up to an integer degree.

However, since factors need to be differentiated (for example, when computing elasticities), only the most essential functions are supported:

- `C` - Mark a variable as categorical. See `patsy.builtins.C()`. Arguments are not supported.
- `I` - Encapsulate mathematical operations. See `patsy.builtins.I()`.
- `log` - Natural logarithm function.

- `exp` - Natural exponential function.

Data associated with variables should generally already be transformed. However, when encapsulated by `I()`, these operators function like normal mathematical operators on numeric variables: `+` adds, `-` subtracts, `*` multiplies, `/` divides, and `**` exponentiates.

Internally, mathematical operations are parsed and evaluated by the `SymPy` package, which is also used to symbolically differentiate terms when derivatives are needed.

Parameters

- **formula** (*str*) – R-style formula used to design a matrix. Variable names will be validated when this formulation and data are passed to a function that uses them. By default, an intercept is included, which can be removed with `0` or `-1`. If `absorb` is specified, intercepts are ignored.
- **absorb** (*str, optional*) – R-style formula used to design a matrix of categorical variables representing fixed effects, which will be absorbed into the matrix designed by `formula`. Fixed effect absorption is only supported for some matrices. Unlike `formula`, intercepts are ignored. Only categorical variables are supported.
- **absorb_method** (*str or Iteration, optional*) – The method with which fixed effects will be absorbed. One of the following:
 - `'simple'` (default for one fixed effect) - Use simple de-meaning. This method is very unlikely to fully absorb more than one fixed effect.
 - `'memory'` (default for two fixed effects) - Use the *Somaini and Wolak (2016)* algorithm, which only works for two-way fixed effects, and which requires inversion of a dense matrix with dimensions equal to the smaller number of fixed effect groups.
 - `'speed'` - Use the same *Somaini and Wolak (2016)* algorithm but pre-compute the *A* matrix, which is a dense matrix with dimensions equal to the larger number of fixed effect groups. Again, this method only works for two-way fixed effects.
 - `Iteration` (default for more than two fixed effects) - Use the method of alternating projections described, for example, in *Guimarães and Portugal (2010)*. By default, `Iteration('simple', {'atol': 1e-12})` is used to iteratively de-mean the matrix within each fixed effect level until convergence. This method is equivalent to `'simple'` for one fixed effect, and it will also work for two fixed effects, although either variant of the *Somaini and Wolak (2016)* algorithm is often more performant.

Examples

The [online version](#) of the following section may be easier to read.

Formulation Example

```
import pyblp

pyblp.__version__

'0.8.1'
```

In this example, we'll design a matrix without an intercept, but with both prices and another numeric size variable.

```
formulation = pyblp.Formulation('0 + prices + size')
formulation

prices + size
```

Next, we'll design a second matrix with an intercept, with first- and second-degree size terms, with categorical product IDs and years, and with the interaction of the last two. The first formulation will include the fixed effects as indicator variables, and the second will absorb them.

```
formulation1 = pyblp.Formulation('size + I(size ** 2) + C(product) * C(year)')
formulation1

1 + size + I(size ** 2) + C(product) + C(year) + C(product):C(year)
```

```
formulation2 = pyblp.Formulation('size + I(size ** 2)', absorb='C(product) * C(year)')
formulation2

size + I(size ** 2) + Absorb[C(product)] + Absorb[C(year)] + Absorb[C(product):C(year)]
```

Finally, we'll design a third matrix with an intercept and with a yearly trend interacted with the natural logarithm of income and categorical education. Absorption of continuous variables is not supported, so we need to use dummy variables.

```
formulation = pyblp.Formulation('year:(log(income) + C(education))')
formulation

1 + year:log(income) + year:C(education)
```

5.1.2 pyblp.Integration

class `pyblp.Integration` (*specification, size, seed=None*)
Configuration for building integration nodes and weights.

Parameters

- **specification** (*str*) – How to build nodes and weights. One of the following:
 - 'monte_carlo' - Draw from a pseudo-random standard multivariate normal distribution. Integration weights are $1 / \text{size}$.
 - 'product' - Generate nodes and weights according to the level-size Gauss-Hermite product rule.
 - 'nested_product' - Generate nodes and weights according to the level-size nested Gauss-Hermite product rule. Weights can be negative.
 - 'grid' - Generate a sparse grid of nodes and weights according to the level-size Gauss-Hermite quadrature rule. Weights can be negative.
 - 'nested_grid' - Generate a sparse grid of nodes and weights according to the level size nested Gauss-Hermite quadrature rule. Weights can be negative.

Best practice for low dimensions is probably to use 'product' to a relatively high degree of polynomial accuracy. In higher dimensions, 'grid' appears to scale the best. For more information, see *Judd and Skrainka (2011)* and *Conlon and Gortmaker (2019)*.

Sparse grids are constructed in analogously to the Matlab function `nwspgr` created by Florian Heiss and Viktor Winschel. For more information, see *Heiss and Winschel (2008)*.

- **size** (*int*) – The number of draws if `specification` is 'monte_carlo', and the level of the quadrature rule otherwise.
- **seed** (*int, optional*) – Passed to `numpy.random.mtrand.RandomState` when `specification` is 'monte_carlo' to seed the random number generator before building nodes. By default, a seed is not passed to the random number generator.

Examples

The [online version](#) of the following section may be easier to read.

Integration Example

```
import pyblp

pyblp.__version__

'0.8.1'
```

In this example, we'll build a Monte Carlo configuration with 1,000 draws for each market and a fixed seed.

```
integration = pyblp.Integration('monte_carlo', size=1000, seed=0)
integration

Configured to construct nodes and weights with Monte Carlo simulation.
```

Depending on the dimension of the integration problem, a level six sparse grid configuration may have a similar number of nodes. However, even if there are fewer nodes, it is likely to perform better in the BLP problem. Sparse grid construction is deterministic, so a seed is not needed to fix the grid every time we use this configuration.

```
integration = pyblp.Integration('grid', size=7)
integration

Configured to construct nodes and weights in a sparse grid according to the level-7 Gauss-Hermite rule.
```

5.1.3 pyblp.Iteration

class `pyblp.Iteration` (*method*, *method_options=None*, *compute_jacobian=False*)

Configuration for solving fixed point problems.

Parameters

- **method** (*str or callable*) – The fixed point iteration routine that will be used. The following routines do not use analytic Jacobians:
 - 'simple' - Non-accelerated iteration.
 - 'squarem' - SQUAREM acceleration method of *Varadhan and Roland (2008)* and considered in the context of the BLP problem in *Reynaerts, Varadhan, and Nash (2012)*. This implementation uses a first-order squared non-monotone extrapolation scheme. If there are any errors during the acceleration step, it uses the last values for the next iteration of the algorithm.
 - 'broyden1' - Use the `scipy.optimize.root()` Broyden's first Jacobian approximation method, known as Broyden's good method.
 - 'broyden2' - Use the `scipy.optimize.root()` Broyden's second Jacobian approximation method, known as Broyden's bad method.
 - 'anderson' - Use the `scipy.optimize.root()` Anderson method.
 - 'krylov' - Use the `scipy.optimize.root()` Krylov approximation for inverse Jacobian method.
 - 'diagbroyden' - Use the `scipy.optimize.root()` diagonal Broyden Jacobian approximation method.
 - 'df-sane' - Use the `scipy.optimize.root()` derivative-free spectral method.

The following routines can use analytic Jacobians:

- 'hybr' - Use the `scipy.optimize.root()` modification of the Powell hybrid method implemented in MINIPACK.
- 'lm' - Uses the `scipy.optimize.root()` modification of the Levenberg-Marquardt algorithm implemented in MINIPACK.

The following trivial routine can be used to simply return the initial values:

- 'return' - Assume that the initial values are the optimal ones.

Also accepted is a custom callable method with the following form:

```
method(initial, contraction, callback, **options) -> (final,
↳ converged)
```

where `initial` is an array of initial values, `contraction` is a callable contraction mapping of the form specified below, `callback` is a function that should be called without any arguments after each major iteration (it is used to record the number of major iterations), `options` are specified below, `final` is an array of final values, and `converged` is a flag for whether the routine converged.

The `contraction` function has the following form:

```
contraction(x0) -> (x1, weights, jacobian)
```

where `weights` are either `None` or a vector of weights that should multiply `x1 - x` before computing the norm of the differences, and `jacobian` is `None` if `compute_jacobian` is `False`.

Regardless of the chosen routine, if there are any computational issues that create infinities or null values, `final` will be the second to last iteration's values.

- **method_options** (*dict, optional*) – Options for the fixed point iteration routine.

For routines other than `'simple'`, `'squarem'`, and `'return'`, these options will be passed to `options` in `scipy.optimize.root()`. Refer to the SciPy documentation for information about which options are available. By default, the `tol_norm` option is configured to use the infinity norm for SciPy methods other than `'hybr'` and `'lm'`, for which a norm cannot be specified.

The `'simple'` and `'squarem'` methods support the following options:

- **max_evaluations** : (*int*) - Maximum number of contraction mapping evaluations. The default value is 5000.
- **atol** : (*float*) - Absolute tolerance for convergence of the configured norm. The default value is $1e-14$. To use only a relative tolerance, set this to zero.
- **rtol** (*float*) - Relative tolerance for convergence of the configured norm. The default value is zero; that is, only absolute tolerance is used by default.
- **norm** : (*callable*) - The norm to be used. By default, the ℓ^∞ -norm is used. If specified, this should be a function that accepts an array of differences and that returns a scalar norm.

The `'squarem'` routine accepts additional options that mirror those in the [SQUAREM](#) package, written in R by Ravi Varadhan, which identifies the step length with $-\alpha$ from [Varadhan and Roland \(2008\)](#):

- **scheme** : (*int*) - The default value is 3, which corresponds to S3 in [Varadhan and Roland \(2008\)](#). Other acceptable schemes are 1 and 2, which correspond to S1 and S2.
- **step_min** : (*float*) - The initial value for the minimum step length. The default value is 1.0.
- **step_max** : (*float*) - The initial value for the maximum step length. The default value is 1.0.
- **step_factor** : (*float*) - When the step length exceeds `step_max`, it is set equal to `step_max`, but `step_max` is scaled by this factor. Similarly, if `step_min` is negative and the step length is below `step_min`, it is set equal to `step_min` and `step_min` is scaled by this factor. The default value is 4.0.
- **compute_jacobian** (*bool, optional*) – Whether to compute an analytic Jacobian during iteration, which must be `False` if `method` does not use analytic Jacobians. By default, analytic Jacobians are not computed, and if a `method` is selected that supports analytic Jacobians, they will by default be numerically approximated.

Examples

The [online version](#) of the following section may be easier to read.

Iteration Example

```
import pyblp
import numpy as np

pyblp.__version__

'0.8.1'
```

In this example, we'll build a SQUAREM configuration with a ℓ^2 -norm and use scheme S1 from *Varadhan and Roland (2008)*.

```
iteration = pyblp.Iteration('squarem', {'norm': np.linalg.norm, 'scheme': 1})
iteration
```

```
Configured to iterate using the SQUAREM acceleration method without analytic Jacobians with options {atol: +1.000000E-
↪14, rtol: 0, max_evaluations: 5000, norm: numpy.linalg.norm, scheme: 1, step_min: +1.000000E+00, step_max: +1.
↪000000E+00, step_factor: +4.000000E+00}.
```

Next, instead of using a built-in routine, we'll create a custom method that implements a version of simple iteration, which, for the sake of having a nontrivial example, arbitrarily identifies a major iteration with three objective evaluations.

```
def custom_method(initial, contraction, callback, max_evaluations, tol, norm):
    x = initial
    evaluations = 0
    while evaluations < max_evaluations:
        x0, (x, weights, _) = x, contraction(x)
        evaluations += 1
        if evaluations % 3 == 0:
            callback()
        if weights is None:
            difference = norm(x - x0)
        else:
            difference = norm(weights * (x - x0))
        if difference < tol:
```

(continues on next page)

(continued from previous page)

```
        break
    return x, evaluations < max_evaluations
```

We can then use this custom method to build a custom iteration configuration.

```
iteration = pyblp.Iteration(custom_method)
iteration
```

```
Configured to iterate using a custom method without analytic Jacobians with options {}.
```

5.1.4 pyblp.Optimization

`class pyblp.Optimization` (*method*, *method_options=None*, *compute_gradient=True*, *universal_display=True*)

Configuration for solving optimization problems.

Parameters

- **method** (*str or callable*) – The optimization routine that will be used. The following routines support parameter bounds and use analytic gradients:
 - 'knitro' - Uses an installed version of [Artleys Knitro](#). Python 3 is supported by Knitro version 10.3 and newer. A number of environment variables most likely need to be configured properly, such as `KNITRODIR`, `ARTELYS_LICENSE`, `LD_LIBRARY_PATH` (on Linux), and `DYLD_LIBRARY_PATH` (on Mac OS X). For more information, refer to the [Knitro installation guide](#).
 - 'slsqp' - Uses the `scipy.optimize.minimize()` SLSQP routine.
 - 'trust-constr' - Uses the `scipy.optimize.minimize()` trust-region routine.
 - 'l-bfgs-b' - Uses the `scipy.optimize.minimize()` L-BFGS-B routine.
 - 'tnc' - Uses the `scipy.optimize.minimize()` TNC routine.

The following routines also use analytic gradients but will ignore parameter bounds (not bounding the problem may create issues if the optimizer tries out large parameter values that create overflow errors):

- 'cg' - Uses the `scipy.optimize.minimize()` CG routine.
- 'bfgs' - Uses the `scipy.optimize.minimize()` BFGS routine.
- 'newton-cg' - Uses the `scipy.optimize.minimize()` Newton-CG routine.

The following routines do not use analytic gradients and will also ignore parameter bounds (without analytic gradients, optimization will likely be much slower):

- 'nelder-mead' - Uses the `scipy.optimize.minimize()` Nelder-Mead routine.
- 'powell' - Uses the `scipy.optimize.minimize()` Powell routine.

The following trivial routine can be used to evaluate an objective at specific parameter values:

- 'return' - Assume that the initial parameter values are the optimal ones.

Also accepted is a custom callable method with the following form:

```
method(initial, bounds, objective_function, iteration_callback,
↳ **options) -> (final, converged)
```

where `initial` is an array of initial parameter values, `bounds` is a list of (`min`, `max`) pairs for each element in `initial`, `objective_function` is a callable objective function of the form specified below, `iteration_callback` is a function that should be called without any arguments after each major iteration (it is used to record the number of major iterations), `options` are specified below, `final` is an array of optimized parameter values, and `converged` is a flag for whether the routine converged.

The `objective_function` has the following form:

```
objective_function(theta) -> (objective, gradient)
```

where `gradient` is `None` if `compute_gradient` is `False`.

- **method_options** (*dict, optional*) – Options for the optimization routine.

For any non-custom method other than 'knitro' and 'return', these options will be passed to options in `scipy.optimize.minimize()`. Refer to the SciPy documentation for information about which options are available for each optimization routine.

If method is 'knitro', these options should be [Knitro user options](#). The non-standard `knitro_dir` option can also be specified. The following options have non-standard default values:

- **knitro_dir** : (*str*) - By default, the KNITRODIR environment variable is used. Otherwise, this option should point to the installation directory of Knitro, which contains direct subdirectories such as 'examples' and 'lib'. For example, on Windows this option could be '/Program Files/Artleys3/Knitro 10.3.0'.
- **algorithm** : (*int*) - The optimization algorithm to be used. The default value is 1, which corresponds to the Interior/Direct algorithm.
- **gradopt** : (*int*) - How the objective's gradient is computed. The default value is 1 if `compute_gradient` is `True` and is 2 otherwise, which corresponds to estimating the gradient with finite differences.
- **hessopt** : (*int*) - How the objective's Hessian is computed. The default value is 2, which corresponds to computing a quasi-Newton BFGS Hessian.
- **honorbnds** : (*int*) - Whether to enforce satisfaction of simple variable bounds. The default value is 1, which corresponds to enforcing that the initial point and all subsequent solution estimates satisfy the bounds.
- **compute_gradient** (*bool, optional*) – Whether to compute an analytic objective gradient during optimization, which must be `False` if method does not use analytic gradients, and must be `True` if method is 'newton-cg', which requires an analytic gradient. By default, analytic gradients are computed. Not using an analytic gradient will likely slow down estimation a good deal. If `False`, an analytic gradient may still be computed once at the end of optimization to compute optimization results.
- **universal_display** (*bool, optional*) – Whether to format optimization progress such that the display looks the same for all routines. By default, the universal display is used and some `method_options` are used to prevent default displays from showing up.

Examples

The [online version](#) of the following section may be easier to read.

Optimization Example

```
import pyblp
import numpy as np

pyblp.__version__

'0.8.1'
```

In this example, we'll build a L-BFGS-B configuration with a non-default tolerance.

```
optimization = pyblp.Optimization('l-bfgs-b', {'gtol': 1e-3})
optimization
```

```
Configured to optimize using the L-BFGS-B algorithm implemented in SciPy with analytic gradients and options {gtol: +1.
↪000000E-03}.
```

Next, instead of using a non-custom routine, we'll create a custom method that implements a grid search over parameter values between specified bounds.

```
from itertools import product
def custom_method(initial, bounds, objective_function, iteration_callback):
    best_values = initial
    best_objective = np.inf
    for values in product(*(np.linspace(l, u, 10) for l, u in bounds)):
        objective, _ = objective_function(values)
        if objective < best_objective:
            best_values = values
            best_objective = objective
        iteration_callback()
    return best_values, True
```

We can then use this custom method to build an optimization configuration.

```
optimization = pyblp.Optimization(custom_method, compute_gradient=False)
optimization
```

```
Configured to optimize using a custom method without analytic gradients and options {}.
```

5.2 Data Construction Functions

There are also a number of convenience functions that can be used to construct common components of product and agent data.

| | |
|---|--|
| <code>build_matrix(formulation, data)</code> | Construct a matrix according to a formulation. |
| <code>build_blp_instruments(formulation, product_data)</code> | Construct “sums of characteristics” excluded BLP instruments. |
| <code>build_differentiation_instruments(...[, ...])</code> | Construct excluded differentiation instruments. |
| <code>build_id_data(T, J, F)</code> | Build a balanced panel of market and firm IDs. |
| <code>build_ownership(product_data[, ...])</code> | Build ownership matrices, <i>O</i> . |
| <code>build_integration(integration, dimensions)</code> | Build nodes and weights for integration over agent choice probabilities. |

5.2.1 pyblp.build_matrix

`pyblp.build_matrix(formulation, data)`
Construct a matrix according to a formulation.

Parameters

- **formulation** (*Formulation*) – *Formulation* configuration for the matrix. Variable names should correspond to fields in *data*. The `absorb` argument of *Formulation* can be used to absorb fixed effects after the matrix has been constructed.
- **data** (*structured array-like*) – Fields can be used as variables in *formulation*.

Returns The built matrix.

Return type *ndarray*

Examples

The [online version](#) of the following section may be easier to read.

Building a Matrix Example

```
import pyblp
import pandas as pd
```

```
pyblp.__version__
```

```
'0.8.1'
```

In this example, we'll load the fake cereal data from *Nevo (2000)* and create a simple matrix involving a constant, prices, and shares.

```
formulation = pyblp.Formulation('1 + prices + shares')
```

```
formulation
```

```
1 + prices + shares
```

```
product_data = pd.read_csv(pyblp.data.NEVO_PRODUCTS_LOCATION)
```

```
product_data.head()
```

| | market_ids | city_ids | quarter | product_ids | firm_ids | brand_ids | shares | \ |
|---|------------|----------|---------|-------------|----------|-----------|----------|---|
| 0 | C01Q1 | 1 | 1 | F1B04 | 1 | 4 | 0.012417 | |
| 1 | C01Q1 | 1 | 1 | F1B06 | 1 | 6 | 0.007809 | |
| 2 | C01Q1 | 1 | 1 | F1B07 | 1 | 7 | 0.012995 | |
| 3 | C01Q1 | 1 | 1 | F1B09 | 1 | 9 | 0.005770 | |
| 4 | C01Q1 | 1 | 1 | F1B11 | 1 | 11 | 0.017934 | |

| | prices | sugar | mushy | ... | demand_instruments10 | demand_instruments11 | \ |
|---|----------|-------|-------|-----|----------------------|----------------------|---|
| 0 | 0.072088 | 2 | 1 | ... | 2.116358 | -0.154708 | |
| 1 | 0.114178 | 18 | 1 | ... | -7.374091 | -0.576412 | |
| 2 | 0.132391 | 4 | 1 | ... | 2.187872 | -0.207346 | |
| 3 | 0.130344 | 3 | 0 | ... | 2.704576 | 0.040748 | |
| 4 | 0.154823 | 12 | 0 | ... | 1.261242 | 0.034836 | |

| | demand_instruments12 | demand_instruments13 | demand_instruments14 | \ |
|---|----------------------|----------------------|----------------------|---|
| 0 | -0.005796 | 0.014538 | 0.126244 | |
| 1 | 0.012991 | 0.076143 | 0.029736 | |

(continues on next page)

(continued from previous page)

```

2          0.003509          0.091781          0.163773
3         -0.003724          0.094732          0.135274
4         -0.000568          0.102451          0.130640

   demand_instruments15 demand_instruments16 demand_instruments17 \
0          0.067345          0.068423          0.034800
1          0.087867          0.110501          0.087784
2          0.111881          0.108226          0.086439
3          0.088090          0.101767          0.101777
4          0.084818          0.101075          0.125169

   demand_instruments18 demand_instruments19
0          0.126346          0.035484
1          0.049872          0.072579
2          0.122347          0.101842
3          0.110741          0.104332
4          0.133464          0.121111

[5 rows x 30 columns]

```

```

matrix = pyblp.build_matrix(formulation, product_data)
matrix
array([[1.          , 0.07208794, 0.01241721],
       [1.          , 0.11417849, 0.00780939],
       [1.          , 0.13239066, 0.01299451],
       ...,
       [1.          , 0.13701741, 0.00222918],
       [1.          , 0.10017433, 0.01146267],
       [1.          , 0.12755747, 0.02620832]])

```

For various reasons, we may want to absorb fixed effects into the matrix. This can be done with the `absorb` argument of *Formulation*. We'll now re-create the matrix, absorbing product-specific fixed effects. Note that the constant column is now ignored.

```

absorb_formulation = pyblp.Formulation('prices + shares', absorb='product_ids')
absorb_formulation
prices + shares + Absorb[product_ids]

```



```
demeaned_matrix = pyblp.build_matrix(absorb_formulation, product_data)
demeaned_matrix
```

```
array([[ -0.01124832, -0.00052161],
       [ -0.00713476, -0.03144549],
       [  0.02367765, -0.01664996],
       ...,
       [  0.03371995, -0.00779841],
       [-0.00417404, -0.0117508 ],
       [-0.01195648,  0.00666695]])
```

5.2.2 pyblp.build_blp_instruments

`pyblp.build_blp_instruments` (*formulation, product_data*)

Construct “sums of characteristics” excluded BLP instruments.

Traditional “sums of characteristics” BLP instruments are

$$Z^{BLP}(X) = [Z^{BLP,Other}(X), Z^{BLP,Rival}(X)], \quad (5.1)$$

in which X is a matrix of product characteristics, $Z^{BLP,Other}(X)$ is a second matrix that consists of sums over characteristics of non-rival goods, and $Z^{BLP,Rival}(X)$ is a third matrix that consists of sums over rival goods. All three matrices have the same dimensions.

Note: To construct simpler, firm-agnostic instruments that are sums over characteristics of other goods, specify a constant column of firm IDs and keep only the first half of the instrument columns.

Let x_{jt} be the vector of characteristics in X for product j in market t , which is produced by firm f . That is, $j \in \mathcal{J}_{ft}$. Then,

$$\begin{aligned} Z_{jt}^{BLP,Other}(X) &= \sum_{k \in \mathcal{J}_{ft} \setminus \{j\}} x_{kt}, \\ Z_{jt}^{BLP,Rival}(X) &= \sum_{k \notin \mathcal{J}_{ft}} x_{kt}. \end{aligned} \quad (5.2)$$

Note: Usually, any supply or demand shifters are added to these excluded instruments, depending on whether they are meant to be used for demand- or supply-side estimation.

Parameters

- **formulation** (*Formulation*) – *Formulation* configuration for X , the matrix of product characteristics used to build excluded instruments. Variable names should correspond to fields in `product_data`.
- **product_data** (*structured array-like*) – Each row corresponds to a product. Markets can have differing numbers of products. The following fields are required:
 - **market_ids** : (*object*) - IDs that associate products with markets.
 - **firm_ids** : (*object*) - IDs that associate products with firms.

Along with `market_ids` and `firm_ids`, the names of any additional fields can be used as variables in `formulation`.

Returns Traditional “sums of characteristics” BLP instruments, $Z^{BLP}(X)$.

Return type `ndarray`

Examples

The [online version](#) of the following section may be easier to read.

Building “Sums of Characteristics” BLP Instruments Example

```
import pyblp
import pandas as pd
```

```
pyblp.__version__
```

```
'0.8.1'
```

In this example, we’ll load the automobile product data from *Berry, Levinsohn, and Pakes (1995)* and build some very simple excluded demand-side instruments for the problem. These instruments are different from the pre-built ones included in the automobile product data file, which used principal component analysis to alleviate a collinearity problem

```
formulation = pyblp.Formulation('1 + hpwt + air + mpd + space')
formulation
```

```
1 + hpwt + air + mpd + space
```

```
product_data = pd.read_csv(pyblp.data.BLP_PRODUCTS_LOCATION)
product_data.head()
```

| | market_ids | clustering_ids | car_ids | firm_ids | region | shares | prices | \ |
|---|------------|----------------|---------|----------|--------|----------|----------|---|
| 0 | 1971 | AMGREG71 | 129 | 15 | US | 0.001051 | 4.935802 | |
| 1 | 1971 | AMHORN71 | 130 | 15 | US | 0.000670 | 5.516049 | |
| 2 | 1971 | AMJAVL71 | 132 | 15 | US | 0.000341 | 7.108642 | |
| 3 | 1971 | AMMATA71 | 134 | 15 | US | 0.000522 | 6.839506 | |
| 4 | 1971 | AMAMBS71 | 136 | 15 | US | 0.000442 | 8.928395 | |

| | hpwt | air | mpd | ... | supply_instruments11 | supply_instruments12 | \ |
|---|----------|-----|----------|-----|----------------------|----------------------|---|
| 0 | 0.528997 | 0 | 1.888146 | ... | 2.024221 | -0.251338 | |
| 1 | 0.494324 | 0 | 1.935989 | ... | 2.026706 | -0.205305 | |
| 2 | 0.467613 | 0 | 1.716799 | ... | 1.882594 | -0.417634 | |
| 3 | 0.426540 | 0 | 1.687871 | ... | 1.846079 | -0.445630 | |
| 4 | 0.452489 | 0 | 1.504286 | ... | 1.760927 | -0.610203 | |


```
supply_instruments13  supply_instruments14  supply_instruments15  \
```

(continues on next page)

(continued from previous page)

```
0          0.741272          0.820281          3.691881
1          0.675468          1.004707          3.628727
2          0.946640          0.429942          4.001801
3          0.982922          0.392283          4.061161
4          1.201745          -0.085373         4.340853

  supply_instruments16  supply_instruments17  supply_instruments18  \
0          0.566070          -2.328128          0.385197
1          0.605683          -2.268975          0.363954
2          0.253461          -2.581003         -0.113763
3          0.183331          -2.626074         -0.253065
4          -0.043107         -2.785686         -0.556135

  supply_instruments19  supply_instruments20
0          0.648526          1.052606
1          0.643205          0.906429
2          0.770387          1.511333
3          0.802037          1.583208
4          0.931611          2.177230
```

```
[5 rows x 47 columns]
```

```
instruments = pyblp.build_blp_instruments(formulation, product_data)
instruments.shape
```

```
(2217, 10)
```

5.2.3 pyblp.build_differentiation_instruments

`pyblp.build_differentiation_instruments` (*formulation*, *product_data*, *version='local'*, *interact=False*)

Construct excluded differentiation instruments.

Differentiation instruments in the spirit of *Gandhi and Houde (2017)* are

$$Z^{Diff}(X) = [Z^{Diff,Other}(X), Z^{Diff,Rival}(X)], \quad (5.3)$$

in which X is a matrix of product characteristics, $Z^{Diff,Other}(X)$ is a second matrix that consists of sums over functions of differences between non-rival goods, and $Z^{Diff,Rival}(X)$ is a third matrix that consists of sums over rival goods. Without optional interaction terms, all three matrices have the same dimensions.

Note: To construct simpler, firm-agnostic instruments that are sums over functions of differences between all different goods, specify a constant column of firm IDs and keep only the first half of the instrument columns.

Let $x_{j t \ell}$ be characteristic ℓ in X for product j in market t , which is produced by firm f . That is, $j \in \mathcal{J}_{f t}$. Then in the “local” version of $Z^{Diff}(X)$,

$$\begin{aligned} Z_{j t \ell}^{Diff,Other,Local}(X) &= \sum_{k \in \mathcal{J}_{f t} \setminus \{j\}} 1(|d_{j k t \ell}| < SD_{\ell}), \\ Z_{j t \ell}^{Diff,Rival,Local}(X) &= \sum_{k \notin \mathcal{J}_{f t}} 1(|d_{j k t \ell}| < SD_{\ell}), \end{aligned} \quad (5.4)$$

where $d_{j k t \ell} = x_{k t \ell} - x_{j t \ell}$ is the difference between products j and k in terms of characteristic ℓ , SD_{ℓ} is the standard deviation of these pairwise differences computed across all markets, and $1(|d_{j k t \ell}| < SD_{\ell})$ indicates that products j and k are close to each other in terms of characteristic ℓ .

The intuition behind this “local” version is that demand for products is often most influenced by a small number of other goods that are very similar. For the “quadratic” version of $Z^{Diff}(X)$, which uses a more continuous measure of the distance between goods,

$$\begin{aligned} Z_{j t k}^{Diff,Other,Quad}(X) &= \sum_{k \in \mathcal{J}_{f t} \setminus \{j\}} d_{j k t \ell}^2, \\ Z_{j t k}^{Diff,Rival,Quad}(X) &= \sum_{k \notin \mathcal{J}_{f t}} d_{j k t \ell}^2. \end{aligned} \quad (5.5)$$

With interaction terms, which reflect covariances between different characteristics, the summands for the “local” versions are $1(|d_{j k t \ell}| < SD_{\ell}) \times d_{j k t \ell'}$ for all characteristics ℓ' , and the summands for the “quadratic” versions are $d_{j k t \ell} \times d_{j k t \ell'}$ for all $\ell' \geq \ell$.

Note: Usually, any supply or demand shifters are added to these excluded instruments, depending on whether they are meant to be used for demand- or supply-side estimation.

Parameters

- **formulation** (*Formulation*) – *Formulation* configuration for X , the matrix of product characteristics used to build excluded instruments. Variable names should correspond to fields in `product_data`.
- **product_data** (*structured array-like*) – Each row corresponds to a product. Markets can have differing numbers of products. The following fields are required:

- **market_ids** : (*object*) - IDs that associate products with markets.
- **firm_ids** : (*object*) - IDs that associate products with firms.

Along with `market_ids` and `firm_ids`, the names of any additional fields can be used as variables in `formulation`.

- **version** (*str, optional*) – The version of differentiation instruments to construct:
 - 'local' (default) - Construct the instruments in (5.4) that consider only the characteristics of “close” products in each market.
 - 'quadratic' - Construct the more continuous instruments in (5.5) that consider all products in each market.
- **interact** (*bool, optional*) – Whether to include interaction terms between different product characteristics, which can help capture covariances between product characteristics.

Returns Excluded differentiation instruments, $Z^{Diff}(X)$.

Return type *ndarray*

Examples

The [online version](#) of the following section may be easier to read.

Building Differentiation Instruments Example

```
import pyblp
import pandas as pd
```

```
pyblp.__version__
```

```
'0.8.1'
```

In this example, we'll load the automobile product data from *Berry, Levinsohn, and Pakes (1995)* and build some very simple excluded demand-side instruments for the problem in the spirit of *Gandhi and Houde (2017)*.

```
formulation = pyblp.Formulation('0 + hpwt + air + mpd + space')
```

```
formulation
```

```
hpwt + air + mpd + space
```

```
product_data = pd.read_csv(pyblp.data.BLP_PRODUCTS_LOCATION)
```

```
product_data.head()
```

| | market_ids | clustering_ids | car_ids | firm_ids | region | shares | prices | \ |
|---|------------|----------------|---------|----------|--------|----------|----------|---|
| 0 | 1971 | AMGREM71 | 129 | 15 | US | 0.001051 | 4.935802 | |
| 1 | 1971 | AMHORN71 | 130 | 15 | US | 0.000670 | 5.516049 | |
| 2 | 1971 | AMJAVL71 | 132 | 15 | US | 0.000341 | 7.108642 | |
| 3 | 1971 | AMMATA71 | 134 | 15 | US | 0.000522 | 6.839506 | |
| 4 | 1971 | AMAMBS71 | 136 | 15 | US | 0.000442 | 8.928395 | |

| | hpwt | air | mpd | ... | supply_instruments11 | supply_instruments12 | \ |
|---|----------|-----|----------|-----|----------------------|----------------------|---|
| 0 | 0.528997 | 0 | 1.888146 | ... | 2.024221 | -0.251338 | |
| 1 | 0.494324 | 0 | 1.935989 | ... | 2.026706 | -0.205305 | |
| 2 | 0.467613 | 0 | 1.716799 | ... | 1.882594 | -0.417634 | |
| 3 | 0.426540 | 0 | 1.687871 | ... | 1.846079 | -0.445630 | |
| 4 | 0.452489 | 0 | 1.504286 | ... | 1.760927 | -0.610203 | |

| | supply_instruments13 | supply_instruments14 | supply_instruments15 | \ |
|---|----------------------|----------------------|----------------------|---|
| 0 | 0.741272 | 0.820281 | 3.691881 | |
| 1 | 0.675468 | 1.004707 | 3.628727 | |

(continues on next page)

(continued from previous page)

```

2          0.946640          0.429942          4.001801
3          0.982922          0.392283          4.061161
4          1.201745          -0.085373          4.340853

  supply_instruments16  supply_instruments17  supply_instruments18  \
0          0.566070          -2.328128          0.385197
1          0.605683          -2.268975          0.363954
2          0.253461          -2.581003          -0.113763
3          0.183331          -2.626074          -0.253065
4          -0.043107          -2.785686          -0.556135

  supply_instruments19  supply_instruments20
0          0.648526          1.052606
1          0.643205          0.906429
2          0.770387          1.511333
3          0.802037          1.583208
4          0.931611          2.177230

[5 rows x 47 columns]
```

Note that we’re excluding the constant column because it yields collinear constant columns of differentiation instruments.

We’ll first build “local” differentiation instruments, which are constructed by default, and which consist of counts of “close” rival and non-rival products in each market.

```

local_instruments = pyblp.build_differentiation_instruments(
    formulation,
    product_data
)
local_instruments.shape

(2217, 8)
```

Next, we’ll build a more continuous “quadratic” version of the instruments, which consist of sums over squared differences between rival and non-rival products in each market.

```

quadratic_instruments = pyblp.build_differentiation_instruments(
    formulation,
    product_data,
    version='quadratic'
```

(continues on next page)

(continued from previous page)

```
)  
quadratic_instruments.shape  
(2217, 8)
```

We could also use `interact=True` to include interaction terms in either version of instruments, which would help capture covariances between different product characteristics.

5.2.4 pyblp.build_id_data

`pyblp.build_id_data(T, J, F)`

Build a balanced panel of market and firm IDs.

This function can be used to build `id_data` for *Simulation* initialization.

Parameters

- **T** (*int*) – Number of markets.
- **J** (*int*) – Number of products in each market.
- **F** (*int*) – Number of firms. If J is divisible by F , firms produce J / F products in each market. Otherwise, firms with smaller IDs will produce excess products.

Returns

IDs that associate products with markets and firms. Each of the $T * J$ rows corresponds to a product. Fields:

- **market_ids** : (*object*) - Market IDs that take on values from 0 to $T - 1$.
- **firm_ids** : (*object*) - Firm IDs that take on values from 0 to $F - 1$.

Return type *recarray*

Examples

The [online version](#) of the following section may be easier to read.

Building ID Data Example

```
import pyblp
import numpy as np

np.set_printoptions(linewidth=1)
pyblp.__version__

'0.8.1'
```

In this example, we'll build a small panel of market and firm IDs.

```
id_data = pyblp.build_id_data(T=2, J=5, F=4)
id_data

rec.array([(0, [0]),
          (0, [0]),
          (0, [1]),
          (0, [2]),
          (0, [3]),
          (1, [0]),
          (1, [0]),
          (1, [1]),
          (1, [2]),
          (1, [3])],
          dtype=[('market_ids', 'O', (1,)), ('firm_ids', 'O', (1,))])
```

5.2.5 pyblp.build_ownership

`pyblp.build_ownership` (*product_data*, *kappa_specification=None*)

Build ownership matrices, O .

Ownership matrices are defined by their cooperation matrix counterparts, κ . For each market t , $O_{jk} = \kappa_{fg}$ where $j \in \mathcal{J}_{ft}$, the set of products produced by firm f in the market, and similarly, $g \in \mathcal{J}_{gt}$.

Parameters

- **product_data** (*structured array-like*) – Each row corresponds to a product. Markets can have differing numbers of products. The following fields are required:
 - **market_ids** : (*object*) - IDs that associate products with markets.
 - **firm_ids** : (*object*) - IDs that associate products with firms.
- **kappa_specification** (*callable, optional*) – A function that specifies each market's cooperation matrix, κ . The function is of the following form:

```
kappa(f, g) -> value
```

where $value$ is O_{jk} and both f and g are firm IDs from the `firm_ids` field of `product_data`.

The default specification, `lambda: f, g: int(f == g)`, constructs traditional ownership matrices. That is, $\kappa = I$, the identify matrix, implies that O_{jk} is 1 if the same firm produces products j and k , and is 0 otherwise.

If `firm_ids` happen to be indices for an actual κ matrix, `lambda f, g: kappa[f, g]` will build ownership matrices according to the matrix `kappa`.

Returns Stacked $J_t \times J_t$ ownership matrices, O , for each market t . If a market has fewer products than others, extra columns will contain `numpy.nan`.

Return type `ndarray`

Examples

The [online version](#) of the following section may be easier to read.

Building Ownership Matrices Example

```
import pyblp
import numpy as np

np.set_printoptions(threshold=100)
pyblp.__version__

'0.8.1'
```

In this example, we'll use the IDs created in the *building ID data example* to build a stack of standard ownership matrices. We'll delete the first data row to demonstrate what ownership matrices should look like when markets have varying numbers of products.

```
id_data = pyblp.build_id_data(T=2, J=5, F=4)
id_data = id_data[1:]
standard_ownership = pyblp.build_ownership(id_data)
standard_ownership

array([[ 1.,  0.,  0.,  0., nan],
       [ 0.,  1.,  0.,  0., nan],
       [ 0.,  0.,  1.,  0., nan],
       [ 0.,  0.,  0.,  1., nan],
       [ 1.,  1.,  0.,  0.,  0.],
       [ 1.,  1.,  0.,  0.,  0.],
       [ 0.,  0.,  1.,  0.,  0.],
       [ 0.,  0.,  0.,  1.,  0.],
       [ 0.,  0.,  0.,  0.,  1.]])
```

We'll now modify the default κ specification so that the elements associated with firm IDs 0 and 1 are equal to 0.5.

```
def kappa_specification(f, g):
    if f == g:
        return 1
    return 0.5 if f < 2 and g < 2 else 0
```

Finally, we'll use this specification to build a stack of alternative ownership matrices.

```
alternative_ownership = pyblp.build_ownership(id_data, kappa_specification)
alternative_ownership
```

```
array([[1. , 0.5, 0. , 0. , nan],
       [0.5, 1. , 0. , 0. , nan],
       [0. , 0. , 1. , 0. , nan],
       [0. , 0. , 0. , 1. , nan],
       [1. , 1. , 0.5, 0. , 0. ],
       [1. , 1. , 0.5, 0. , 0. ],
       [0.5, 0.5, 1. , 0. , 0. ],
       [0. , 0. , 0. , 1. , 0. ],
       [0. , 0. , 0. , 0. , 1. ]])
```

5.2.6 pyblp.build_integration

`pyblp.build_integration` (*integration, dimensions*)

Build nodes and weights for integration over agent choice probabilities.

This function can be used to build custom `agent_data` for `Problem` initialization. Specifically, this function affords more flexibility than passing an `Integration` configuration directly to `Problem`. For example, if agents have unobserved tastes over only a subset of nonlinear product characteristics (i.e., if `sigma` in `Problem.solve()` has columns of zeros), this function can be used to build agent data with fewer columns of integration nodes than the number of unobserved product characteristics, K_2 . This function can also be used to construct nodes that can be transformed into demographic variables.

To build nodes and weights for multiple markets, this function can be called multiple times, once for each market.

Parameters

- **integration** (*Integration*) – `Integration` configuration for how to build nodes and weights for integration.
- **dimensions** (*int*) – Number of dimensions over which to integrate, or equivalently, the number of columns of integration nodes. When an `Integration` configuration is passed directly to `Problem`, this is the number of nonlinear product characteristics, K_2 .

Returns

Nodes and weights for integration over agent utilities. Fields:

- **weights** : (*numeric*) - Integration weights, w .
- **nodes** : (*numeric*) - Unobserved agent characteristics called integration nodes, ν .

Return type *recarray*

Examples

The [online version](#) of the following section may be easier to read.

Building Nodes and Weights for Integration Example

```
import pyblp

pyblp.__version__

'0.8.1'
```

In this example, we'll build nodes and weights for integration over agent choice probabilities according to a *Integration* configuration. We'll construct a sparse grid of nodes and weights according to a level-5 Gauss-Hermite quadrature rule.

```
integration = pyblp.Integration('grid', 5)
integration
```

Configured to construct nodes and weights in a sparse grid according to the level-5 Gauss-Hermite rule.

Usually, this configuration should be passed directly to *Problem*, which will create a sparse grid of dimension K_2 , the number of nonlinear product characteristics. Alternatively, we can build the sparse grid ourselves and pass the constructed agent data to *Problem*, possibly after modifying the nodes and weights. If we want to allow agents to have heterogeneous tastes over 2 product characteristics, we'll need a grid of dimension 2.

```
agent_data = pyblp.build_integration(integration, 2)
agent_data.nodes.shape

(53, 2)
```

```
agent_data.weights.shape

(53, 1)
```

If we wanted to construct nodes and weights for each market, we could call *build_integration* once for each market, add a column of market IDs, and stack the arrays.

5.3 Problem Class

Given data and appropriate configurations, a BLP-type problem can be structured by initializing the following class.

| | |
|--|---------------------|
| <code>Problem(product_formulations, product_data)</code> | A BLP-type problem. |
|--|---------------------|

5.3.1 pyblp.Problem

```
class pyblp.Problem(product_formulations, product_data, agent_formulation=None,
                    agent_data=None, integration=None, costs_type='linear')
```

A BLP-type problem.

This class is initialized with relevant data and solved with `Problem.solve()`.

Parameters

- **product_formulations** (*Formulation or sequence of Formulation*) – *Formulation* configuration or a sequence of up to three *Formulation* configurations for the matrix of linear product characteristics, X_1 , for the matrix of nonlinear product characteristics, X_2 , and for the matrix of cost characteristics, X_3 , respectively. If the formulation for X_3 is not specified or is `None`, a supply side will not be estimated. Similarly, if the formulation for X_2 is not specified or is `None`, the logit (or nested logit) model will be estimated.

Variable names should correspond to fields in `product_data`. The `shares` variable should not be included in any of the formulations and `prices` should be included in the formulation for X_1 or X_2 (or both). The `absorb` argument of *Formulation* can be used to absorb fixed effects into X_1 and X_3 , but not X_2 . Characteristics in X_2 should generally be included in X_1 . The typical exception is characteristics that are collinear with fixed effects that have been absorbed into X_1 .

Characteristics in X_1 that do not involve `prices`, X_1^f , will be combined with excluded demand-side instruments (specified below) to create the full set of demand-side instruments, Z_D . Any fixed effects absorbed into X_1 will also be absorbed into Z_D . Similarly, characteristics in X_3 will be combined with the excluded supply-side instruments to create Z_S , and any fixed effects absorbed into X_3 will also be absorbed into Z_S .

Warning: Characteristics that involve prices, p , should always be formulated with the `prices` variable. If another name is used, *Problem* will not understand that the characteristic is endogenous, so it will be erroneously included in Z_D , and derivatives computed with respect to prices will likely be wrong. For example, to include a p^2 characteristic, include `I(prices**2)` in a formula instead of manually including a `prices_squared` variable in `product_data` and a formula.

- **product_data** (*structured array-like*) – Each row corresponds to a product. Markets can have differing numbers of products. The following fields are required:
 - **market_ids** : (*object*) - IDs that associate products with markets.
 - **shares** : (*numeric*) - Marketshares, s , which should be between zero and one, exclusive. Outside shares should also be between zero and one. Shares in each market should sum to less than one.
 - **prices** : (*numeric*) - Product prices, p .

If a formulation for X_3 is specified in `product_formulations`, firm IDs are also required, since they will be used to estimate the supply side of the problem:

- **firm_ids** : (*object, optional*) - IDs that associate products with firms.

Excluded instruments should generally be specified with the following fields:

- **demand_instruments** : (*numeric*) - Excluded demand-side instruments, which, together with the formulated exogenous linear product characteristics, X_1^x , constitute the full set of demand-side instruments, Z_D .
- **supply_instruments** : (*numeric, optional*) - Excluded supply-side instruments, which, together with the formulated cost characteristics, X_3 , constitute the full set of supply-side instruments, Z_S .

The recommendation in *Conlon and Gortmaker (2019)* is to start with differentiation instruments of *Gandhi and Houde (2017)*, which can be built with `build_differentiation_instruments()`, and then compute feasible optimal instruments with `ProblemResults.compute_optimal_instruments()` in the second stage.

If `firm_ids` are specified, custom ownership matrices can be specified as well:

- **ownership** : (*numeric, optional*) - Custom stacked $J_t \times J_t$ ownership matrices, O , for each market t , which can be built with `build_ownership()`. By default, standard ownership matrices are built only when they are needed to reduce memory usage. If specified, there should be as many columns as there are products in the market with the most products. Rightmost columns in markets with fewer products will be ignored.

Note: Fields that can have multiple columns (`demand_instruments`, `supply_instruments`, and `ownership`) can either be matrices or can be broken up into multiple one-dimensional fields with column index suffixes that start at zero. For example, if there are three columns of excluded demand-side instruments, a `demand_instruments` field with three columns can be replaced by three one-dimensional fields: `demand_instruments0`, `demand_instruments1`, and `demand_instruments2`.

To estimate a nested logit or random coefficients nested logit (RCNL) model, nesting groups must be specified:

- **nesting_ids** (*object, optional*) - IDs that associate products with nesting groups. When these IDs are specified, `rho` must be specified in `Problem.solve()` as well.

Finally, clustering groups can be specified to account for within-group correlation while updating the weighting matrix and estimating standard errors:

- **clustering_ids** (*object, optional*) - Cluster group IDs, which will be used if `W_type` or `se_type` in `Problem.solve()` is 'clustered'.

Along with `market_ids`, `firm_ids`, `nesting_ids`, `clustering_ids`, and `prices`, the names of any additional fields can typically be used as variables in `product_formulations`. However, there are a few variable names such as 'X1', which are reserved for use by *Products*.

- **agent_formulation** (*Formulation, optional*) – *Formulation* configuration for the matrix of observed agent characteristics called demographics, d , which will only be included in the model if this formulation is specified. Since demographics are only used if there are nonlinear product characteristics, this formulation should only be specified if X_2 is

formulated in `product_formulations`. Variable names should correspond to fields in `agent_data`.

- **agent_data** (*structured array-like, optional*) – Each row corresponds to an agent. Markets can have differing numbers of agents. Since simulated agents are only used if there are nonlinear product characteristics, agent data should only be specified if X_2 is formulated in `product_formulations`. If agent data are specified, market IDs are required:
 - **market_ids** : (*object*) - IDs that associate agents with markets. The set of distinct IDs should be the same as the set in `product_data`. If `integration` is specified, there must be at least as many rows in each market as the number of nodes and weights that are built for the market.

If `integration` is not specified, the following fields are required:

- **weights** : (*numeric, optional*) - Integration weights, w , for integration over agent choice probabilities.
- **nodes** : (*numeric, optional*) - Unobserved agent characteristics called integration nodes, ν . If there are more than K_2 columns (the number of nonlinear product characteristics), only the first K_2 will be retained.

The convenience function `build_integration()` can be useful when constructing custom nodes and weights.

Note: If `nodes` has multiple columns, it can be specified as a matrix or broken up into multiple one-dimensional fields with column index suffixes that start at zero. For example, if there are three columns of nodes, a `nodes` field with three columns can be replaced by three one-dimensional fields: `nodes0`, `nodes1`, and `nodes2`.

Along with `market_ids`, the names of any additional fields can be typically be used as variables in `agent_formulation`. The exception is the name 'demographics', which is reserved for use by *Agents*.

- **integration** (*Integration, optional*) – `Integration` configuration for how to build nodes and weights for integration over agent choice probabilities, which will replace any `nodes` and `weights` fields in `agent_data`. This configuration is required if `nodes` and `weights` in `agent_data` are not specified. It should not be specified if X_2 is not formulated in `product_formulations`.

If this configuration is specified, K_2 columns of nodes (the number of nonlinear product characteristics) will be built. However, if `sigma` in `Problem.solve()` is left unspecified or specified with columns fixed at zero, fewer columns will be used.

- **costs_type** (*str, optional*) – Functional form of the marginal cost function $\tilde{c} = f(c)$ in (3.9). The following specifications are supported:
 - 'linear' (default) - Linear specification: $\tilde{c} = c$.
 - 'log' - Log-linear specification: $\tilde{c} = \log c$.

This specification is only relevant if X_3 is formulated.

product_formulations

Formulation configurations for X_1 , X_2 , and X_3 , respectively.

Type *Formulation or sequence of Formulation*

agent_formulation

Formulation configuration for d .

Type Formulation**products**

Product data structured as *Products*, which consists of data taken from `product_data` along with matrices built according to `Problem.product_formulations`.

Type Products**agents**

Agent data structured as *Agents*, which consists of data taken from `agent_data` or built by integration along with any demographics built according to `Problem.agent_formulation`.

Type Agents**unique_market_ids**

Unique market IDs in product and agent data.

Type ndarray**unique_firm_ids**

Unique firm IDs in product data.

Type ndarray**unique_nesting_ids**

Unique nesting group IDs in product data.

Type ndarray**costs_type**

Functional form of the marginal cost function $\tilde{c} = f(c)$.

Type str**T**

Number of markets, T .

Type int**N**

Number of products across all markets, N .

Type int**F**

Number of firms across all markets, F .

Type int**I**

Number of agents across all markets, I .

Type int**K1**

Number of linear product characteristics, K_1 .

Type int**K2**

Number of nonlinear product characteristics, K_2 .

Type int**K3**

Number of cost product characteristics, K_3 .

Type *int***D**Number of demographic variables, D .**Type** *int***MD**Number of demand-side instruments, M_D , which is the number of excluded demand-side instruments plus the number of exogenous linear product characteristics, K_1^x .**Type** *int***MS**Number of supply-side instruments, M_S , which is the number of excluded supply-side instruments plus the number of cost product characteristics, K_3 .**Type** *int***ED**Number of absorbed dimensions of demand-side fixed effects, E_D .**Type** *int***ES**Number of absorbed dimensions of supply-side fixed effects, E_S .**Type** *int***H**Number of nesting groups, H .**Type** *int*

Examples

- [Tutorial](#)

Methods

| | |
|--|--------------------|
| <code>solve([sigma, pi, rho, beta, gamma, ...])</code> | Solve the problem. |
|--|--------------------|

Once initialized, the following method solves the problem.

| | |
|--|--------------------|
| <code>Problem.solve([sigma, pi, rho, beta, gamma, ...])</code> | Solve the problem. |
|--|--------------------|

5.3.2 pyblp.Problem.solve

`Problem.solve` (*sigma=None, pi=None, rho=None, beta=None, gamma=None, sigma_bounds=None, pi_bounds=None, rho_bounds=None, beta_bounds=None, gamma_bounds=None, delta=None, method='2s', optimization=None, check_optimality='both', error_behavior='revert', error_punishment=1, delta_behavior='first', iteration=None, fp_type='safe_linear', costs_bounds=None, W=None, center_moments=True, W_type='robust', se_type='robust', micro_moments=(), extra_micro_covariances=None*)

Solve the problem.

The problem is solved in one or more GMM steps. During each step, any parameters in $\hat{\theta}$ are optimized to minimize the GMM objective value. If there are no parameters in $\hat{\theta}$ (for example, in the logit model there are no nonlinear parameters and all linear parameters can be concentrated out), the objective is evaluated once during the step.

If there are nonlinear parameters, the mean utility, $\delta(\hat{\theta})$ is computed market-by-market with fixed point iteration. Otherwise, it is computed analytically according to the solution of the logit model. If a supply side is to be estimated, marginal costs, $c(\hat{\theta})$, are also computed market-by-market. Linear parameters are then estimated, which are used to recover structural error terms, which in turn are used to form the objective value. By default, the objective gradient is computed as well.

Note: This method supports `parallel()` processing. If multiprocessing is used, market-by-market computation of $\delta(\hat{\theta})$ (and $\tilde{c}(\hat{\theta})$ if a supply side is estimated), along with associated Jacobians, will be distributed among the processes.

Parameters

- **sigma** (*array-like, optional*) – Configuration for which elements in the lower-triangular Cholesky root of the covariance matrix for unobserved taste heterogeneity, Σ , are fixed at zero and starting values for the other elements, which, if not fixed by `sigma_bounds`, are in the vector of unknown elements, θ .

Rows and columns correspond to columns in X_2 , which is formulated according `product_formulations` in *Problem*. If X_2 was not formulated, this should not be specified, since the logit model will be estimated.

Values above the diagonal are ignored. Zeros are assumed to be zero throughout estimation and nonzeros are, if not fixed by `sigma_bounds`, starting values for unknown elements in θ . If any columns are fixed at zero, only the first few columns of integration nodes (specified in *Problem*) will be used.

- **pi** (*array-like, optional*) – Configuration for which elements in the matrix of parameters that measures how agent tastes vary with demographics, Π , are fixed at zero and starting values for the other elements, which, if not fixed by `pi_bounds`, are in the vector of unknown elements, θ .

Rows correspond to the same product characteristics as in `sigma`. Columns correspond to columns in d , which is formulated according to `agent_formulation` in *Problem*. If d was not formulated, this should not be specified.

Zeros are assumed to be zero throughout estimation and nonzeros are, if not fixed by `pi_bounds`, starting values for unknown elements in θ .

- **rho** (*array-like, optional*) – Configuration for which elements in the vector of parameters that measure within nesting group correlation, ρ , are fixed at zero and starting values for the other elements, which, if not fixed by `rho_bounds`, are in the vector of unknown elements, θ .

If this is a scalar, it corresponds to all groups defined by the `nesting_ids` field of `product_data` in *Problem*. If this is a vector, it must have H elements, one for each nesting group. Elements correspond to group IDs in the sorted order of *Problem.unique_nesting_ids*. If nesting IDs were not specified, this should not be specified either.

Zeros are assumed to be zero throughout estimation and nonzeros are, if not fixed by `rho_bounds`, starting values for unknown elements in θ .

- **beta** (*array-like, optional*) – Configuration for which elements in the vector of demand-side linear parameters, β , are concentrated out of the problem. Usually, this is left unspecified, unless there is a supply side, in which case parameters on endogenous product characteristics cannot be concentrated out of the problem. Values specify which elements are fixed at zero and starting values for the other elements, which, if not fixed by `beta_bounds`, are in the vector of unknown elements, θ .

Elements correspond to columns in X_1 , which is formulated according to `product_formulations` in *Problem*.

Both `None` and `numpy.nan` indicate that the parameter should be concentrated out of the problem. That is, it will be estimated, but does not have to be included in θ . Zeros are assumed to be zero throughout estimation and nonzeros are, if not fixed by `beta_bounds`, starting values for unknown elements in θ .

- **gamma** (*array-like, optional*) – Configuration for which elements in the vector of supply-side linear parameters, γ , are concentrated out of the problem. Usually, this is left unspecified. Values specify which elements are fixed at zero and starting values for the other elements, which, if not fixed by `gamma_bounds`, are in the vector of unknown elements, θ .

Elements correspond to columns in X_3 , which is formulated according to `product_formulations` in *Problem*. If X_3 was not formulated, this should not be specified.

Both `None` and `numpy.nan` indicate that the parameter should be concentrated out of the problem. That is, it will be estimated, but does not have to be included in θ . Zeros are assumed to be zero throughout estimation and nonzeros are, if not fixed by `gamma_bounds`, starting values for unknown elements in θ .

- **sigma_bounds** (*tuple, optional*) – Configuration for Σ bounds of the form `(lb, ub)`, in which both `lb` and `ub` are of the same size as `sigma`. Each element in `lb` and `ub` determines the lower and upper bound for its counterpart in `sigma`. If `optimization` does not support bounds, these will be ignored.

By default, if bounds are supported, the diagonal of `sigma` is bounded from below by zero. Conditional on X_2 , μ , and an initial estimate of μ , default bounds for off-diagonal parameters are chosen to reduce the need for overflow safety precautions.

Values above the diagonal are ignored. Lower and upper bounds corresponding to zeros in `sigma` are set to zero. Setting a lower bound equal to an upper bound fixes the corresponding element, removing it from θ . Both `None` and `numpy.nan` are converted to `-numpy.inf` in `lb` and to `numpy.inf` in `ub`.

- **pi_bounds** (*tuple, optional*) – Configuration for Π bounds of the form `(lb, ub)`, in which both `lb` and `ub` are of the same size as `pi`. Each element in `lb` and `ub` determines the lower and upper bound for its counterpart in `pi`. If `optimization` does not support bounds, these will be ignored.

By default, if bounds are supported, conditional on X_2 , d , and an initial estimate of μ , default bounds are chosen to reduce the need for overflow safety precautions.

Lower and upper bounds corresponding to zeros in `pi` are set to zero. Setting a lower bound equal to an upper bound fixes the corresponding element, removing it from θ . Both `None` and `numpy.nan` are converted to `-numpy.inf` in `lb` and to `numpy.inf` in `ub`.

- **rho_bounds** (*tuple, optional*) – Configuration for ρ bounds of the form `(lb, ub)`, in which both `lb` and `ub` are of the same size as `rho`. Each element in `lb` and `ub` determines the lower and upper bound for its counterpart in `rho`. If `optimization` does not support bounds, these will be ignored.

By default, if bounds are supported, all elements are bounded from below by 0, which corresponds to the simple logit model. Conditional on an initial estimate of μ , upper bounds are chosen to reduce the need for overflow safety precautions, and are less than 1 because larger values are inconsistent with utility maximization.

Lower and upper bounds corresponding to zeros in `rho` are set to zero. Setting a lower bound equal to an upper bound fixes the corresponding element, removing it from θ . Both `None` and `numpy.nan` are converted to `-numpy.inf` in `lb` and to `numpy.inf` in `ub`.

- **beta_bounds** (*tuple, optional*) – Configuration for β bounds of the form `(lb, ub)`, in which both `lb` and `ub` are of the same size as `beta`. Each element in `lb` and `ub` determines the lower and upper bound for its counterpart in `beta`. If `optimization` does not support bounds, these will be ignored.

Usually, this is left unspecified, unless there is a supply side, in which case parameters on endogenous product characteristics cannot be concentrated out of the problem. It is generally a good idea to constrain such parameters to be nonzero so that the intra-firm Jacobian of shares with respect to prices does not become singular.

By default, all non-concentrated out parameters are unbounded. Bounds should only be specified for parameters that are included in θ ; that is, those with initial values specified in `beta`.

Lower and upper bounds corresponding to zeros in `beta` are set to zero. Setting a lower bound equal to an upper bound fixes the corresponding element, removing it from θ . Both `None` and `numpy.nan` are converted to `-numpy.inf` in `lb` and to `numpy.inf` in `ub`.

- **gamma_bounds** (*tuple, optional*) – Configuration for γ bounds of the form `(lb, ub)`, in which both `lb` and `ub` are of the same size as `gamma`. Each element in `lb` and `ub` determines the lower and upper bound for its counterpart in `gamma`. If `optimization` does not support bounds, these will be ignored.

By default, all non-concentrated out parameters are unbounded. Bounds should only be specified for parameters that are included in θ ; that is, those with initial values specified in `gamma`.

Lower and upper bounds corresponding to zeros in `gamma` are set to zero. Setting a lower bound equal to an upper bound fixes the corresponding element, removing it from θ . Both `None` and `numpy.nan` are converted to `-numpy.inf` in `lb` and to `numpy.inf` in `ub`.

- **delta** (*array-like, optional*) – Initial values for the mean utility, δ . If there are any non-linear parameters, these are the values at which the fixed point iteration routine will start during the first objective evaluation. By default, the solution to the logit model in (3.40) is used. If ρ is specified, the solution to the nested logit model in (3.41) under the initial `rho` is used instead.

- **method** (*str, optional*) – The estimation routine that will be used. The following methods are supported:

- `'1s'` - One-step GMM.
- `'2s'` (default) - Two-step GMM.

Iterated GMM can be manually implemented by executing single GMM steps in a loop, in which after the first iteration, nonlinear parameters and weighting matrices from the last `ProblemResults` are passed as arguments.

- **optimization** (*Optimization, optional*) – `Optimization` configuration for how to solve the optimization problem in each GMM step, which is only used if there are unfixed nonlinear parameters over which to optimize. By default, `Optimization('l-bfgs-b', {'ftol': 0, 'gtol': 1e-8})` is used. If

available, `Optimization('knitro')` may be preferable. Generally, it is recommended to consider a number of different optimization routines and starting values, verifying that $\hat{\theta}$ satisfies both the first and second order conditions. Routines that do not support bounds will ignore `sigma_bounds` and `pi_bounds`. Choosing a routine that does not use analytic gradients will often down estimation.

- **check_optimality** (*str, optional*) – How to check for optimality (first and second order conditions) after the optimization routine finishes. The following configurations are supported:
 - 'gradient' - Analytically compute the gradient after optimization finishes, but do not compute the Hessian. Since Jacobians needed to compute standard errors will already be computed, gradient computation will not take a long time. This option may be useful if Hessian computation takes a long time when, for example, there are a large number of parameters.
 - 'both' (default) - Also compute the Hessian with central finite differences after optimization finishes. Specifically, analytically compute the gradient $2P$ times, perturbing each of the P parameters by $\pm\sqrt{\epsilon^{mach}/2}$ where ϵ^{mach} is the machine precision.
- **error_behavior** (*str, optional*) – How to handle any errors. For example, there can sometimes be overflow or underflow when computing $\delta(\hat{\theta})$ at a large $\hat{\theta}$. The following behaviors are supported:
 - 'revert' (default) - Revert problematic values to their last computed values. If there are problematic values during the first objective evaluation, revert values in $\delta(\hat{\theta})$ to their starting values; in $\tilde{c}(\hat{\theta})$, to prices; in the objective, to $1e10$; and in other matrices such as Jacobians, to zeros.
 - 'punish' - Set the objective to 1 and its gradient to all zeros. This option along with a large `error_punishment` can be helpful for routines that do not use analytic gradients.
 - 'raise' - Raise an exception.
- **error_punishment** (*float, optional*) – How to scale the GMM objective value after an error. By default, the objective value is not scaled.
- **delta_behavior** (*str, optional*) – Configuration for the values at which the fixed point computation of $\delta(\hat{\theta})$ in each market will start. This configuration is only relevant if there are unfixed nonlinear parameters over which to optimize. The following behaviors are supported:
 - 'first' (default) - Start at the values configured by `delta` during the first GMM step, and at the values computed by the last GMM step for each subsequent step.
 - 'last' - Start at the values of $\delta(\hat{\theta})$ computed during the last objective evaluation, or, if this is the first evaluation, at the values configured by `delta`. This behavior tends to speed up computation but may introduce some instability into estimation.
- **iteration** (*Iteration, optional*) – *Iteration* configuration for how to solve the fixed point problem used to compute $\delta(\hat{\theta})$ in each market. This configuration is only relevant if there are nonlinear parameters, since δ can be estimated analytically in the logit model. By default, `Iteration('squarem', {'atol': 1e-14})` is used. Newton-based routines such as `Iteration('lm'')` that compute the Jacobian can often be faster (especially when there are nesting parameters), but the non-Jacobian SQUAREM routine is used by default because its speed is often comparable and in practice it can be slightly more stable.

- **fp_type** (*str, optional*) – Configuration for the type of contraction mapping used to compute $\delta(\hat{\theta})$. The following types are supported:
 - 'safe_linear' (default) - The standard linear contraction mapping in (3.13) (or (3.39) when there is nesting) with safeguards against numerical overflow. Specifically, $\max_j V_{jti}$ (or $\max_j V_{jti}/(1 - \rho_h(j))$ when there is nesting) is subtracted from V_{jti} and the logit expression for choice probabilities in (3.5) (or (3.37)) is re-scaled accordingly. Such re-scaling is known as the log-sum-exp trick.
 - 'linear' - The standard linear contraction mapping without safeguards against numerical overflow. This option may be preferable to 'safe_linear' if utilities are reasonably small and unlikely to create overflow problems.
 - 'nonlinear' - Iteration over $\exp(\delta_{jt})$ instead of δ_{jt} . This can be faster than 'linear' because it involves fewer logarithms. Also, following *Brunner, Heiss, Romahn, and Weiser (2017)*, the $\exp(\delta_{jt})$ term can be cancelled out of the expression because it also appears in the numerator of (3.5) in the definition of $s_{jt}(\delta, \hat{\theta})$. This second trick only works when there are no nesting parameters.
 - 'safe_nonlinear' - Exponentiated version with minimal safeguards against numerical overflow. Specifically, $\max_j \mu_{jti}$ is subtracted from μ_{jti} . This helps with stability but is less helpful than subtracting from the full V_{jti} , so this version is less stable than 'safe_linear'.

This option is only relevant if `sigma` or `pi` are specified because δ can be estimated analytically in the logit model with (3.40) and in the nested logit model with (3.41).

- **costs_bounds** (*tuple, optional*) – Configuration for c bounds of the form (lb, ub), in which both lb and ub are floats. This is only relevant if X_3 was formulated by `product_formulations` in *Problem*. By default, marginal costs are unbounded.

When `costs_type` in *Problem* is 'log', nonpositive $c(\hat{\theta})$ values can create problems when computing $\tilde{c}(\hat{\theta}) = \log c(\hat{\theta})$. One solution is to set lb to a small number. Rows in Jacobians associated with clipped marginal costs will be zero.

Both `None` and `numpy.nan` are converted to `-numpy.inf` in lb and to `numpy.inf` in ub.

- **W** (*array-like, optional*) – Starting values for the weighting matrix, W . By default, the 2SLS weighting matrix in (3.23) is used.

If there are any `micro_moments`, the initial weighting matrix will by default be block-diagonal with an identity matrix for the micro moment block. This micro moment block should usually be replaced by a matrix that better reflects micro moment covariances and the size of the micro dataset relative to N .

- **center_moments** (*bool, optional*) – Whether to center each column of the demand- and supply-side moments g before updating the weighting matrix W according to (3.24). By default, the moments are centered. This has no effect if `w_type` is 'unadjusted'.
- **W_type** (*str, optional*) – How to update the weighting matrix. This has no effect if `method` is 'ls'. Usually, `se_type` should be the same. The following types are supported:
 - 'robust' (default) - Heteroscedasticity robust weighting matrix defined in (3.24) and (3.25).
 - 'clustered' - Clustered weighting matrix defined in (3.24) and (3.26). Clusters must be defined by the `clustering_ids` field of `product_data` in *Problem*.
 - 'unadjusted' - Homoskedastic weighting matrix defined in (3.24) and (3.28).

This only affects the standard demand- and supply-side block of the updated weighting matrix. If there are micro moments, this matrix will be block-diagonal with a micro moment block equal to the inverse of the covariance matrix defined in (3.35) plus any `extra_micro_covariances`.

- **se_type** (*str, optional*) – How to compute parameter covariances and standard errors. Usually, `w_type` should be the same. The following types are supported:
 - 'robust' (default) - Heteroscedasticity robust covariances defined in (3.30) and (3.25).
 - 'clustered' - Clustered covariances defined in (3.30) and (3.26). Clusters must be defined by the `clustering_ids` field of `product_data` in *Problem*.
 - 'unadjusted' - Homoskedastic covariances defined in (3.31), which are computed under the assumption that the weighting matrix is optimal.

This only affects the standard demand- and supply-side block of the matrix of averaged moment covariances. If there are micro moments, the S matrix defined in the expressions referenced above will be block-diagonal with a micro moment block equal to the covariance matrix defined in (3.35) plus any `extra_micro_covariances`.

- **micro_moments** (*sequence of FirstChoiceCovarianceMoment, optional*) – Configurations for the M_M micro moments that will be added to the standard set of moments. The only type of micro moment currently supported is the *FirstChoiceCovarianceMoment*. By default, no micro moments are used, so $M_M = 0$.

If micro moments are specified, the micro moment block in W should usually be replaced by a matrix that better reflects micro moment covariances and the size of the micro dataset relative to N . If micro moments were computed with substantial sampling error, `extra_micro_covariances` can be specified to account for this additional source of error.

- **extra_micro_covariances** (*array-like, optional*) – Covariance matrix that is added on to the $M_M \times M_M$ matrix of micro moments covariances defined in (3.35), which is used to update the weighting matrix and compute standard errors. By default, this matrix is assumed to be zero. It should be specified if, for example, micro moments were computed with substantial sampling error.

Returns *ProblemResults* of the solved problem.

Return type *ProblemResults*

Examples

- *Tutorial*

5.4 Micro Moment Classes

Micro moment configurations can be passed to `Problem.solve()`. Only one type of micro moment is currently supported.

FirstChoiceCovarianceMoment(X2_index,
...[,...])

Configuration for micro moments that match covariances between product and agent characteristics, conditional on purchasing non-outside goods.

5.4.1 pyblp.FirstChoiceCovarianceMoment

class pyblp.**FirstChoiceCovarianceMoment** (*X2_index*, *demographics_index*, *value*, *market_ids=None*)

Configuration for micro moments that match covariances between product and agent characteristics, conditional on purchasing non-outside goods.

For example, survey data can often be used to compute the covariance σ_{xy} between a product characteristic, x_{jt} , and an agent demographic such as income, y_{it} , conditional on purchasing a non-outside good. With this covariance, a micro moment m in market t for agent i can be defined by

$$g_{M,mti} = (z_{it} - \bar{z}_t)(y_{it} - \bar{y}_t) - \sigma_{xy} \quad (5.6)$$

where $\bar{z}_t = \sum_i w_{it} z_{it}$, $\bar{y}_t = \sum_i w_{it} y_{it}$, and conditional on purchasing a non-outside good, the expected value of x_{jt} for agent i is

$$z_{it} = \sum_{j=1}^{J_t} x_{jt} s_{j(-0)ti} \quad (5.7)$$

where $s_{j(-0)ti}$ is the probability of choosing j when the outside option is removed from the choice set.

Integrals of these micro moments are approximated within and averaged across a set \mathcal{T}_m of markets in which the micro data used to compute σ_{xy} is relevant, which gives $\bar{g}_{M,m}$ in (3.34).

Parameters

- **x2_index** (*int*) – Column index of x_{jt} in the matrix of nonlinear product characteristics, X_2 . This should be between zero and $K_2 - 1$, inclusive.
- **demographics_index** (*int*) – Column index of the demographic y_{it} (which can be any demographic, not just income) in the matrix of agent demographics, d . This should be between zero and $D - 1$, inclusive.
- **value** (*float*) – Value of the covariance σ_{xy} estimated from micro data.
- **market_ids** (*array-like, optional*) – Distinct market IDs over which the micro moments will be averaged to get $\bar{g}_{M,m}$. These are also the only markets in which the moments will be computed. By default, the moments are computed for and averaged across all markets. That is, by default, it is assumed that the specified `value` is relevant for and on average the same for all markets.

Examples

- *Tutorial*

Methods

5.5 Problem Results Class

Solved problems return the following results class.

*ProblemResults*Results of a solved BLP problem.

5.5.1 pyblp.ProblemResults

class `pyblp.ProblemResults`

Results of a solved BLP problem.

Many results are class attributes. Other post-estimation outputs be computed by calling class methods.

Note: Methods in this class that compute one or more post-estimation output per market support `parallel()` processing. If multiprocessing is used, market-by-market computation of each post-estimation output will be distributed among the processes.

problem*Problem* that created these results.**Type** *Problem***last_results***ProblemResults* from the last GMM step.**Type** *ProblemResults***step**

GMM step that created these results.

Type *int***optimization_time**

Number of seconds it took the optimization routine to finish.

Type *float***cumulative_optimization_time**Sum of *ProblemResults.optimization_time* for this step and all prior steps.**Type** *float***total_time**Sum of *ProblemResults.optimization_time* and the number of seconds it took to set up the GMM step and compute results after optimization had finished.**Type** *float***cumulative_total_time**Sum of *ProblemResults.total_time* for this step and all prior steps.**Type** *float***converged**

Whether the optimization routine converged.

Type *bool***cumulative_converged**

Whether the optimization routine converged for this step and all prior steps.

Type *bool*

optimization_iterations

Number of major iterations completed by the optimization routine.

Type *int*

cumulative_optimization_iterations

Sum of *ProblemResults.optimization_iterations* for this step and all prior steps.

Type *int*

objective_evaluations

Number of GMM objective evaluations.

Type *int*

cumulative_objective_evaluations

Sum of *ProblemResults.objective_evaluations* for this step and all prior steps.

Type *int*

fp_converged

Flags for convergence of the iteration routine used to compute $\delta(\hat{\theta})$ in each market during each objective evaluation. Rows are in the same order as *Problem.unique_market_ids* and column indices correspond to objective evaluations.

Type *ndarray*

cumulative_fp_converged

Concatenation of *ProblemResults.fp_converged* for this step and all prior steps.

Type *ndarray*

fp_iterations

Number of major iterations completed by the iteration routine used to compute $\delta(\hat{\theta})$ in each market during each objective evaluation. Rows are in the same order as *Problem.unique_market_ids* and column indices correspond to objective evaluations.

Type *ndarray*

cumulative_fp_iterations

Concatenation of *ProblemResults.fp_iterations* for this step and all prior steps.

Type *ndarray*

contraction_evaluations

Number of times the contraction used to compute $\delta(\hat{\theta})$ was evaluated in each market during each objective evaluation. Rows are in the same order as *Problem.unique_market_ids* and column indices correspond to objective evaluations.

Type *ndarray*

cumulative_contraction_evaluations

Concatenation of *ProblemResults.contraction_evaluations* for this step and all prior steps.

Type *ndarray*

parameters

Stacked parameters in the following order: $\hat{\theta}$, concentrated out elements of $\hat{\beta}$, and concentrated out elements of $\hat{\gamma}$.

Type *ndarray*

parameter_covariances

Estimated covariance matrix of the stacked parameters, from which standard errors are extracted. Parameter covariances are not estimated during the first step of two-step GMM.

Type *ndarray*

theta

Estimated unfixed parameters, $\hat{\theta}$, in the following order: $\hat{\Sigma}$, $\hat{\Pi}$, $\hat{\rho}$, non-concentrated out elements from $\hat{\beta}$, and non-concentrated out elements from $\hat{\gamma}$.

Type *ndarray*

sigma

Estimated Cholesky root of the covariance matrix for unobserved taste heterogeneity, $\hat{\Sigma}$.

Type *ndarray*

pi

Estimated parameters that measures how agent tastes vary with demographics, $\hat{\Pi}$.

Type *ndarray*

rho

Estimated parameters that measure within nesting group correlations, $\hat{\rho}$.

Type *ndarray*

beta

Estimated demand-side linear parameters, $\hat{\beta}$.

Type *ndarray*

gamma

Estimated supply-side linear parameters, $\hat{\gamma}$.

Type *ndarray*

sigma_se

Estimated standard errors for $\hat{\Sigma}$, which are not estimated in the first step of two-step GMM.

Type *ndarray*

pi_se

Estimated standard errors for $\hat{\Pi}$, which are not estimated in the first step of two-step GMM.

Type *ndarray*

rho_se

Estimated standard errors for $\hat{\rho}$, which are not estimated in the first step of two-step GMM.

Type *ndarray*

beta_se

Estimated standard errors for $\hat{\beta}$, which are not estimated in the first step of two-step GMM.

Type *ndarray*

gamma_se

Estimated standard errors for $\hat{\gamma}$, which are not estimated in the first step of two-step GMM.

Type *ndarray*

sigma_bounds

Bounds for Σ that were used during optimization, which are of the form (lb, ub).

Type *tuple*

pi_bounds

Bounds for Π that were used during optimization, which are of the form (lb, ub).

Type *tuple*

rho_bounds

Bounds for ρ that were used during optimization, which are of the form (lb, ub).

Type *tuple*

beta_bounds

Bounds for β that were used during optimization, which are of the form (lb, ub).

Type *tuple*

gamma_bounds

Bounds for γ that were used during optimization, which are of the form (lb, ub).

Type *tuple*

delta

Estimated mean utility, $\delta(\hat{\theta})$.

Type *ndarray*

tilde_costs

Estimated transformed marginal costs, $\tilde{c}(\hat{\theta})$ from (3.9). If `costs_bounds` were specified in `Problem.solve()`, `c` may have been clipped.

Type *ndarray*

clipped_costs

Vector of booleans indicating whether the associated marginal costs were clipped. All elements will be `False` if `costs_bounds` in `Problem.solve()` was not specified.

Type *ndarray*

xi

Estimated unobserved demand-side product characteristics, $\xi(\hat{\theta})$, or equivalently, the demand-side structural error term. When there are demand-side fixed effects, this is $\Delta\xi(\hat{\theta})$ in (3.32). That is, fixed effects are not included.

Type *ndarray*

omega

Estimated unobserved supply-side product characteristics, $\omega(\hat{\theta})$, or equivalently, the supply-side structural error term. When there are supply-side fixed effects, this is $\Delta\omega(\hat{\theta})$ in (3.32). That is, fixed effects are not included.

Type *ndarray*

micro

Averaged micro moments, \bar{g}_M , in (3.34).

Type *ndarray*

objective

GMM objective value, $q(\hat{\theta})$, defined in (3.10). Note that in some of the BLP literature (and earlier versions of this package), this expression was previously scaled by N^2 .

Type *float*

xi_by_theta_jacobian

Estimated $\frac{\partial \xi}{\partial \theta} = \frac{\partial \delta}{\partial \theta}$, which is used to compute the gradient and standard errors.

Type *ndarray*

omega_by_theta_jacobian

Estimated $\frac{\partial \omega}{\partial \theta} = \frac{\partial \tilde{c}}{\partial \theta}$, which is used to compute the gradient and standard errors.

Type *ndarray*

micro_by_theta_jacobian

Estimated $\frac{\partial \hat{q}_M}{\partial \theta}$, which is used to compute the gradient and standard errors.

Type *ndarray*

gradient

Gradient of the GMM objective, $\nabla q(\hat{\theta})$, defined in (3.18). This is computed after the optimization routine finishes even if the routine was configured to not use analytic gradients.

Type *ndarray*

projected_gradient

Projected gradient of the GMM objective. When there are no parameter bounds, this will always be equal to *ProblemResults.gradient*. Otherwise, if an element in $\hat{\theta}$ is equal to its lower (upper) bound, the corresponding projected gradient value will be truncated at a maximum (minimum) of zero.

Type *ndarray*

projected_gradient_norm

Infinity norm of *ProblemResults.projected_gradient*.

Type *ndarray*

hessian

Estimated Hessian of the GMM objective. By default, this is computed with finite central differences after the optimization routine finishes.

Type *ndarray*

reduced_hessian

Reduced Hessian of the GMM objective. When there are no parameter bounds, this will always be equal to *ProblemResults.hessian*. Otherwise, if an element in $\hat{\theta}$ is equal to either its lower or upper bound, the corresponding row and column in the reduced Hessian will be all zeros.

Type *ndarray*

reduced_hessian_eigenvalues

Eigenvalues of *ProblemResults.reduced_hessian*.

Type *ndarray*

W

Weighting matrix, W , used to compute these results.

Type *ndarray*

updated_W

Weighting matrix updated according to (3.24).

Type *ndarray*

Examples

- [Tutorial](#)

Methods

| | |
|--|---|
| <code>bootstrap([draws, seed, iteration])</code> | Use a parametric bootstrap to create an empirical distribution of results. |
| <code>compute_aggregate_elasticities([factor, ...])</code> | Estimate aggregate elasticities of demand, \mathcal{E} , with respect to a variable, x . |
| <code>compute_approximate_prices([firm_ids, ...])</code> | Approximate equilibrium prices after firm or cost changes, p^* , under the assumption that shares and their price derivatives are unaffected by such changes. |
| <code>compute_consumer_surpluses([prices, market_id])</code> | Estimate population-normalized consumer surpluses, CS. |
| <code>compute_costs([market_id])</code> | Estimate marginal costs, c . |
| <code>compute_diversion_ratios([name, market_id])</code> | Estimate matrices of diversion ratios, \mathcal{D} , with respect to a variable, x . |
| <code>compute_elasticities([name, market_id])</code> | Estimate matrices of elasticities of demand, ε , with respect to a variable, x . |
| <code>compute_hhi([firm_ids, shares, market_id])</code> | Estimate Herfindahl-Hirschman Indices, HHI. |
| <code>compute_long_run_diversion_ratios([market_id])</code> | Estimate the matrices of long-run diversion ratios, \mathcal{D} . |
| <code>compute_markup([prices, costs, market_id])</code> | Estimate markups, \mathcal{M} . |
| <code>compute_optimal_instruments([method, draws, ...])</code> | Estimate feasible optimal or efficient instruments, Z_D^{Opt} and Z_S^{Opt} . |
| <code>compute_prices([firm_ids, ownership, costs, ...])</code> | Estimate equilibrium prices after firm or cost changes, p^* . |
| <code>compute_probabilities([market_id])</code> | Estimate matrices of choice probabilities. |
| <code>compute_profits([prices, shares, costs, ...])</code> | Estimate population-normalized gross expected profits, π . |
| <code>compute_shares([prices, market_id])</code> | Estimate shares evaluated at specified prices. |
| <code>extract_diagonal_means(matrices)</code> | Extract means of diagonals from stacked $J_t \times J_t$ matrices for each market t . |
| <code>extract_diagonals(matrices)</code> | Extract diagonals from stacked $J_t \times J_t$ matrices for each market t . |
| <code>run_distance_test(unrestricted)</code> | Test the validity of model restrictions with the distance test. |
| <code>run_hansen_test()</code> | Test the validity of overidentifying restrictions with the Hansen J test. |
| <code>run_lm_test()</code> | Test the validity of model restrictions with the Lagrange multiplier test. |
| <code>run_wald_test(restrictions, ...)</code> | Test the validity of model restrictions with the Wald test. |
| <code>to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |

The results can be converted into a dictionary.

| | |
|---|--|
| <code>ProblemResults.to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |
|---|--|

5.5.2 pyblp.ProblemResults.to_dict

`ProblemResults.to_dict` (*attributes*=(`'step'`, `'optimization_time'`, `'cumulative_optimization_time'`, `'total_time'`, `'cumulative_total_time'`, `'converged'`, `'cumulative_converged'`, `'optimization_iterations'`, `'cumulative_optimization_iterations'`, `'objective_evaluations'`, `'cumulative_objective_evaluations'`, `'fp_converged'`, `'cumulative_fp_converged'`, `'fp_iterations'`, `'cumulative_fp_iterations'`, `'contraction_evaluations'`, `'cumulative_contraction_evaluations'`, `'parameters'`, `'parameter_covariances'`, `'theta'`, `'sigma'`, `'pi'`, `'rho'`, `'beta'`, `'gamma'`, `'sigma_se'`, `'pi_se'`, `'rho_se'`, `'beta_se'`, `'gamma_se'`, `'sigma_bounds'`, `'pi_bounds'`, `'rho_bounds'`, `'beta_bounds'`, `'gamma_bounds'`, `'delta'`, `'tilde_costs'`, `'clipped_costs'`, `'xi'`, `'omega'`, `'micro'`, `'objective'`, `'xi_by_theta_jacobian'`, `'omega_by_theta_jacobian'`, `'micro_by_theta_jacobian'`, `'gradient'`, `'projected_gradient'`, `'projected_gradient_norm'`, `'hessian'`, `'reduced_hessian'`, `'reduced_hessian_eigenvalues'`, `'W'`, `'updated_W'`))

Convert these results into a dictionary that maps attribute names to values.

Once converted to a dictionary, these results can be saved to a file with `pickle.dump()`.

Parameters *attributes* (sequence of str, optional) – Name of attributes that will be added to the dictionary. By default, all `ProblemResults` attributes are added except for `ProblemResults.problem` and `ProblemResults.last_results`.

Returns Mapping from attribute names to values.

Return type *dict*

Examples

- [Tutorial](#)

The following methods test the validity of overidentifying and model restrictions.

| | |
|--|---|
| <code>ProblemResults.run_hansen_test()</code> | Test the validity of overidentifying restrictions with the Hansen J test. |
| <code>ProblemResults.run_distance_test(unrestricted)</code> | Test the validity of model restrictions with the distance test. |
| <code>ProblemResults.run_lm_test()</code> | Test the validity of model restrictions with the Lagrange multiplier test. |
| <code>ProblemResults.run_wald_test(restrictions, ...)</code> | Test the validity of model restrictions with the Wald test. |

5.5.3 pyblp.ProblemResults.run_hansen_test

`ProblemResults.run_hansen_test()`

Test the validity of overidentifying restrictions with the Hansen J test.

Following *Hansen (1982)*, the J statistic is

$$J = N\bar{g}(\hat{\theta})'W\bar{g}(\hat{\theta}) \quad (5.8)$$

where $\bar{g}(\hat{\theta})$ is defined in (3.11) and W is the optimal weighting matrix in (3.24).

Note: The statistic can equivalently be written as $J = Nq(\hat{\theta})$ where the GMM objective value is defined in (3.10).

When the overidentifying restrictions in this model are valid, the J statistic is asymptotically χ^2 with degrees of freedom equal to the number of overidentifying restrictions. This requires that there are more moments than parameters.

Warning: This test requires `ProblemResults.W` to be an optimal weighting matrix, so it should typically be run only after two-step GMM or after one-step GMM with a pre-specified optimal weighting matrix.

Returns The J statistic.

Return type *float*

Examples

- [Tutorial](#)

5.5.4 pyblp.ProblemResults.run_distance_test

`ProblemResults.run_distance_test` (*unrestricted*)

Test the validity of model restrictions with the distance test.

Following *Newey and West (1987)*, the distance or likelihood ratio-like statistic is

$$LR = J(\hat{\theta}^r) - J(\hat{\theta}^u) \tag{5.9}$$

where $J(\hat{\theta}^r)$ is the J statistic defined in (5.8) for this restricted model and $J(\hat{\theta}^u)$ is the J statistic for the unrestricted model.

Note: The statistic can equivalently be written as $LR = N[q(\hat{\theta}^r) - q(\hat{\theta}^u)]$ where the GMM objective value is defined in (3.10).

If the restrictions in this model are valid, the distance statistic is asymptotically χ^2 with degrees of freedom equal to the number of restrictions.

Warning: This test requires each model's `ProblemResults.W` to be the optimal weighting matrix, so it should typically be run only after two-step GMM or after one-step GMM with pre-specified optimal weighting matrices.

Parameters `unrestricted` (`ProblemResults`) – `ProblemResults` for the unrestricted model.

Returns The distance statistic.

Return type *float*

Examples

- [Tutorial](#)

5.5.5 pyblp.ProblemResults.run_lm_test

`ProblemResults.run_lm_test()`

Test the validity of model restrictions with the Lagrange multiplier test.

Following *Newey and West (1987)*, the Lagrange multiplier or score statistic is

$$LM = N\bar{g}(\hat{\theta})'W\bar{G}(\hat{\theta})V\bar{G}(\hat{\theta})'W\bar{g}(\hat{\theta}) \quad (5.10)$$

where $\bar{g}(\hat{\theta})$ is defined in (3.11), $\bar{G}(\hat{\theta})$ is defined in (3.19), W is the optimal weighting matrix in (3.24), and V is the covariance matrix of parameters in (3.30).

If the restrictions in this model are valid, the Lagrange multiplier statistic is asymptotically χ^2 with degrees of freedom equal to the number of restrictions.

Warning: This test requires `ProblemResults.W` to be an optimal weighting matrix, so it should typically be run only after two-step GMM or after one-step GMM with a pre-specified optimal weighting matrix.

Returns The Lagrange multiplier statistic.

Return type *float*

Examples

- *Tutorial*

5.5.6 pyblp.ProblemResults.run_wald_test

`ProblemResults.run_wald_test(restrictions, restrictions_jacobian)`

Test the validity of model restrictions with the Wald test.

Following *Newey and West (1987)*, the Wald statistic is

$$\text{Wald} = Nr(\hat{\theta})'[R(\hat{\theta})VR(\hat{\theta})']^{-1}r(\hat{\theta}) \quad (5.11)$$

where the restrictions are $r(\theta) = 0$ under the test's null hypothesis, their Jacobian is $R(\theta) = \frac{\partial r(\theta)}{\partial \theta}$, and V is the covariance matrix of parameters in (3.30).

If the restrictions are valid, the Wald statistic is asymptotically χ^2 with degrees of freedom equal to the number of restrictions.

Parameters

- **restrictions** (*array-like*) – Column vector of the model restrictions evaluated at the estimated parameters, $r(\hat{\theta})$.
- **restrictions_jacobian** (*array-like*) – Estimated Jacobian of the restrictions with respect to all parameters, $R(\hat{\theta})$. This matrix should have as many rows as `restrictions` and as many columns as `ProblemResults.parameter_covariances`.

Returns The Wald statistic.

Return type *float*

Examples

- *Tutorial*

In addition to class attributes, other post-estimation outputs can be estimated market-by-market with the following methods, which each return an array.

| | |
|--|---|
| <code>ProblemResults.compute_aggregate_elasticities([name, ...])</code> | Estimate aggregate elasticities of demand, \mathcal{E} , with respect to a variable, x . |
| <code>ProblemResults.compute_elasticities([name, ...])</code> | Estimate matrices of elasticities of demand, ε , with respect to a variable, x . |
| <code>ProblemResults.compute_diversion_ratios([name, ...])</code> | Estimate matrices of diversion ratios, \mathcal{D} , with respect to a variable, x . |
| <code>ProblemResults.compute_long_run_diversion_ratios([name, ...])</code> | Estimate matrices of long-run diversion ratios, \mathcal{D} . |
| <code>ProblemResults.compute_probabilities([market_id, ...])</code> | Estimate matrices of choice probabilities. |
| <code>ProblemResults.extract_diagonals(matrices)</code> | Extract diagonals from stacked $J_t \times J_t$ matrices for each market t . |
| <code>ProblemResults.extract_diagonal_means(matrices)</code> | Extract means of diagonals from stacked $J_t \times J_t$ matrices for each market t . |
| <code>ProblemResults.compute_costs([market_id])</code> | Estimate marginal costs, c . |
| <code>ProblemResults.compute_approximate_prices([firm_ids, ...])</code> | Approximate equilibrium prices after firm or cost changes, p^* , under the assumption that shares and their price derivatives are unaffected by such changes. |
| <code>ProblemResults.compute_prices([firm_ids, ...])</code> | Estimate equilibrium prices after firm or cost changes, p^* . |
| <code>ProblemResults.compute_shares([prices, ...])</code> | Estimate shares evaluated at specified prices. |
| <code>ProblemResults.compute_hhi([firm_ids, ...])</code> | Estimate Herfindahl-Hirschman Indices, HHI. |
| <code>ProblemResults.compute_markup([prices, ...])</code> | Estimate markups, \mathcal{M} . |
| <code>ProblemResults.compute_profits([prices, ...])</code> | Estimate population-normalized gross expected profits, π . |
| <code>ProblemResults.compute_consumer_surpluses([prices, ...])</code> | Estimate population-normalized consumer surpluses, CS. |

5.5.7 pyblp.ProblemResults.compute_aggregate_elasticities

`ProblemResults.compute_aggregate_elasticities` (*factor=0.1*, *name='prices'*, *market_id=None*)

Estimate aggregate elasticities of demand, \mathcal{E} , with respect to a variable, x .

In market t , the aggregate elasticity of demand is

$$\mathcal{E} = \sum_{j=1}^{J_t} \frac{s_{jt}(x + \Delta x) - s_{jt}}{\Delta}, \quad (5.12)$$

in which Δ is a scalar factor and $s_{jt}(x + \Delta x)$ is the share of product j in market t , evaluated at the scaled values of the variable.

Parameters

- **factor** (*float, optional*) – The scalar factor, Δ .
- **name** (*str, optional*) – Name of the variable, x . By default, $x = p$, prices.

- **market_id** (*object, optional*) – ID of the market in which to compute aggregate elasticities. By default, aggregate elasticities are computed in all markets and stacked.

Returns Estimates of aggregate elasticities of demand, \mathcal{E} . If `market_id` was not specified, rows are in the same order as `Problem.unique_market_ids`.

Return type `ndarray`

Examples

- [Tutorial](#)

5.5.8 pyblp.ProblemResults.compute_elasticities

`ProblemResults.compute_elasticities` (*name='prices', market_id=None*)

Estimate matrices of elasticities of demand, ε , with respect to a variable, x .

For each market, the value in row j and column k of ε is

$$\varepsilon_{jk} = \frac{x_k}{s_j} \frac{\partial s_j}{\partial x_k}. \quad (5.13)$$

Parameters

- **name** (*str, optional*) – Name of the variable, x . By default, $x = p$, prices.
- **market_id** (*object, optional*) – ID of the market in which to compute elasticities. By default, elasticities are computed in all markets and stacked.

Returns Estimated $J_t \times J_t$ matrices of elasticities of demand, ε . If `market_id` was not specified, matrices are estimated in each market t and stacked. Columns for a market are in the same order as products for the market. If a market has fewer products than others, extra columns will contain `numpy.nan`.

Return type `ndarray`

Examples

- [Tutorial](#)

5.5.9 pyblp.ProblemResults.compute_diversion_ratios

`ProblemResults.compute_diversion_ratios` (*name='prices', market_id=None*)

Estimate matrices of diversion ratios, \mathcal{D} , with respect to a variable, x .

For each market, the value in row j and column $k \neq j$ is

$$\mathcal{D}_{jk} = -\frac{\partial s_k}{\partial x_j} \bigg/ \frac{\partial s_j}{\partial x_j}. \quad (5.14)$$

Diversion ratios for the outside good are reported on diagonals:

$$\mathcal{D}_{jj} = -\frac{\partial s_0}{\partial x_j} \bigg/ \frac{\partial s_j}{\partial x_j}. \quad (5.15)$$

Unlike `ProblemResults.compute_long_run_diversion_ratios()`, this gives the marginal treatment effect (MTE) version of the diversion ratio. For more information, see [Conlon and Mortimer \(2018\)](#).

Parameters

- **name** (*str, optional*) – Name of the variable, x . By default, $x = p$, prices.
- **market_id** (*object, optional*) – ID of the market in which to compute diversion ratios. By default, diversion ratios are computed in all markets and stacked.

Returns Estimated $J_t \times J_t$ matrices of diversion ratios, \mathcal{D} . If `market_id` was not specified, matrices are estimated in each market t and stacked. Columns for a market are in the same order as products for the market. If a market has fewer products than others, extra columns will contain `numpy.nan`.

Return type *ndarray*

Examples

- *Tutorial*

5.5.10 pyblp.ProblemResults.compute_long_run_diversion_ratios

`ProblemResults.compute_long_run_diversion_ratios` (*market_id=None*)

Estimate matrices of long-run diversion ratios, $\bar{\mathcal{D}}$.

For each market, the value in row j and column $k \neq j$ is

$$\bar{\mathcal{D}}_{jk} = \frac{s_{k(-j)} - s_k}{s_j}, \quad (5.16)$$

in which $s_{k(-j)}$ is the share of product k computed with j removed from the choice set. Long-run diversion ratios for the outside good are reported on diagonals:

$$\bar{\mathcal{D}}_{jj} = \frac{s_{0(-j)} - s_0}{s_j}. \quad (5.17)$$

Unlike `ProblemResults.compute_diversion_ratios()`, this gives the average treatment effect (ATE) version of the diversion ratio. For more information, see [Conlon and Mortimer \(2018\)](#).

Parameters `market_id` (*object, optional*) – ID of the market in which to compute long-run diversion ratios. By default, long-run diversion ratios are computed in all markets and stacked.

Returns Estimated $J_t \times J_t$ matrices of long-run diversion ratios, $\bar{\mathcal{D}}$. If `market_id` was not specified, matrices are estimated in each market t and stacked. Columns for a market are in the same order as products for the market. If a market has fewer products than others, extra columns will contain `numpy.nan`.

Return type *ndarray*

Examples

- *Tutorial*

5.5.11 pyblp.ProblemResults.compute_probabilities

`ProblemResults.compute_probabilities` (*market_id=None*)

Estimate matrices of choice probabilities.

For each market, the value in row j and column i is given by (3.5) when there are random coefficients, and by (3.37) when there is additionally a nested structure. For the logit and nested logit models, choice probabilities are marketshares.

Parameters `market_id` (*object, optional*) – ID of the market in which to compute choice probabilities. By default, choice probabilities are computed in all markets and stacked.

Returns Estimated $J_t \times I_t$ matrices of choice probabilities. If `market_id` was not specified, matrices are estimated in each market t and stacked. Columns for a market are in the same order as agents for the market. If a market has fewer agents than others, extra columns will contain `numpy.nan`.

Return type `ndarray`

Examples

- *Tutorial*

5.5.12 `pyblp.ProblemResults.extract_diagonals`

`ProblemResults.extract_diagonals` (*matrices*)

Extract diagonals from stacked $J_t \times J_t$ matrices for each market t .

Parameters `matrices` (*array-like*) – Stacked matrices, such as estimates of ε , computed by `ProblemResults.compute_elasticities()`; \mathcal{D} , computed by `ProblemResults.compute_diversion_ratios()`; $\bar{\mathcal{D}}$, computed by `ProblemResults.compute_long_run_diversion_ratios()`; or s_{jti} computed by `ProblemResults.compute_probabilities()`.

Returns Stacked matrix diagonals. If the matrices are estimates of ε , a diagonal is a market's own elasticities of demand; if they are estimates of \mathcal{D} or $\bar{\mathcal{D}}$, a diagonal is a market's diversion ratios to the outside good.

Return type `ndarray`

Examples

- *Tutorial*

5.5.13 `pyblp.ProblemResults.extract_diagonal_means`

`ProblemResults.extract_diagonal_means` (*matrices*)

Extract means of diagonals from stacked $J_t \times J_t$ matrices for each market t .

Parameters `matrices` (*array-like*) – Stacked matrices, such as estimates of ε , computed by `ProblemResults.compute_elasticities()`; \mathcal{D} , computed by `ProblemResults.compute_diversion_ratios()`; $\bar{\mathcal{D}}$, computed by `ProblemResults.compute_long_run_diversion_ratios()`; or s_{jti} computed by `ProblemResults.compute_probabilities()`.

Returns Stacked diagonal means. If the matrices are estimates of ε , the mean of a diagonal is a market's mean own elasticity of demand; if they are estimates of \mathcal{D} or $\bar{\mathcal{D}}$, the mean of a diagonal is a market's mean diversion ratio to the outside good. Rows are in the same order as `Problem.unique_market_ids`.

Return type *ndarray*

Examples

- *Tutorial*

5.5.14 pyblp.ProblemResults.compute_costs

`ProblemResults.compute_costs` (*market_id=None*)

Estimate marginal costs, *c*.

Marginal costs are computed with the η -markup equation in (3.7):

$$c = p - \eta. \quad (5.18)$$

Parameters `market_id` (*object, optional*) – ID of the market in which to compute marginal costs.

By default, marginal costs are computed in all markets and stacked.

Returns Marginal costs, *c*.

Return type *ndarray*

Examples

- *Tutorial*

5.5.15 pyblp.ProblemResults.compute_approximate_prices

`ProblemResults.compute_approximate_prices` (*firm_ids=None, ownership=None, costs=None, market_id=None*)

Approximate equilibrium prices after firm or cost changes, p^* , under the assumption that shares and their price derivatives are unaffected by such changes.

This approximation is discussed in, for example, *Nevo (1997)*. Prices in each market are computed according to the η -markup equation in (3.7):

$$p^* = c^* + \eta^*, \quad (5.19)$$

in which the markup term is approximated with

$$\eta^* \approx - \left(O^* \odot \frac{\partial s}{\partial p} \right)^{-1} s \quad (5.20)$$

where O^* is the ownership matrix associated with firm changes.

Parameters

- **firm_ids** (*array-like, optional*) – Potentially changed firm IDs. By default, the unchanged `firm_ids` field of `product_data` in *Problem* will be used.
- **ownership** (*array-like, optional*) – Potentially changed ownership matrices. By default, standard ownership matrices based on `firm_ids` will be used unless the `ownership` field of `product_data` in *Problem* was specified.
- **costs** (*array-like, optional*) – Potentially changed marginal costs, c^* . By default, unchanged marginal costs are computed with `ProblemResults.compute_costs()`.

- **market_id** (*object, optional*) – ID of the market in which to compute approximate equilibrium prices. By default, approximate equilibrium prices are computed in all markets and stacked.

Returns Approximation of equilibrium prices after any firm or cost changes, p^* .

Return type *ndarray*

Examples

- *Tutorial*

5.5.16 pyblp.ProblemResults.compute_prices

`ProblemResults.compute_prices` (*firm_ids=None, ownership=None, costs=None, prices=None, iteration=None, market_id=None*)

Estimate equilibrium prices after firm or cost changes, p^* .

Note: To compute equilibrium prices (and shares) associated with a more complicated counterfactual, a *Simulation* for the counterfactual can be initialized with the estimated parameters, structural errors, and marginal costs from these results, and then solved with *Simulation.replace_endogenous()*. The returned *SimulationResults* gives more information about the contraction than this method, such as the number of contraction evaluations.

Prices are computed in each market by iterating over the ζ -markup contraction in (3.46):

$$p^* \leftarrow c^* + \zeta^*(p^*), \quad (5.21)$$

in which the markup term from (3.43) is

$$\zeta^*(p^*) = \Lambda^{-1}(p^*)[O^* \odot \Gamma(p^*)]'(p^* - c^*) - \Lambda^{-1}(p^*) \quad (5.22)$$

where O^* is the ownership matrix associated with firm changes.

Parameters

- **firm_ids** (*array-like, optional*) – Potentially changed firm IDs. By default, the unchanged `firm_ids` field of `product_data` in *Problem* will be used.
- **ownership** (*array-like, optional*) – Potentially changed ownership matrices. By default, standard ownership matrices based on `firm_ids` will be used unless the `ownership` field of `product_data` in *Problem* was specified.
- **costs** (*array-like*) – Potentially changed marginal costs, c^* . By default, unchanged marginal costs are computed with *ProblemResults.compute_costs()*.
- **prices** (*array-like, optional*) – Prices at which the fixed point iteration routine will start. By default, unchanged prices, p , are used as starting values. Other reasonable starting prices include the approximate equilibrium prices computed by *ProblemResults.compute_approximate_prices()*.
- **iteration** (*Iteration, optional*) – *Iteration* configuration for how to solve the fixed point problem in each market. By default, `Iteration('simple', {'atol': 1e-12})` is used. Analytic Jacobians are not supported for solving this system.
- **market_id** (*object, optional*) – ID of the market in which to compute equilibrium prices. By default, equilibrium prices are computed in all markets and stacked.

Returns Estimates of equilibrium prices after any firm or cost changes, p^* .

Return type *ndarray*

Examples

- *Tutorial*

5.5.17 pyblp.ProblemResults.compute_shares

`ProblemResults.compute_shares` (*prices=None, market_id=None*)

Estimate shares evaluated at specified prices.

Note: To compute equilibrium shares (and prices) associated with a more complicated counterfactual, a *Simulation* for the counterfactual can be initialized with the estimated parameters, structural errors, and marginal costs from these results, and then solved with *Simulation.replace_endogenous()*.

Parameters

- **prices** (*array-like*) – Prices at which to evaluate shares, such as equilibrium prices, p^* , computed by *ProblemResults.compute_prices()*. By default, unchanged prices are used.
- **market_id** (*object, optional*) – ID of the market in which to compute shares. By default, shares are computed in all markets and stacked.

Returns Estimates of shares evaluated at the specified prices.

Return type *ndarray*

Examples

- *Tutorial*

5.5.18 pyblp.ProblemResults.compute_hhi

`ProblemResults.compute_hhi` (*firm_ids=None, shares=None, market_id=None*)

Estimate Herfindahl-Hirschman Indices, HHI.

The index in market t is

$$\text{HHI} = 10,000 \times \sum_{f=1}^{F_t} \left(\sum_{j \in \mathcal{J}_{ft}} s_j \right)^2. \quad (5.23)$$

Parameters

- **firm_ids** (*array-like, optional*) – Firm IDs. By default, the unchanged `firm_ids` field of `product_data` in *Problem* will be used.
- **shares** (*array-like, optional*) – Shares, s , such as those computed by *ProblemResults.compute_shares()*. By default, unchanged shares are used.
- **market_id** (*object, optional*) – ID of the market in which to compute the index. By default, indices are computed in all markets and stacked.

Returns Estimated Herfindahl-Hirschman Indices, HHI. If `market_ids` was not specified, rows are in the same order as `Problem.unique_market_ids`.

Return type `ndarray`

Examples

- *Tutorial*

5.5.19 `pyblp.ProblemResults.compute_markup`s

`ProblemResults.compute_markup`s (`prices=None, costs=None, market_id=None`)

Estimate markups, \mathcal{M} .

The markup of product j in market t is

$$\mathcal{M}_{jt} = \frac{p_{jt} - c_{jt}}{p_{jt}}. \quad (5.24)$$

Parameters

- **prices** (*array-like, optional*) – Prices, p , such as equilibrium prices, p^* , computed by `ProblemResults.compute_prices()`. By default, unchanged prices are used.
- **costs** (*array-like*) – Marginal costs, c . By default, marginal costs are computed with `ProblemResults.compute_costs()`.
- **market_id** (*object, optional*) – ID of the market in which to compute markups. By default, markups are computed in all markets and stacked.

Returns Estimated markups, \mathcal{M} .

Return type `ndarray`

Examples

- *Tutorial*

5.5.20 `pyblp.ProblemResults.compute_profits`

`ProblemResults.compute_profits` (`prices=None, shares=None, costs=None, market_id=None`)

Estimate population-normalized gross expected profits, π .

The profit from product j in market t is

$$\pi_{jt} = (p_{jt} - c_{jt})s_{jt}. \quad (5.25)$$

Parameters

- **prices** (*array-like, optional*) – Prices, p , such as equilibrium prices, p^* , computed by `ProblemResults.compute_prices()`. By default, unchanged prices are used.
- **shares** (*array-like, optional*) – Shares, s , such as those computed by `ProblemResults.compute_shares()`. By default, unchanged shares are used.
- **costs** (*array-like*) – Marginal costs, c . By default, marginal costs are computed with `ProblemResults.compute_costs()`.

- **market_id** (*object, optional*) – ID of the market in which to compute profits. By default, profits are computed in all markets and stacked.

Returns Estimated population-normalized gross expected profits, π .

Return type *ndarray*

Examples

- *Tutorial*

5.5.21 pyblp.ProblemResults.compute_consumer_surpluses

`ProblemResults.compute_consumer_surpluses` (*prices=None, market_id=None*)

Estimate population-normalized consumer surpluses, CS.

Assuming away nonlinear income effects, the surplus in market t is

$$CS = \sum_{i=1}^{I_t} w_{it} CS_{it}, \quad (5.26)$$

in which the consumer surplus for individual i is

$$CS_{it} = \log \left(1 + \sum_{j=1}^{J_t} \exp V_{jti} \right) / \left(-\frac{\partial V_{1ti}}{\partial p_{1t}} \right), \quad (5.27)$$

or with nesting parameters,

$$CS_{it} = \log \left(1 + \sum_{h=1}^H \exp V_{hti} \right) / \left(-\frac{\partial V_{1ti}}{\partial p_{1t}} \right) \quad (5.28)$$

where V_{jti} is defined in (3.1) and V_{hti} is defined in (3.38).

Warning: $\frac{\partial V_{1ti}}{\partial p_{1t}}$ is the derivative of utility for the first product with respect to its price. The first product is chosen arbitrarily because this method assumes that there are no nonlinear income effects, which implies that this derivative is the same for all products. Computed consumer surpluses will likely be incorrect if prices are formulated in a nonlinear fashion like `log(prices)`.

Parameters

- **prices** (*array-like, optional*) – Prices at which utilities and price derivatives will be evaluated, such as equilibrium prices, p^* , computed by `ProblemResults.compute_prices()`. By default, unchanged prices are used.
- **market_id** (*object, optional*) – ID of the market in which to compute consumer surplus. By default, consumer surpluses are computed in all markets and stacked.

Returns Estimated population-normalized consumer surpluses, CS. If `market_ids` was not specified, rows are in the same order as `Problem.unique_market_ids`.

Return type *ndarray*

Examples

- *Tutorial*

A parametric bootstrap can be used, for example, to compute standard errors for the above post-estimation outputs. The following method returns a results class with all of the above methods, which returns a distribution of post-estimation outputs corresponding to different bootstrapped samples.

| | |
|---|--|
| <code>ProblemResults.bootstrap([draws, seed, ...])</code> | Use a parametric bootstrap to create an empirical distribution of results. |
|---|--|

5.5.22 pyblp.ProblemResults.bootstrap

`ProblemResults.bootstrap` (*draws=1000, seed=None, iteration=None*)

Use a parametric bootstrap to create an empirical distribution of results.

The constructed `BootstrappedResults` can be used just like `ProblemResults` to compute various post-estimation outputs for different markets. The only difference is that `BootstrappedResults` methods return arrays with an extra first dimension, along which bootstrapped results are stacked. These stacked results can be used to construct, for example, confidence intervals for post-estimation outputs.

For each bootstrap draw, parameters are drawn from the estimated multivariate normal distribution of all parameters defined by `ProblemResults.parameters` and `ProblemResults.parameter_covariances`. Any bounds configured in `Problem.solve()` will also bound parameter draws. Each parameter draw is used to compute the implied mean utility, δ , and shares, s . If a supply side was estimated, the implied marginal costs, c , and prices, p , are computed as well by iterating over the ζ -markup contraction in (3.46).

Note: By default, parametric bootstrapping may use a lot of memory. This is because all bootstrapped results (for all `draws`) are stored in memory at the same time. Memory usage can be reduced by calling this method in a loop with `draws = 1`. In each iteration of the loop, compute the desired post-estimation output with the proper method of the returned `BootstrappedResults` class and store these outputs.

Parameters

- **draws** (*int, optional*) – The number of draws that will be taken from the joint distribution of the parameters. The default value is 1000.
- **seed** (*int, optional*) – Passed to `numpy.random.mtrand.RandomState` to seed the random number generator before any draws are taken. By default, a seed is not passed to the random number generator.
- **iteration** (*Iteration, optional*) – `Iteration` configuration used to compute bootstrapped prices by iterating over the ζ -markup equation in (3.46). By default, if a supply side was estimated, this is `Iteration('simple', {'atol': 1e-12})`. Analytic Jacobians are not supported for solving this system. This configuration is not used if a supply side was not estimated.

Returns Computed `BootstrappedResults`.

Return type `BootstrappedResults`

Examples

- *Tutorial*

Optimal instruments, which also return a results class instead of an array, can be estimated with the following method.

```
ProblemResults.compute_optimal_instruments
```

Estimate feasible optimal or efficient instruments, Z_D^{Opt} and Z_S^{Opt} .

5.5.23 pyblp.ProblemResults.compute_optimal_instruments

`ProblemResults.compute_optimal_instruments` (*method='approximate', draws=1, seed=None, expected_prices=None, iteration=None*)

Estimate feasible optimal or efficient instruments, Z_D^{Opt} and Z_S^{Opt} .

Optimal instruments have been shown, for example, by *Reynaert and Verboven (2014)* and *Conlon and Gortmaker (2019)*, to reduce bias, improve efficiency, and enhance stability of BLP estimates.

Optimal instruments in the spirit of *Amemiya (1977)* or *Chamberlain (1987)* are defined by

$$\begin{bmatrix} Z_{D,jt}^{Opt} \\ Z_{S,jt}^{Opt} \end{bmatrix} = \Sigma_{\xi\omega}^{-1} E \left[\begin{array}{c} \frac{\partial \xi_{jt}}{\partial \theta} \\ \frac{\partial \omega_{jt}}{\partial \theta} \end{array} \middle| Z \right], \quad (5.29)$$

in which Z are all exogenous variables.

Feasible optimal instruments are estimated by evaluating this expression at an estimated $\hat{\theta}$. The expectation is taken by approximating an integral over the joint density of ξ and ω . For each error term realization, if not already estimated, equilibrium prices and shares are computed by iterating over the ζ -markup contraction in (3.46).

The expected Jacobians are estimated with the average over all computed Jacobian realizations. The 2×2 normalizing matrix $\Sigma_{\xi\omega}$ is estimated with the sample covariance matrix of the error terms.

Optimal instruments for linear parameters not included in θ are simple product characteristics, so they are not computed here but are rather included in the final set of instruments by `OptimalInstrumentResults.to_problem()`.

Note: When both a supply and demand side are estimated, there are usually collinear rows in (5.29) because of overlapping product characteristics in X_1 and X_3 . The expression can be corrected by multiplying it with a conformable matrix of ones and zeros that remove the collinearity problem. The question of which rows to exclude is addressed in `OptimalInstrumentResults.to_problem()`.

Parameters

- **method** (*str, optional*) – The method by which the integral over the joint density of ξ and ω is approximated. The following methods are supported:
 - 'approximate' (default) - Evaluate the Jacobians at the expected value of the error terms: zero (draws will be ignored).
 - 'normal' - Draw from the normal approximation to the joint distribution of the error terms and take the average over the computed Jacobians (draws determines the number of draws).

- 'empirical' - Draw with replacement from the empirical joint distribution of the error terms and take the average over the computed Jacobians (draws determines the number of draws).
- **draws** (*int, optional*) – The number of draws that will be taken from the joint distribution of the error terms. This is ignored if method is 'approximate'. Because the default method is 'approximate', the default number of draws is 1, even though it will be ignored. For 'normal' or empirical, larger numbers such as 100 or 1000 are recommended.
- **seed** (*int, optional*) – Passed to `numpy.random.mtrand.RandomState` to seed the random number generator before any draws are taken. By default, a seed is not passed to the random number generator.
- **expected_prices** (*array-like, optional*) – Vector of expected prices conditional on all exogenous variables, $E[p | Z]$. By default, if a supply side was estimated, `iteration` is used. If only a demand side was estimated, this is by default estimated with the fitted values from a reduced form regression of endogenous prices onto Z_D .
- **iteration** (*Iteration, optional*) – `Iteration` configuration used to estimate expected prices by iterating over the ζ -markup contraction in (3.46). By default, if a supply side was estimated, this is `Iteration('simple', {'atol': 1e-12})`. Analytic Jacobians are not supported for solving this system. This configuration is not used if `expected_prices` is specified.

Returns Computed *OptimalInstrumentResults*.

Return type *OptimalInstrumentResults*

Examples

- *Tutorial*

5.6 Bootstrapped Problem Results Class

Parametric bootstrap computation returns the following class.

BootstrappedResults

Bootstrapped results of a solved problem.

5.6.1 pyblp.BootstrappedResults

class `pyblp.BootstrappedResults`
 Bootstrapped results of a solved problem.

This class has same methods as *ProblemResults* that compute post-estimation outputs in one or more markets, but not other methods like *ProblemResults.compute_optimal_instruments()* that do not make sense in a bootstrapped dataset. The only other difference is that methods return arrays with an extra first dimension along which bootstrapped results are stacked (these stacked results can be used to construct, for example, confidence intervals for post-estimation outputs). Similarly, arrays of data (except for firm IDs and ownership matrices) passed as arguments to methods should have an extra first dimension of size *BootstrappedResults.draws*.

problem_results

ProblemResults that was used to compute these bootstrapped results.

Type *ProblemResults*

bootstrapped_sigma

Bootstrapped Cholesky decomposition of the covariance matrix for unobserved taste heterogeneity, Σ .

Type *ndarray*

bootstrapped_pi

Bootstrapped parameters that measures how agent tastes vary with demographics, Π .

Type *ndarray*

bootstrapped_rho

Bootstrapped parameters that measure within nesting group correlations, ρ .

Type *ndarray*

bootstrapped_beta

Bootstrapped demand-side linear parameters, β .

Type *ndarray*

bootstrapped_gamma

Bootstrapped supply-side linear parameters, γ .

Type *ndarray*

bootstrapped_prices

Bootstrapped prices, p . If a supply side was not estimated, these are unchanged prices. Otherwise, they are equilibrium prices implied by each draw.

Type *ndarray*

bootstrapped_shares

Bootstrapped marketshares, s , implied by each draw.

Type *ndarray*

bootstrapped_delta

Bootstrapped mean utility, δ , implied by each draw.

Type *ndarray*

computation_time

Number of seconds it took to compute the bootstrapped results.

Type *float*

draws

Number of bootstrap draws.

Type *int*

fp_converged

Flags for convergence of the iteration routine used to compute equilibrium prices in each market. Rows are in the same order as *Problem.unique_market_ids* and column indices correspond to draws.

Type *ndarray*

fp_iterations

Number of major iterations completed by the iteration routine used to compute equilibrium prices in each market for each draw. Rows are in the same order as *Problem.unique_market_ids* and column indices correspond to draws.

Type *ndarray*

contraction_evaluations

Number of times the contraction used to compute equilibrium prices was evaluated in each market for each draw. Rows are in the same order as `Problem.unique_market_ids` and column indices correspond to draws.

Type `ndarray`

Examples

- *Tutorial*

This class has all of the same methods as `ProblemResults`, except for `ProblemResults.bootstrap()` and `ProblemResults.compute_optimal_instruments()`. It can also be converted into a dictionary.

| | |
|--|--|
| <code>BootstrappedResults.to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |
|--|--|

5.6.2 pyblp.BootstrappedResults.to_dict

```
BootstrappedResults.to_dict(attributes=('bootstrapped_sigma', 'bootstrapped_pi', 'bootstrapped_rho', 'bootstrapped_beta', 'bootstrapped_gamma', 'bootstrapped_prices', 'bootstrapped_shares', 'bootstrapped_delta', 'computation_time', 'draws', 'fp_converged', 'fp_iterations', 'contraction_evaluations'))
```

Convert these results into a dictionary that maps attribute names to values.

Once converted to a dictionary, these results can be saved to a file with `pickle.dump()`.

Parameters `attributes` (*sequence of str, optional*) – Name of attributes that will be added to the dictionary. By default, all `BootstrappedResults` attributes are added except for `BootstrappedResults.problem_results`.

Returns Mapping from attribute names to values.

Return type `dict`

Examples

- *Tutorial*

5.7 Optimal Instrument Results Class

Optimal instrument computation returns the following results class.

| | |
|---------------------------------------|--|
| <code>OptimalInstrumentResults</code> | Results of optimal instrument computation. |
|---------------------------------------|--|

5.7.1 pyblp.OptimalInstrumentResults

```
class pyblp.OptimalInstrumentResults
    Results of optimal instrument computation.
```

The `OptimalInstrumentResults.to_problem()` method can be used to update the original `Problem` with the computed optimal instruments.

problem_results

`ProblemResults` that was used to compute these optimal instrument results.

Type `ProblemResults`

demand_instruments

Estimated optimal demand-side instruments for θ , Z_D^{Opt} .

Type `ndarray`

supply_instruments

Estimated optimal supply-side instruments for θ , Z_S^{Opt} .

Type `ndarray`

supply_shifter_formulation

`Formulation` configuration for supply shifters that will by default be included in the full set of optimal demand-side instruments. This is only constructed if a supply side was estimated, and it can be changed in `OptimalInstrumentResults.to_problem()`. By default, this is the formulation for X_3 from `Problem` excluding any variables in the formulation for X_1 .

Type `Formulation or None`

demand_shifter_formulation

`Formulation` configuration for demand shifters that will by default be included in the full set of optimal supply-side instruments. This is only constructed if a supply side was estimated, and it can be changed in `OptimalInstrumentResults.to_problem()`. By default, this is the formulation for X_1^x from `Problem` excluding any variables in the formulation for X_3 .

Type `Formulation or None`

inverse_covariance_matrix

Inverse of the sample covariance matrix of the estimated ξ and ω , which is used to normalize the expected Jacobians. If a supply side was not estimated, this is simply the sample estimate of $1/\sigma_\xi^2$.

Type `ndarray`

expected_xi_by_theta_jacobian

Estimated $E[\frac{\partial \xi}{\partial \theta} | Z]$.

Type `ndarray`

expected_omega_by_theta_jacobian

Estimated $E[\frac{\partial \omega}{\partial \theta} | Z]$.

Type `ndarray`

expected_prices

Vector of expected prices conditional on all exogenous variables, $E[p | Z]$, which may have been specified in `ProblemResults.compute_optimal_instruments()`.

Type `ndarray`

computation_time

Number of seconds it took to compute optimal excluded instruments.

Type `float`

draws

Number of draws used to approximate the integral over the error term density.

Type `int`

fp_converged

Flags for convergence of the iteration routine used to compute equilibrium prices in each market. Rows are in the same order as `Problem.unique_market_ids` and column indices correspond to draws.

Type `ndarray`

fp_iterations

Number of major iterations completed by the iteration routine used to compute equilibrium prices in each market for each error term draw. Rows are in the same order as `Problem.unique_market_ids` and column indices correspond to draws.

Type `ndarray`

contraction_evaluations

Number of times the contraction used to compute equilibrium prices was evaluated in each market for each error term draw. Rows are in the same order as `Problem.unique_market_ids` and column indices correspond to draws.

Type `ndarray`

Examples

- [Tutorial](#)

Methods

| | |
|--|--|
| <code>to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |
| <code>to_problem([supply_shifter_formulation, ...])</code> | Re-create the problem with estimated feasible optimal instruments. |

The results can be converted into a dictionary.

| | |
|---|--|
| <code>OptimalInstrumentResults.to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |
|---|--|

5.7.2 pyblp.OptimalInstrumentResults.to_dict

`OptimalInstrumentResults.to_dict` (`attributes=('demand_instruments', 'supply_instruments', 'inverse_covariance_matrix', 'expected_xi_by_theta_jacobian', 'expected_omega_by_theta_jacobian', 'expected_prices', 'computation_time', 'draws', 'fp_converged', 'fp_iterations', 'contraction_evaluations')`)

Convert these results into a dictionary that maps attribute names to values.

Once converted to a dictionary, these results can be saved to a file with `pickle.dump()`.

Parameters `attributes` (*sequence of str, optional*) – Name of attributes that will be added to the dictionary. By default, all `OptimalInstrumentResults` attributes are added except for `OptimalInstrumentResults.problem_results`, `OptimalInstrumentResults.supply_shifter_formulation`, and `OptimalInstrumentResults.demand_shifter_formulation`.

Returns Mapping from attribute names to values.

Return type *dict*

Examples

- *Tutorial*

They can also be converted into a *Problem* with the following method.

| | |
|---|--|
| <code>OptimalInstrumentResults.to_problem(...)</code> | Re-create the problem with estimated feasible optimal instruments. |
|---|--|

5.7.3 pyblp.OptimalInstrumentResults.to_problem

`OptimalInstrumentResults.to_problem` (*supply_shifter_formulation=None*, *de-*
mand_shifter_formulation=None)

Re-create the problem with estimated feasible optimal instruments.

The re-created problem will be exactly the same, except that instruments will be replaced with estimated feasible optimal instruments.

Note: Most of the explanation here is only important if a supply side was estimated.

The optimal excluded demand-side instruments consist of the following:

1. Estimated optimal demand-side instruments for θ , Z_D^{Opt} , excluding columns of instruments for any exogenous linear parameters that were not concentrated out, but rather included in θ by `Problem.solve()`.
2. Optimal instruments for any linear demand-side parameters on endogenous product characteristics, α , which were concentrated out and hence not included in θ . These optimal instruments are simply an integral of the endogenous product characteristics, X_1^p , over the joint density of ξ and ω . It is only possible to concentrate out α when there isn't a supply side, so the approximation of these optimal instruments is simply X_1^p evaluated at the constant vector of expected prices, $E[p | Z]$, specified in `ProblemResults.compute_optimal_instruments()`.
3. If a supply side was estimated, any supply shifters, which are by default formulated by `OptimalInstrumentResults.supply_shifter_formulation`: all characteristics in X_3 not in X_1 .

Similarly, if a supply side was estimated, the optimal excluded supply-side instruments consist of the following:

1. Estimated optimal supply-side instruments for θ , Z_S^{Opt} , excluding columns of instruments for any exogenous linear parameters that were not concentrated out, but rather included in θ by `Problem.solve()`.
2. If a supply side was estimated, any demand shifters, which are by default formulated by `OptimalInstrumentResults.demand_shifter_formulation`: all characteristics in X_1^x not in X_3 .

As usual, the excluded demand-side instruments will be supplemented with X_1^x and the excluded supply-side instruments will be supplemented with X_3 . The same fixed effects configured in `Problem` will be absorbed.

Warning: If a supply side was estimated, the addition of supply- and demand-shifters may create collinearity issues. Make sure to check that shifters and other product characteristics are not collinear.

Parameters

- **supply_shifter_formulation** (*Formulation, optional*) – *Formulation* configuration for supply shifters to be included in the set of optimal demand-side instruments. This is only used if a supply side was estimated. Intercepts will be ignored. By default, *OptimalInstrumentResults.supply_shifter_formulation* is used.
- **demand_shifter_formulation** (*Formulation, optional*) – *Formulation* configuration for demand shifters to be included in the set of optimal supply-side instruments. This is only used if a supply side was estimated. Intercepts will be ignored. By default, *OptimalInstrumentResults.demand_shifter_formulation* is used.

Returns *OptimalInstrumentProblem*, which is a *Problem* updated to use the estimated optimal instruments.

Return type *OptimalInstrumentProblem*

Examples

- *Tutorial*

This method returns the following class, which behaves exactly like a *Problem*.

| | |
|---------------------------------|--|
| <i>OptimalInstrumentProblem</i> | A BLP problem updated with optimal excluded instruments. |
|---------------------------------|--|

5.7.4 pyblp.OptimalInstrumentProblem

class `pyblp.OptimalInstrumentProblem`

A BLP problem updated with optimal excluded instruments.

This class can be used exactly like *Problem*.

5.8 Simulation Class

The following class allows for evaluation of more complicated counterfactuals than is possible with *ProblemResults* methods, or for simulation of synthetic data from scratch.

| | |
|---|--|
| <i>Simulation</i> (<code>product_formulations, ...[, ...]</code>) | Simulation of data in BLP-type models. |
|---|--|

5.8.1 pyblp.Simulation

class `pyblp.Simulation`(`product_formulations, product_data, beta, sigma=None, pi=None, gamma=None, rho=None, agent_formulation=None, agent_data=None, integration=None, xi=None, omega=None, xi_variance=1, omega_variance=1, correlation=0.9, costs_type='linear', seed=None`)

Simulation of data in BLP-type models.

Any data left unspecified are simulated during initialization. Simulated prices and shares can be replaced by *Simulation.replace_endogenous()* with equilibrium values that are consistent with true parameters. Less commonly, simulated exogenous variables can be replaced instead by *Simulation.replace_exogenous()*. Simulations are typically used for two purposes:

1. Solving for equilibrium prices and shares under more complicated counterfactuals than is possible with `ProblemResults.compute_prices()` and `ProblemResults.compute_shares()`. For example, this class can be initialized with estimated parameters, structural errors, and marginal costs from a `ProblemResults()`, but with changed data (fewer products, new products, different characteristics, etc.) and `Simulation.replace_endogenous()` can be used to compute the corresponding prices and shares.
2. Simulation of BLP-type models from scratch. For example, a model with fixed true parameters can be simulated many times, converted into problems with `SimulationResults.to_problem()`, and solved with `Problem.solve()` to evaluate in a Monte Carlo study how well the true parameters can be recovered.

If data for variables (used to formulate product characteristics in X_1 , X_2 , and X_3 , as well as agent demographics, d , and endogenous prices and marketshares p and s) are not provided, the values for each unspecified variable are drawn independently from the standard uniform distribution. In each market t , marketshares are divided by the number of products in the market J_t . Typically, `Simulation.replace_endogenous()` is used to replace prices and shares with equilibrium values that are consistent with true parameters.

If data for unobserved demand-and supply-side product characteristics, ξ and ω , are not provided, they are by default drawn from a mean-zero bivariate normal distribution.

After variables are loaded or simulated, any unspecified integration nodes and weights, ν and w , are constructed according to a specified `Integration` configuration.

Parameters

- **product_formulations** (*Formulation or sequence of Formulation*) – `Formulation` configuration or a sequence of up to three `Formulation` configurations for the matrix of linear product characteristics, X_1 , for the matrix of nonlinear product characteristics, X_2 , and for the matrix of cost characteristics, X_3 , respectively. If the formulation for X_2 is not specified or is `None`, the logit (or nested logit) model will be simulated.

The `shares` variable should not be included in any of the formulations and `prices` should be included in the formulation for X_1 or X_2 (or both). Variables that cannot be loaded from `product_data` will be drawn from independent standard uniform distributions. Unlike in `Problem`, fixed effect absorption is not supported during simulation.

Warning: Characteristics that involve prices, p , should always be formulated with the `prices` variable. If another name is used, `Simulation` will not understand that the characteristic is endogenous. For example, to include a p^2 characteristic, include `I(prices**2)` in a formula instead of manually including a `prices_squared` variable in `product_data` and a formula.

- **product_data** (*structured array-like*) – Each row corresponds to a product. Markets can have differing numbers of products. The convenience function `build_id_data()` can be used to construct the following required ID data:

- **market_ids** : (*object*) - IDs that associate products with markets.
- **firm_ids** : (*object*) - IDs that associate products with firms.

Custom ownership matrices can be specified as well:

- **ownership** : (*numeric, optional*) - *Custom stacked* :math: 'J_t \text{ times } J_t ownership matrices, O , for each market t , which can be built with `build_ownership()`. By default, standard ownership matrices are built only when they are needed to reduce memory usage. If specified, there should be as many columns as there are products in the market with the most products. Rightmost columns in markets with fewer products will be ignored.

Note: The `ownership` field can either be a matrix or can be broken up into multiple one-dimensional fields with column index suffixes that start at zero. For example, if there are three products in each market, a `ownership` field with three columns can be replaced by three one-dimensional fields: `ownership0`, `ownership1`, and `ownership2`.

To simulate a nested logit or random coefficients nested logit (RCNL) model, nesting groups must be specified:

- **nesting_ids** (*object, optional*) - IDs that associate products with nesting groups. When these IDs are specified, `rho` must be specified as well.

Along with `market_ids`, `firm_ids`, and `nesting_ids`, the names of any additional fields can typically be used as variables in `product_formulations`. However, there are a few variable names such as 'X1', which are reserved for use by *Products*.

- **beta** (*array-like*) – Vector of demand-side linear parameters, β . Elements correspond to columns in X_1 , which is formulated by `product_formulations`.
- **sigma** (*array-like, optional*) – Lower-triangular Cholesky root of the covariance matrix for unobserved taste heterogeneity, Σ . Rows and columns correspond to columns in X_2 , which is formulated by `product_formulations`. If X_2 is not formulated, this should not be specified, since the logit model will be simulated.
- **pi** (*array-like, optional*) – Parameters that measure how agent tastes vary with demographics, Π . Rows correspond to the same product characteristics as in `sigma`. Columns correspond to columns in d , which is formulated by `agent_formulation`. If d is not formulated, this should not be specified.
- **gamma** (*array-like, optional*) – Vector of supply-side linear parameters, γ . Elements correspond to columns in X_3 , which is formulated by `product_formulations`. If X_3 is not formulated, this should not be specified.
- **rho** (*array-like, optional*) – Parameters that measure within nesting group correlation, ρ . If this is a scalar, it corresponds to all groups defined by the `nesting_ids` field of `product_data`. If this is a vector, it must have H elements, one for each nesting group. Elements correspond to group IDs in the sorted order of `Simulation.unique_nesting_ids`. If nesting IDs are not specified, this should not be specified either.
- **agent_formulation** (*Formulation, optional*) – `Formulation` configuration for the matrix of observed agent characteristics called demographics, d , which will only be included in the model if this formulation is specified. Any variables that cannot be loaded from `agent_data` will be drawn from independent standard uniform distributions.
- **agent_data** (*structured array-like, optional*) – Each row corresponds to an agent. Markets can have differing numbers of agents. Since simulated agents are only used if there are nonlinear product characteristics, agent data should only be specified if X_2 is formulated in `product_formulations`. If agent data are specified, market IDs are required:
 - **market_ids** : (*object, optional*) - IDs that associate agents with markets. The set of distinct IDs should be the same as the set in `product_data`. If integration is specified, there must be at least as many rows in each market as the number of nodes and weights that are built for the market.

If `integration` is not specified, the following fields are required:

- **weights** : (*numeric, optional*) - Integration weights, w , for integration over agent choice probabilities.

- **nodes** : (*numeric, optional*) - Unobserved agent characteristics called integration nodes, ν . If there are more than K_2 columns (the number of nonlinear product characteristics), only the first K_2 will be used.

The convenience function `build_integration()` can be useful when constructing custom nodes and weights.

Note: If `nodes` has multiple columns, it can be specified as a matrix or broken up into multiple one-dimensional fields with column index suffixes that start at zero. For example, if there are three columns of nodes, a `nodes` field with three columns can be replaced by three one-dimensional fields: `nodes0`, `nodes1`, and `nodes2`.

Along with `market_ids`, the names of any additional fields can typically be used as variables in `agent_formulation`. The exception is the name 'demographics', which is reserved for use by *Agents*.

- **integration** (*Integration, optional*) – `Integration` configuration for how to build nodes and weights for integration over agent choice probabilities, which will replace any `nodes` and `weights` fields in `agent_data`. This configuration is required if `nodes` and `weights` in `agent_data` are not specified. It should not be specified if X_2 is not formulated in `product_formulations`.

If this configuration is specified, K_2 columns of nodes (the number of nonlinear product characteristics) will be built. However, if `sigma` is left unspecified or is specified with columns fixed at zero, fewer columns will be used.

- **xi** (*array-like, optional*) – Unobserved demand-side product characteristics, ξ . By default, if X_3 is formulated, each pair of unobserved characteristics in this vector and ω is drawn from a mean-zero bivariate normal distribution. This must be specified if X_3 is not formulated or if `omega` is specified.
- **omega** (*array-like, optional*) – Unobserved supply-side product characteristics, ω . By default, if X_3 is formulated, each pair of unobserved characteristics in this vector and ξ is drawn from a mean-zero bivariate normal distribution. This must be specified if X_3 is formulated and `xi` is specified. It is ignored if X_3 is not formulated.
- **xi_variance** (*float, optional*) – Variance of ξ . The default value is 1.0. This is ignored if `xi` or `omega` is specified.
- **omega_variance** (*float, optional*) – Variance of ω . The default value is 1.0. This is ignored if `xi` or `omega` is specified.
- **correlation** (*float, optional*) – Correlation between ξ and ω . The default value is 0.9. This is ignored if `xi` or `omega` is specified.
- **costs_type** (*str, optional*) – Specification of the marginal cost function $\tilde{c} = f(c)$ in (3.9). The following specifications are supported:
 - 'linear' (default) - Linear specification: $\tilde{c} = c$.
 - 'log' - Log-linear specification: $\tilde{c} = \log c$.
- **seed** (*int, optional*) – Passed to `numpy.random.mtrand.RandomState` to seed the random number generator before data are simulated. By default, a seed is not passed to the random number generator.

product_formulations

Formulation configurations for X_1 , X_2 , and X_3 , respectively.

Type *tuple*

agent_formulation

Formulation configuration for *d*.

Type *tuple*

product_data

Synthetic product data that were loaded or simulated during initialization. Typically, *Simulation.replace_endogenous()* is used to replace prices and shares with equilibrium values that are consistent with true parameters.

Type *recarray*

agent_data

Synthetic agent data that were loaded or simulated during initialization.

Type *recarray*

integration

Integration configuration for how any nodes and weights were built during initialization.

Type *Integration*

products

Product data structured as *Products*, which consists of data taken from *Simulation.product_data* along with matrices built according to *Simulation.product_formulations*.

Type *Products*

agents

Agent data structured as *Agents*, which consists of data taken from *Simulation.agent_data* or built by *Simulation.integration* along with any demographics formulated by *Simulation.agent_formulation*.

Type *Agents*

unique_market_ids

Unique market IDs in product and agent data.

Type *ndarray*

unique_firm_ids

Unique firm IDs in product data.

Type *ndarray*

unique_nesting_ids

Unique nesting IDs in product data.

Type *ndarray*

beta

Demand-side linear parameters, β .

Type *ndarray*

sigma

Cholesky root of the covariance matrix for unobserved taste heterogeneity, Σ .

Type *ndarray*

gamma

Supply-side linear parameters, γ .

Type *ndarray*

- pi**
Parameters that measures how agent tastes vary with demographics, Π .
Type *ndarray*
- rho**
Parameters that measure within nesting group correlation, ρ .
Type *ndarray*
- xi**
Unobserved demand-side product characteristics, ξ .
Type *ndarray*
- omega**
Unobserved supply-side product characteristics, ω .
Type *ndarray*
- costs_type**
Functional form of the marginal cost function $\tilde{c} = f(c)$.
Type *str*
- T**
Number of markets, T .
Type *int*
- N**
Number of products across all markets, N .
Type *int*
- F**
Number of firms across all markets, F .
Type *int*
- I**
Number of agents across all markets, I .
Type *int*
- K1**
Number of linear product characteristics, K_1 .
Type *int*
- K2**
Number of nonlinear product characteristics, K_2 .
Type *int*
- K3**
Number of cost product characteristics, K_3 .
Type *int*
- D**
Number of demographic variables, D .
Type *int*

MD

Number of demand-side instruments, M_D , which is always zero because instruments are added or constructed in `SimulationResults.to_problem()`.

Type *int*

MS

Number of supply-side instruments, M_S , which is always zero because instruments are added or constructed in `SimulationResults.to_problem()`.

Type *int*

ED

Number of absorbed dimensions of demand-side fixed effects, E_D , which is always zero because simulations do not support fixed effect absorption.

Type *int*

ES

Number of absorbed dimensions of supply-side fixed effects, E_S , which is always zero because simulations do not support fixed effect absorption.

Type *int*

H

Number of nesting groups, H .

Type *int*

Examples

- *Tutorial*

Methods

| | |
|--|---|
| <code>replace_endogenous([costs, prices, ...])</code> | Replace simulated prices and marketshares with equilibrium values that are consistent with true parameters. |
| <code>replace_exogenous(X1_name[, X3_name, delta, ...])</code> | Replace exogenous product characteristics with values that are consistent with true parameters. |

Once initialized, the following method replaces prices and shares with equilibrium values that are consistent with true parameters.

| | |
|--|---|
| <code>Simulation.replace_endogenous([costs, ...])</code> | Replace simulated prices and marketshares with equilibrium values that are consistent with true parameters. |
|--|---|

5.8.2 pyblp.Simulation.replace_endogenous

`Simulation.replace_endogenous` (*costs=None, prices=None, iteration=None, error_behavior='raise'*)

Replace simulated prices and marketshares with equilibrium values that are consistent with true parameters.

This method is the standard way of solving the simulation. Prices and marketshares are computed in each market

by iterating over the ζ -markup contraction in (3.46):

$$p \leftarrow c + \zeta(p). \tag{5.30}$$

Note: To create a simulation under perfect (instead of Bertrand) competition, use an *Iteration* configuration with `method='return'`. This will set prices equal to the default starting values for the iteration routine, which are marginal costs.

Note: This method supports *parallel()* processing. If multiprocessing is used, market-by-market computation of prices and shares will be distributed among the processes.

Parameters

- **costs** (*array-like, optional*) – Marginal costs, c . By default, $c = X_3\gamma + \omega$. Marginal costs must be specified if X_3 was not formulated in *Simulation*.
- **prices** (*array-like, optional*) – Prices at which the fixed point iteration routine will start. By default, `costs`, are used as starting values.
- **iteration** (*Iteration, optional*) – *Iteration* configuration for how to solve the fixed point problem. By default, `Iteration('simple', {'atol': 1e-12})` is used. Analytic Jacobians are not supported for solving this system.
- **error_behavior** (*str, optional*) – How to handle errors when computing prices and shares. For example, the fixed point routine may not converge if the effects of nonlinear parameters on price overwhelm the linear parameter on price, which should be sufficiently negative. The following behaviors are supported:
 - 'raise' (default) - Raise an exception.
 - 'warn' - Use the last computed prices and shares. If the fixed point routine fails to converge, these are the last prices and shares computed by the routine. If there are other issues, these are the starting prices and their associated shares.

Returns *SimulationResults* of the solved simulation.

Return type *SimulationResults*

Examples

- *Tutorial*

A less common way to solve the simulation is to assume simulated prices and shares represent an equilibrium and to replace exogenous variables instead.

| | |
|--|---|
| <code><i>Simulation.replace_exogenous</i>(X1_name[, ...])</code> | Replace exogenous product characteristics with values that are consistent with true parameters. |
|--|---|

5.8.3 pyblp.Simulation.replace_exogenous

`Simulation.replace_exogenous(X1_name, X3_name=None, delta=None, iteration=None, fp_type='safe_linear', error_behavior='raise')`

Replace exogenous product characteristics with values that are consistent with true parameters.

This method implements a less common way of solving the simulation. It may be preferable to `Simulation.replace_endogenous()` when for some reason it is desirable to retain the prices and marketshares from `Simulation`, which are assumed to be in equilibrium.

For this method of solving the simulation to be used, there must be an exogenous product characteristic v that shows up only in X_1 , and if there is a supply side, another product characteristic w that shows up only in X_3 . These characteristics will be replaced with values that are consistent with true parameters.

First, the mean utility δ is computed in each market by iterating over the contraction in (3.13) and $(\delta - \xi - X_1\beta)\beta_v^{-1}$ is added to the v from `Simulation`. Here, β_v is the linear parameter in β on v .

With a supply side, the marginal cost function \tilde{c} is computed according to (3.7) and (3.9) and $(\tilde{c} - \omega - X_3\gamma)\gamma_w^{-1}$ is added to the w from `Simulation`. Here, γ_w is the linear parameter in γ on w .

Note: This method supports `parallel()` processing. If multiprocessing is used, market-by-market computation of prices and shares will be distributed among the processes.

Parameters

- **X1_name** (*str*) – The name of the variable v in X_1 that will be replaced. It should show up only once in the formulation for X_1 from `Simulation`, and it should not be transformed in any way.
- **X3_name** (*str, optional*) – The name of the variable w in X_3 that will be replaced. It should show up only once in the formulation for X_3 from `Simulation`, and it should not be transformed in any way. This will only be used if there is a supply side.
- **delta** (*array-like, optional*) – Initial values for the mean utility, δ , which the fixed point iteration routine will start at. By default, the solution to the logit model in (3.40) is used. If there is a nesting structure, solution to the nested logit model in (3.41) under the initial `rho` is used instead.
- **iteration** (*Iteration, optional*) – `Iteration` configuration for how to solve the fixed point problem used to compute δ in each market. This configuration is only relevant if there are nonlinear parameters, since δ can be estimated analytically in the logit model. By default, `Iteration('squarem', {'atol': 1e-14})` is used. For more information, refer to the same argument in `Problem.solve()`.
- **fp_type** (*str, optional*) – Configuration for the type of contraction mapping used to compute δ . For information about the different types, refer to the same argument in `Problem.solve()`.
- **error_behavior** (*str, optional*) – How to handle errors when computing δ and \tilde{c} . The following behaviors are supported:
 - 'raise' (default) - Raise an exception.
 - 'warn' - Use the last computed δ and \tilde{c} . If the fixed point routine fails to converge, these are the last δ and the associated \tilde{c} by the routine. If there are other issues, these are the starting δ values and their associated \tilde{c} .

Returns `SimulationResults` of the solved simulation.

Return type `SimulationResults`

Examples

- *Tutorial*

5.9 Simulation Results Class

Solved simulations return the following results class.

*SimulationResults*Results of a solved simulation of synthetic BLP data.

5.9.1 pyblp.SimulationResults

class pyblp.SimulationResults

Results of a solved simulation of synthetic BLP data.

The *SimulationResults.to_problem()* method can be used to convert the full set of simulated data (along with some basic default instruments) and configured information into a *Problem*.

simulation

Simulation that created these results.

Type *Simulation*

product_data

Simulated *Simulation.product_data* with product characteristics replaced so as to be consistent with the true parameters. If *Simulation.replace_endogenous()* was used to create these results, prices and marketshares were replaced. If *Simulation.replace_exogenous()* was used, exogenous characteristics were replaced instead.

Type *recarray*

computation_time

Number of seconds it took to compute prices and marketshares.

Type *float*

fp_converged

Flags for convergence of the iteration routine used to compute prices or δ (depending on the method used to create these results) in each market. Flags are in the same order as *Simulation.unique_market_ids*.

Type *ndarray*

fp_iterations

Number of major iterations completed by the iteration routine used to compute prices or δ in each market. Counts are in the same order as *Simulation.unique_market_ids*.

Type *ndarray*

contraction_evaluations

Number of times the contraction used to compute prices or δ was evaluated in each market. Counts are in the same order as *Simulation.unique_market_ids*.

Type *ndarray*

Examples

- [Tutorial](#)

Methods

| | |
|--|--|
| <code>compute_micro(micro_moments)</code> | Compute averaged micro moment values, \bar{g}_M . |
| <code>to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |
| <code>to_problem([product_formulations, ...])</code> | Convert the solved simulation into a problem. |

The results can be converted into a dictionary.

| | |
|--|--|
| <code>SimulationResults.to_dict([attributes])</code> | Convert these results into a dictionary that maps attribute names to values. |
|--|--|

5.9.2 pyblp.SimulationResults.to_dict

`SimulationResults.to_dict(attributes=('product_data', 'computation_time', 'fp_converged', 'fp_iterations', 'contraction_evaluations'))`

Convert these results into a dictionary that maps attribute names to values.

Once converted to a dictionary, these results can be saved to a file with `pickle.dump()`.

Parameters `attributes` (*sequence of str, optional*) – Name of attributes that will be added to the dictionary. By default, all `SimulationResults` attributes are added except for `SimulationResults.simulation`.

Returns Mapping from attribute names to values.

Return type `dict`

Examples

- [Tutorial](#)

They can also be converted into a `Problem` with the following method.

| | |
|--|---|
| <code>SimulationResults.to_problem([...])</code> | Convert the solved simulation into a problem. |
|--|---|

5.9.3 pyblp.SimulationResults.to_problem

`SimulationResults.to_problem(product_formulations=None, product_data=None, agent_formulation=None, agent_data=None, integration=None, costs_type=None)`

Convert the solved simulation into a problem.

Arguments are the same as those of `Problem`. By default, the structure of the problem will be the same as that of the solved simulation.

By default, some simple “sums of characteristics” BLP instruments are constructed. Demand-side instruments are constructed by `build_blp_instruments()` from variables in X_1^x , along with any supply shifters

(variables in X_3 but not X_1). Supply side instruments are constructed from variables in X_3 , along with any demand shifters (variables in X_1 but not X_3). Instruments will also be constructed from columns of ones if there is variation in J_t , the number of products per market. Any constant columns will be dropped. For example, if each firm owns exactly one product in each market, the “rival” columns of instruments will be zero and hence dropped.

Note: These excluded instruments are constructed only for convenience. Especially for more complicated problems, they should be replaced with better instruments.

Parameters

- **product_formulations** (*Formulation or sequence of Formulation, optional*) – By default, `Simulation.product_formulations`.
- **product_data** (*structured array-like, optional*) – By default, `SimulationResults.product_data` with excluded instruments.
- **agent_formulation** (*Formulation, optional*) – By default, `Simulation.agent_formulation`.
- **agent_data** (*structured array-like, optional*) – By default, `Simulation.agent_data`.
- **integration** (*Integration, optional*) – By default, this is unspecified.
- **costs_type** (*str, optional*) – By default, `Simulation.costs_type`.

Returns A BLP problem.

Return type *Problem*

Examples

- *Tutorial*

Simulation results can also be used to compute micro moment values.

| | |
|---|---|
| <code>SimulationResults.compute_micro(micro_moments)</code> | Compute averaged micro moment values, \bar{g}_M . |
|---|---|

5.9.4 pyblp.SimulationResults.compute_micro

`SimulationResults.compute_micro(micro_moments)`

Compute averaged micro moment values, \bar{g}_M .

Typically, this method is used to compute the values that micro moments aim to match. This can be done by setting `value=0` in each of the configured `micro_moments`.

Parameters `micro_moments` (*sequence of FirstChoiceCovarianceMoment*) – Configurations for the averaged micro moments that will be computed. The only type of micro moment currently supported is the `FirstChoiceCovarianceMoment`.

Returns Averaged micro moments, \bar{g}_M , in (3.34).

Return type *ndarray*

Examples

- *Tutorial*

5.10 Structured Data Classes

Product and agent data that are passed or constructed by *Problem* and *Simulation* are structured internally into classes with field names that more closely resemble BLP notation. Although these structured data classes are not directly constructable, they can be accessed with *Problem* and *Simulation* class attributes. It can be helpful to compare these structured data classes with the data or configurations used to create them.

| | |
|-----------------|--|
| <i>Products</i> | Product data structured as a record array. |
| <i>Agents</i> | Agent data structured as a record array. |

5.10.1 pyblp.Products

class `pyblp.Products`

Product data structured as a record array.

Attributes in addition to the ones below are the variables underlying X_1 , X_2 , and X_3 .

market_ids

IDs that associate products with markets.

Type *ndarray*

firm_ids

IDs that associate products with firms.

Type *ndarray*

demand_ids

IDs used to create demand-side fixed effects.

Type *ndarray*

supply_ids

IDs used to create supply-side fixed effects.

Type *ndarray*

nesting_ids

IDs that associate products with nesting groups.

Type *ndarray*

clustering_ids

IDs used to compute clustered standard errors.

Type *ndarray*

ownership

Stacked $J_t \times J_t$ ownership matrices, O , for each market t .

Type *ndarray*

shares

Marketshares, s .

Type *ndarray*

prices

Product prices, p .

Type *ndarray*

ZD

Full set of demand-side instruments, Z_D : excluded demand-side instruments and X_1^x .

Type *ndarray*

ZS

Full set of supply-side instruments, Z_S : excluded supply-side instruments and X_3 .

Type *ndarray*

X1

Linear product characteristics, X_1 .

Type *ndarray*

X2

Nonlinear product characteristics, X_2 .

Type *ndarray*

X3

Cost product characteristics, X_3 .

Type *ndarray*

5.10.2 pyblp.Agents

class `pyblp.Agents`

Agent data structured as a record array.

market_ids

IDs that associate agents with markets.

Type *ndarray*

weights

Integration weights, w .

Type *ndarray*

nodes

Unobserved agent characteristics called integration nodes, ν .

Type *ndarray*

demographics

Observed agent characteristics, d .

Type *ndarray*

5.11 Multiprocessing

A context manager can be used to enable parallel processing for methods that perform market-by-market computation.

| | |
|----------------------------------|--|
| <code>parallel(processes)</code> | Context manager used for parallel processing in a <code>with</code> statement context. |
|----------------------------------|--|

5.11.1 pyblp.parallel

`pyblp.parallel` (*processes*)

Context manager used for parallel processing in a `with` statement context.

This manager creates a context in which a pool of Python processes will be used by any method that requires market-by-market computation. These methods will distribute their work among the processes. After the context created by the `with` statement ends, all worker processes in the pool will be terminated. Outside of this context, such methods will not use multiprocessing.

Importantly, multiprocessing will only improve speed if gains from parallelization outweigh overhead from serializing and passing data between processes. For example, if computation for a single market is very fast and there is a lot of data in each market that must be serialized and passed between processes, using multiprocessing may reduce overall speed.

Parameters `processes` (*int*) – Number of Python processes that will be created and used by any method that supports parallel processing.

Examples

The [online version](#) of the following section may be easier to read.

Parallel Processing Example

```
import pyblp
import pandas as pd

pyblp.options.digits = 2
pyblp.options.verbose = False
pyblp.__version__

'0.8.1'
```

In this example, we'll use parallel processing to compute elasticities market-by-market for a simple Logit problem configured with some of the fake cereal data from *Nevo (2000)*.

```
product_data = pd.read_csv(pyblp.data.NEVO_PRODUCTS_LOCATION)
formulation = pyblp.Formulation('0 + prices', absorb='C(product_ids)')
problem = pyblp.Problem(formulation, product_data)
results = problem.solve()
results
```

Problem Results Summary:

```
=====
GMM   Objective   Weighting Matrix
Step   Value       Condition Number
-----
  2    +8.3E-02      +5.7E+07
=====
```

Cumulative Statistics:

```
=====
Computation   Objective
  Time       Evaluations
-----
00:00:00         2
=====
```

Beta Estimates (Robust SEs in Parentheses):

(continues on next page)

```
=====  
prices  
-----  
-3.0E+01  
(+1.0E+00)  
=====
```

```
pyblp.options.verbose = True  
with pyblp.parallel(2):  
    elasticities = results.compute_elasticities()
```

```
Starting a pool of 2 processes ...  
Started the process pool after 00:00:00.  
Computing elasticities with respect to prices ...  
Finished after 00:00:04.
```

```
Terminating the pool of 2 processes ...  
Terminated the process pool after 00:00:00.
```

Solving a Logit problem does not require market-by-market computation, so parallelization does not change its estimation procedure. Although elasticity computation does happen market-by-market, this problem is very small, so in this small example there are no gains from parallelization.

If the problem were much larger, running *Problem.solve* and *ProblemResults.compute_elasticities* under the `with` statement could substantially speed up estimation and elasticity computation.

5.12 Options and Example Data

In addition to classes and functions, there are also two modules that can be used to configure global package options and locate example data that comes with the package.

| | |
|----------------------|---|
| <code>options</code> | Global options. |
| <code>data</code> | Locations of example data that are included in the package for convenience. |

5.12.1 `pyblp.options`

Global options.

`pyblp.options.digits`

Number of digits displayed by status updates. The default number of digits is 7. The number of digits can be changed to, for example, 2, with `pyblp.options.digits = 2`.

Type *int*

`pyblp.options.verbose`

Whether to output status updates. By default, verbosity is turned on. Verbosity can be turned off with `pyblp.options.verbose = False`.

Type *bool*

`pyblp.options.verbose_tracebacks`

Whether to include full tracebacks in error messages. By default, full tracebacks are turned off. These can be useful when attempting to find the source of an error message. Tracebacks can be turned on with `pyblp.options.verbose_tracebacks = True`.

Type *bool*

`pyblp.options.verbose_output`

Function used to output status updates. The default function is simply `print`. The function can be changed, for example, to include an indicator that statuses are from this package, with `pyblp.verbose_output = lambda x: print(f"pyblp: {x}")`.

Type *callable*

`pyblp.options.dtype`

The data type used for internal calculations, which is by default `numpy.float64`. The other recommended option is `numpy.longdouble`, which is the only extended precision floating point type currently supported by NumPy. Although this data type will be used internally, `numpy.float64` will be used when passing arrays to optimization and fixed point routines, which may not support extended precision. The library underlying `scipy.linalg`, which is used for matrix inversion, may also use `numpy.float64`.

One instance in which extended precision can be helpful in the BLP problem is when there are a large number of near zero choice probabilities with small integration weights, which, under standard precision are called zeros when in aggregate they are nonzero. For example, *Skrainka (2012)* finds that using long doubles is sufficient to solve many utility floating point problems.

The precision of `numpy.longdouble` depends on the platform on which NumPy is installed. If the platform in use does not support extended precision, using `numpy.longdouble` may lead to unreliably results. For example, on Windows, NumPy is usually compiled such that `numpy.longdouble` often behaves like `numpy.float64`. Precisions can be compared with `numpy.finfo` by running `np.finfo(np.float64)` and `np.finfo(np.longdouble)`. For more information, refer to [this discussion](#).

If extended precisions is supported, the data type can be switched with `pyblp.options.dtype = np.longdouble`. On Windows, it is often easier to install Linux in a virtual machine than it is to build NumPy from source with a non-standard compiler.

Type *dtype*

`pyblp.options.pseudo_inverses`

Whether to compute Moore-Penrose pseudo-inverses of matrices with `scipy.linalg.pinv()` instead of their classic inverses with `scipy.linalg.inv()`. This is by default `True`, so pseudo-inverses will be used. Up to small numerical differences, the pseudo-inverse is identical to the classic inverse for invertible matrices. Using the pseudo-inverse by default can help alleviate problems from, for example, near-singular weighting matrices.

To always attempt to compute classic inverses first, set `pyblp.options.pseudo_inverses = False`. If a classic inverse cannot be computed, an error will be displayed, and a pseudo-inverse may be computed instead.

Type *bool*

`pyblp.options.collinear_atol`

Absolute tolerance for detecting collinear columns in each matrix of product characteristics and instruments: X_1 , X_2 , X_3 , Z_D , and Z_S .

Each matrix is decomposed into a QR decomposition and an error is raised for any column whose diagonal element in R has a magnitude less than `collinear_atol + collinear_rtol * sd` where `sd` is the column's standard deviation.

The default absolute tolerance is $1e-14$. To disable collinearity checks, set `pyblp.options.collinear_atol = pyblp.options.collinear_rtol = 0`.

Type *float*

`pyblp.options.collinear_rtol`

Relative tolerance for detecting collinear columns, which is by default also $1e-14$.

Type *float*

`pyblp.options.psd_atol`

Absolute tolerance for detecting non-positive semidefinite matrices. For example, this check is applied to any custom weighting matrix, W .

Singular value decomposition factorizes the matrix into $U\Sigma V$ and an error is raised if any element in the original matrix differs in absolute value from $V'\Sigma V$ by more than `psd_atol + psd_rtol * abs` where `abs` is the element's absolute value.

The default tolerance is $1e-8$. To disable positive semidefinite checks, set `pyblp.options.psd_atol = pyblp.options.psd_rtol = np.inf`.

Type *float*

`pyblp.options.psd_rtol`

Relative tolerance for detecting non-positive definite matrices, which is by default also $1e-8$.

Type *float*

5.12.2 pyblp.data

Locations of example data that are included in the package for convenience.

`pyblp.data.NEVO_PRODUCTS_LOCATION`

Location of a CSV file containing the fake cereal product data from *Nevo (2000)*. The file includes the same pre-computed excluded instruments used in the original paper.

Type *str*

`pyblp.data.NEVO_AGENTS_LOCATION`

Location of a CSV file containing the fake cereal agent data. Included in the file are Monte Carlo weights and draws along with demographics, which are used by *Nevo (2000)* to solve the fake cereal problem.

Type *str*

`pyblp.data.BLP_PRODUCTS_LOCATION`

Location of a CSV file containing the automobile product data extracted by *Andrews, Gentzkow, and Shapiro (2017)* from the original GAUSS code for *Berry, Levinsohn, and Pakes (1999)*, which is commonly assumed to be the same data used in *Berry, Levinsohn, and Pakes (1995)*.

The file also includes a set of excluded instruments. First, “sums of characteristics” BLP instruments from the original paper were computed with `build_blp_instruments()`. As in the original paper, the `mpd` variable was added to the set of excluded instruments for supply. Due to a collinearity problem with these original instruments, each set of excluded instruments was then interacted up to the second degree, standardized, replaced with the minimum set of principal components that explained at least 99% of the variance, and standardized again.

Type *str*

`pyblp.data.BLP_AGENTS_LOCATION`

Location of a CSV file containing automobile agent data. Included in the file are 200 Monte Carlo weights and draws for each market, which, unlike in the fake cereal data, are not the same draws used in the original paper.

Also included is an income demographic, which consists of draws from lognormal distributions. The log of income has a common standard deviation, 1.72, and the following market-varying means:

| Year | Mean |
|------|---------|
| 1971 | 2.01156 |
| 1972 | 2.06526 |
| 1973 | 2.07843 |
| 1974 | 2.05775 |
| 1975 | 2.02915 |
| 1976 | 2.05346 |
| 1977 | 2.06745 |
| 1978 | 2.09805 |
| 1979 | 2.10404 |
| 1980 | 2.07208 |
| 1981 | 2.06019 |
| 1982 | 2.06561 |
| 1983 | 2.07672 |
| 1984 | 2.10437 |
| 1985 | 2.12608 |
| 1986 | 2.16426 |
| 1987 | 2.18071 |
| 1988 | 2.18856 |
| 1989 | 2.21250 |
| 1990 | 2.18377 |

These numbers were also extracted by *Andrews, Gentzkow, and Shapiro (2017)* from the original GAUSS code for *Berry, Levinsohn, and Pakes (1999)*.

Type *str*

Examples

The [online version](#) of the following section may be easier to read.

Loading Data Example

```
import pyblp

pyblp.__version__

'0.8.1'
```

Any number of functions can be used to load the example data into memory. In this example, we'll first use [NumPy](#).

```
import numpy as np
blp_product_data = np.recfromcsv(pyblp.data.BLP_PRODUCTS_LOCATION, encoding='utf-8')
blp_agent_data = np.recfromcsv(pyblp.data.BLP_AGENTS_LOCATION, encoding='utf-8')
```

Record arrays can be cumbersome to manipulate. A more flexible alternative is the [pandas DataFrame](#). Unlike NumPy, pyblp does not directly depend on pandas, but it can be useful when manipulating data.

```
import pandas as pd
blp_product_data = pd.read_csv(pyblp.data.BLP_PRODUCTS_LOCATION)
blp_agent_data = pd.read_csv(pyblp.data.BLP_AGENTS_LOCATION)
```

Another benefit of DataFrame objects is that they display nicely in Jupyter notebooks.

```
blp_product_data.head()

  market_ids clustering_ids  car_ids  firm_ids region  shares  prices \
0      1971      AMGREM71      129      15      US  0.001051  4.935802
1      1971      AMHORN71      130      15      US  0.000670  5.516049
2      1971      AMJAVL71      132      15      US  0.000341  7.108642
3      1971      AMMATA71      134      15      US  0.000522  6.839506
4      1971      AMAMBS71      136      15      US  0.000442  8.928395

  hpwt  air  mpd  ...  supply_instruments11  supply_instruments12 \
0  0.528997  0  1.888146  ...                2.024221                -0.251338
1  0.494324  0  1.935989  ...                2.026706                -0.205305
2  0.467613  0  1.716799  ...                1.882594                -0.417634
3  0.426540  0  1.687871  ...                1.846079                -0.445630
```

(continues on next page)

(continued from previous page)

```

4  0.452489    0  1.504286  ...                1.760927                -0.610203

  supply_instruments13  supply_instruments14  supply_instruments15  \
0          0.741272          0.820281          3.691881
1          0.675468          1.004707          3.628727
2          0.946640          0.429942          4.001801
3          0.982922          0.392283          4.061161
4          1.201745          -0.085373          4.340853

  supply_instruments16  supply_instruments17  supply_instruments18  \
0          0.566070          -2.328128          0.385197
1          0.605683          -2.268975          0.363954
2          0.253461          -2.581003          -0.113763
3          0.183331          -2.626074          -0.253065
4          -0.043107          -2.785686          -0.556135

  supply_instruments19  supply_instruments20
0          0.648526          1.052606
1          0.643205          0.906429
2          0.770387          1.511333
3          0.802037          1.583208
4          0.931611          2.177230

```

[5 rows x 47 columns]

blp_agent_data.head()

```

  market_ids  weights  nodes0  nodes1  nodes2  nodes3  nodes4  \
0          1971    0.005  1.764052  0.997845  0.945508 -1.962653  0.330046
1          1971    0.005  0.400157  0.260081  0.422924 -0.834032 -0.000480
2          1971    0.005  0.978738  0.925066 -1.175680  1.993837  0.818116
3          1971    0.005  2.240893  1.476076 -0.204851 -0.666660  0.428214
4          1971    0.005  1.867558 -1.879252  0.956495  0.976290 -2.503947

  income
0    4.172843
1  112.503416
2    2.221875
3    0.814499
4   13.038725

```

5.13 Exceptions

When errors occur, they will either be displayed as warnings or raised as exceptions.

| | |
|---|--|
| <code>exceptions.MultipleErrors</code> | Multiple errors that occurred around the same time. |
| <code>exceptions.NonpositiveCostsError</code> | Encountered nonpositive marginal costs in a log-linear specification. |
| <code>exceptions.NonpositiveSyntheticCostsError</code> | Encountered nonpositive synthetic marginal costs in a log-linear specification. |
| <code>exceptions.InvalidParameterCovariancesError</code> | Failed to compute standard errors because of invalid estimated covariances of GMM parameters. |
| <code>exceptions.InvalidMomentCovariancesError</code> | Failed to compute a weighting matrix because of invalid estimated covariances of GMM moments. |
| <code>exceptions.DeltaNumericalError</code> | Encountered a numerical error when computing δ . |
| <code>exceptions.CostsNumericalError</code> | Encountered a numerical error when computing marginal costs. |
| <code>exceptions.MicroMomentsNumericalError</code> | Encountered a numerical error when computing micro moments. |
| <code>exceptions.XiByThetaJacobianNumericalError</code> | Encountered a numerical error when computing the Jacobian of ξ (equivalently, of δ) with respect to θ . |
| <code>exceptions.OmegaByThetaJacobianNumericalError</code> | Encountered a numerical error when computing the Jacobian of ω (equivalently, of transformed marginal costs) with respect to θ . |
| <code>exceptions.MicroMomentsByThetaJacobianNumericalError</code> | Encountered a numerical error when computing the Jacobian of micro moments with respect to θ . |
| <code>exceptions.MicroMomentCovariancesNumericalError</code> | Encountered a numerical error when computing micro moment covariances. |
| <code>exceptions.SyntheticPricesNumericalError</code> | Encountered a numerical error when computing synthetic prices. |
| <code>exceptions.SyntheticSharesNumericalError</code> | Encountered a numerical error when computing synthetic shares. |
| <code>exceptions.SyntheticDeltaNumericalError</code> | Encountered a numerical error when computing the synthetic δ . |
| <code>exceptions.SyntheticCostsNumericalError</code> | Encountered a numerical error when computing synthetic marginal costs. |
| <code>exceptions.SyntheticMicroMomentsNumericalError</code> | Encountered a numerical error when computing synthetic micro moments. |
| <code>exceptions.EquilibriumRealizationNumericalError</code> | Encountered a numerical error when solving for a realization of equilibrium prices and shares. |
| <code>exceptions.PostEstimationNumericalError</code> | Encountered a numerical error when computing a post-estimation output. |
| <code>exceptions.AbsorptionConvergenceError</code> | An iterative de-meaning procedure failed to converge when absorbing fixed effects. |
| <code>exceptions.ThetaConvergenceError</code> | The optimization routine failed to converge. |
| <code>exceptions.DeltaConvergenceError</code> | The fixed point computation of δ failed to converge. |
| <code>exceptions.SyntheticPricesConvergenceError</code> | The fixed point computation of synthetic prices failed to converge. |
| <code>exceptions.SyntheticDeltaConvergenceError</code> | The fixed point computation of the synthetic δ failed to converge. |
| <code>exceptions.EquilibriumPricesConvergenceError</code> | The fixed point computation of equilibrium prices failed to converge. |

Continued on next page

Table 34 – continued from previous page

| | |
|---|--|
| <code>exceptions.ObjectiveReversionError</code> | Reverted a problematic GMM objective value. |
| <code>exceptions.GradientReversionError</code> | Reverted problematic elements in the GMM objective gradient. |
| <code>exceptions.DeltaReversionError</code> | Reverted problematic elements in δ . |
| <code>exceptions.CostsReversionError</code> | Reverted problematic marginal costs. |
| <code>exceptions.MicroMomentsReversionError</code> | Reverted problematic micro moments. |
| <code>exceptions.XiByThetaJacobianReversionError</code> | Reverted problematic elements in the Jacobian of ξ (equivalently, of δ) with respect to θ . |
| <code>exceptions.OmegaByThetaJacobianReversionError</code> | Reverted problematic elements in the Jacobian of ω (equivalently, of transformed marginal costs) with respect to θ . |
| <code>exceptions.MicroMomentsByThetaJacobianReversionError</code> | Reverted problematic elements in the Jacobian of micro moments with respect to θ . |
| <code>exceptions.AbsorptionInversionError</code> | Failed to invert the A matrix from <i>Somaini and Wolak (2016)</i> when absorbing two-way fixed effects. |
| <code>exceptions.HessianEigenvaluesError</code> | Failed to compute eigenvalues for the GMM objective's (reduced) Hessian matrix. |
| <code>exceptions.FittedValuesInversionError</code> | Failed to invert an estimated covariance when computing fitted values. |
| <code>exceptions.SharesByXiJacobianInversionError</code> | Failed to invert a Jacobian of shares with respect to ξ when computing the Jacobian of ξ (equivalently, of δ) with respect to θ . |
| <code>exceptions.IntraFirmJacobianInversionError</code> | Failed to invert an intra-firm Jacobian of shares with respect to prices when computing η . |
| <code>exceptions.LinearParameterCovariancesInversionError</code> | Failed to invert an estimated covariance matrix of linear parameters. |
| <code>exceptions.GMMParameterCovariancesInversionError</code> | Failed to invert an estimated covariance matrix of GMM parameters. |
| <code>exceptions.GMMMomentCovariancesInversionError</code> | Failed to invert an estimated covariance matrix of GMM moments. |

5.13.1 pyblp.exceptions.MultipleErrors

class `pyblp.exceptions.MultipleErrors`
Multiple errors that occurred around the same time.

5.13.2 pyblp.exceptions.NonpositiveCostsError

class `pyblp.exceptions.NonpositiveCostsError`
Encountered nonpositive marginal costs in a log-linear specification.

This problem can sometimes be mitigated by bounding costs from below, choosing more reasonable initial parameter values, setting more conservative parameter bounds, or using a linear costs specification.

5.13.3 pyblp.exceptions.NonpositiveSyntheticCostsError

class `pyblp.exceptions.NonpositiveSyntheticCostsError`
Encountered nonpositive synthetic marginal costs in a log-linear specification.

This problem can sometimes be mitigated by more reasonable initial parameter values or using a linear costs specification.

5.13.4 `pyblp.exceptions.InvalidParameterCovariancesError`

class `pyblp.exceptions.InvalidParameterCovariancesError`

Failed to compute standard errors because of invalid estimated covariances of GMM parameters.

5.13.5 `pyblp.exceptions.InvalidMomentCovariancesError`

class `pyblp.exceptions.InvalidMomentCovariancesError`

Failed to compute a weighting matrix because of invalid estimated covariances of GMM moments.

5.13.6 `pyblp.exceptions.DeltaNumericalError`

class `pyblp.exceptions.DeltaNumericalError`

Encountered a numerical error when computing δ .

This problem is often due to prior problems, overflow, or nonpositive shares, and can sometimes be mitigated by choosing smaller initial parameter values, setting more conservative bounds, rescaling data, removing outliers, changing the floating point precision, or using different optimization, iteration, or integration configurations.

5.13.7 `pyblp.exceptions.CostsNumericalError`

class `pyblp.exceptions.CostsNumericalError`

Encountered a numerical error when computing marginal costs.

This problem is often due to prior problems or overflow and can sometimes be mitigated by choosing smaller initial parameter values, setting more conservative bounds, rescaling data, removing outliers, changing the floating point precision, or using different optimization or cost configurations.

5.13.8 `pyblp.exceptions.MicroMomentsNumericalError`

class `pyblp.exceptions.MicroMomentsNumericalError`

Encountered a numerical error when computing micro moments.

This problem is often due to prior problems, overflow, or nonpositive shares, and can sometimes be mitigated by choosing smaller initial parameter values, setting more conservative bounds, rescaling data, removing outliers, changing the floating point precision, or using different optimization, iteration, or integration configurations.

5.13.9 `pyblp.exceptions.XiByThetaJacobianNumericalError`

class `pyblp.exceptions.XiByThetaJacobianNumericalError`

Encountered a numerical error when computing the Jacobian of ξ (equivalently, of δ) with respect to θ .

This problem is often due to prior problems, overflow, or nonpositive shares, and can sometimes be mitigated by choosing smaller initial parameter values, setting more conservative bounds, rescaling data, removing outliers, changing the floating point precision, or using different optimization, iteration, or integration configurations.

5.13.10 `pyblp.exceptions.OmegaByThetaJacobianNumericalError`

class `pyblp.exceptions.OmegaByThetaJacobianNumericalError`

Encountered a numerical error when computing the Jacobian of ω (equivalently, of transformed marginal costs) with respect to θ .

This problem is often due to prior problems or overflow, and can sometimes be mitigated by choosing smaller initial parameter values, setting more conservative bounds, rescaling data, removing outliers, changing the floating point precision, or using different optimization or cost configurations.

5.13.11 `pyblp.exceptions.MicroMomentsByThetaJacobianNumericalError`

class `pyblp.exceptions.MicroMomentsByThetaJacobianNumericalError`

Encountered a numerical error when computing the Jacobian of micro moments with respect to θ .

5.13.12 `pyblp.exceptions.MicroMomentCovariancesNumericalError`

class `pyblp.exceptions.MicroMomentCovariancesNumericalError`

Encountered a numerical error when computing micro moment covariances.

5.13.13 `pyblp.exceptions.SyntheticPricesNumericalError`

class `pyblp.exceptions.SyntheticPricesNumericalError`

Encountered a numerical error when computing synthetic prices.

This problem is often due to prior problems or overflow and can sometimes be mitigated by making sure that the specified parameters are reasonable. For example, the parameters on prices should generally imply a downward sloping demand curve.

5.13.14 `pyblp.exceptions.SyntheticSharesNumericalError`

class `pyblp.exceptions.SyntheticSharesNumericalError`

Encountered a numerical error when computing synthetic shares.

This problem is often due to prior problems or overflow and can sometimes be mitigated by making sure that the specified parameters are reasonable. For example, the parameters on prices should generally imply a downward sloping demand curve.

5.13.15 `pyblp.exceptions.SyntheticDeltaNumericalError`

class `pyblp.exceptions.SyntheticDeltaNumericalError`

Encountered a numerical error when computing the synthetic δ .

This problem is often due to prior problems, overflow, or nonpositive shares, and can sometimes be mitigated by making sure that the specified parameters are reasonable.

5.13.16 `pyblp.exceptions.SyntheticCostsNumericalError`

class `pyblp.exceptions.SyntheticCostsNumericalError`
Encountered a numerical error when computing synthetic marginal costs.

This problem is often due to prior problems or overflow and can sometimes be mitigated by making sure that the specified parameters are reasonable.

5.13.17 `pyblp.exceptions.SyntheticMicroMomentsNumericalError`

class `pyblp.exceptions.SyntheticMicroMomentsNumericalError`
Encountered a numerical error when computing synthetic micro moments.

5.13.18 `pyblp.exceptions.EquilibriumRealizationNumericalError`

class `pyblp.exceptions.EquilibriumRealizationNumericalError`
Encountered a numerical error when solving for a realization of equilibrium prices and shares.

5.13.19 `pyblp.exceptions.PostEstimationNumericalError`

class `pyblp.exceptions.PostEstimationNumericalError`
Encountered a numerical error when computing a post-estimation output.

5.13.20 `pyblp.exceptions.AbsorptionConvergenceError`

class `pyblp.exceptions.AbsorptionConvergenceError`
An iterative de-meaning procedure failed to converge when absorbing fixed effects.

This problem can sometimes be mitigated by increasing the maximum number of fixed point iterations, increasing the fixed point tolerance, configuring other iteration settings, or choosing less complicated sets of fixed effects.

5.13.21 `pyblp.exceptions.ThetaConvergenceError`

class `pyblp.exceptions.ThetaConvergenceError`
The optimization routine failed to converge.

This problem can sometimes be mitigated by choosing more reasonable initial parameter values, setting more conservative bounds, or configuring other optimization settings.

5.13.22 `pyblp.exceptions.DeltaConvergenceError`

class `pyblp.exceptions.DeltaConvergenceError`
The fixed point computation of δ failed to converge.

This problem can sometimes be mitigated by increasing the maximum number of fixed point iterations, increasing the fixed point tolerance, choosing more reasonable initial parameter values, setting more conservative bounds, or using different iteration or optimization configurations.

5.13.23 `pyblp.exceptions.SyntheticPricesConvergenceError`

class `pyblp.exceptions.SyntheticPricesConvergenceError`

The fixed point computation of synthetic prices failed to converge.

This problem can sometimes be mitigated by increasing the maximum number of fixed point iterations, increasing the fixed point tolerance, configuring other iteration settings, or making sure the specified parameters are reasonable. For example, the parameters on prices should generally imply a downward sloping demand curve.

5.13.24 `pyblp.exceptions.SyntheticDeltaConvergenceError`

class `pyblp.exceptions.SyntheticDeltaConvergenceError`

The fixed point computation of the synthetic δ failed to converge.

This problem can sometimes be mitigated by increasing the maximum number of fixed point iterations, increasing the fixed point tolerance, choosing more reasonable parameter values, or using a different iteration configuration.

5.13.25 `pyblp.exceptions.EquilibriumPricesConvergenceError`

class `pyblp.exceptions.EquilibriumPricesConvergenceError`

The fixed point computation of equilibrium prices failed to converge.

This problem can sometimes be mitigated by increasing the maximum number of fixed point iterations, increasing the fixed point tolerance, or configuring other iteration settings.

5.13.26 `pyblp.exceptions.ObjectiveReversionError`

class `pyblp.exceptions.ObjectiveReversionError`

Reverted a problematic GMM objective value.

5.13.27 `pyblp.exceptions.GradientReversionError`

class `pyblp.exceptions.GradientReversionError`

Reverted problematic elements in the GMM objective gradient.

5.13.28 `pyblp.exceptions.DeltaReversionError`

class `pyblp.exceptions.DeltaReversionError`

Reverted problematic elements in δ .

5.13.29 `pyblp.exceptions.CostsReversionError`

class `pyblp.exceptions.CostsReversionError`

Reverted problematic marginal costs.

5.13.30 `pyblp.exceptions.MicroMomentsReversionError`

class `pyblp.exceptions.MicroMomentsReversionError`

Reverted problematic micro moments.

5.13.31 `pyblp.exceptions.XiByThetaJacobianReversionError`

class `pyblp.exceptions.XiByThetaJacobianReversionError`
Reverted problematic elements in the Jacobian of ξ (equivalently, of δ) with respect to θ .

5.13.32 `pyblp.exceptions.OmegaByThetaJacobianReversionError`

class `pyblp.exceptions.OmegaByThetaJacobianReversionError`
Reverted problematic elements in the Jacobian of ω (equivalently, of transformed marginal costs) with respect to θ .

5.13.33 `pyblp.exceptions.MicroMomentsByThetaJacobianReversionError`

class `pyblp.exceptions.MicroMomentsByThetaJacobianReversionError`
Reverted problematic elements in the Jacobian of micro moments with respect to θ .

5.13.34 `pyblp.exceptions.AbsorptionInversionError`

class `pyblp.exceptions.AbsorptionInversionError`
Failed to invert the A matrix from *Somains and Wolak (2016)* when absorbing two-way fixed effects.
The formulated fixed effects may be highly collinear.

5.13.35 `pyblp.exceptions.HessianEigenvaluesError`

class `pyblp.exceptions.HessianEigenvaluesError`
Failed to compute eigenvalues for the GMM objective's (reduced) Hessian matrix.

5.13.36 `pyblp.exceptions.FittedValuesInversionError`

class `pyblp.exceptions.FittedValuesInversionError`
Failed to invert an estimated covariance when computing fitted values.
There are probably collinearity issues.

5.13.37 `pyblp.exceptions.SharesByXiJacobianInversionError`

class `pyblp.exceptions.SharesByXiJacobianInversionError`
Failed to invert a Jacobian of shares with respect to ξ when computing the Jacobian of ξ (equivalently, of δ) with respect to θ .

5.13.38 `pyblp.exceptions.IntraFirmJacobianInversionError`

class `pyblp.exceptions.IntraFirmJacobianInversionError`
Failed to invert an intra-firm Jacobian of shares with respect to prices when computing η .

5.13.39 `pyblp.exceptions.LinearParameterCovariancesInversionError`

class `pyblp.exceptions.LinearParameterCovariancesInversionError`

Failed to invert an estimated covariance matrix of linear parameters.

One or more data matrices may be highly collinear.

5.13.40 `pyblp.exceptions.GMMParameterCovariancesInversionError`

class `pyblp.exceptions.GMMParameterCovariancesInversionError`

Failed to invert an estimated covariance matrix of GMM parameters.

One or more data matrices may be highly collinear.

5.13.41 `pyblp.exceptions.GMMMomentCovariancesInversionError`

class `pyblp.exceptions.GMMomentCovariancesInversionError`

Failed to invert an estimated covariance matrix of GMM moments.

One or more data matrices may be highly collinear.

REFERENCES

This page contains a full list of references cited in the documentation, including the original work of *Berry, Levinsohn, and Pakes (1995)*. If you use `pyblp` in your research, we ask that you also cite the below *Conlon and Gortmaker (2019)*, which describes the advances implemented in the package.

6.1 Conlon and Gortmaker (2019)

Conlon, Christopher T., and Jeff Gortmaker (2019). Best practices for differentiated products demand estimation with `pyblp`. Working paper.

6.2 Other References

6.2.1 Amemiya (1977)

Amemiya, Takeshi (1977). A note on a heteroscedastic model. *Journal of Econometrics*, 6 (3), 365-370.

6.2.2 Andrews, Gentzkow, and Shapiro (2017)

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *Quarterly Journal of Economics*, 132 (4), 1553-1592.

6.2.3 Armstrong (2016)

Armstrong, Timothy B. (2016). Large market asymptotics for differentiated product demand estimators with economic models of supply. *Econometrica*, 84 (5), 1961-1980.

6.2.4 Berry (1994)

Berry, Steven (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25 (2), 242-262.

6.2.5 Berry, Levinsohn, and Pakes (1995)

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 63 (4), 841-890.

6.2.6 Berry, Levinsohn, and Pakes (1999)

Berry, Steven, James Levinsohn, and Ariel Pakes (1999). Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review*, 83 (9), 400-430.

6.2.7 Berry, Levinsohn, and Pakes (2004)

Berry, Steven, James Levinsohn, and Ariel Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112 (1), 68-105.

6.2.8 Brenkers and Verboven (2006)

Brenkers, Randy, and Frank Verboven (2006). Liberalizing a distribution system: The European car market. *Journal of the European Economic Association*, 4 (1), 216-251.

6.2.9 Brunner, Heiss, Romahn, and Weiser (2017)

Brunner, Daniel, Florian Heiss, André Romahn, and Constantin Weiser (2017) Reliable estimation of random coefficient logit demand models. DICE Discussion Paper 267.

6.2.10 Cardell (1997)

Cardell, N. Scott (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13 (2), 185-213.

6.2.11 Chamberlain (1987)

Chamberlain, Gary (1987) Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics*, 34 (3), 305-334.

6.2.12 Conlon and Mortimer (2018)

Conlon, Christopher T., and Julie H. Mortimer (2018). Empirical properties of diversion ratios. NBER working paper 24816.

6.2.13 Frisch and Waugh (1933)

Frisch, Ragnar, and Frederick V. Waugh (1933). Partial time regressions as compared with individual trends. *Econometrica*, 1 (4), 387-401.

6.2.14 Gandhi and Houde (2017)

Gandhi, Amit, and Jean-François Houde (2017). Measuring substitution patterns in differentiated products industries. Working paper.

6.2.15 Grigolon and Verboven (2014)

Grigolon, Laura, and Frank Verboven (2014). Nested logit or random coefficients logit? A comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics*, 96 (5), 916-935.

6.2.16 Guimarães and Portugal (2010)

Guimarães, Paulo, and Pedro Portugal (2010). A simple feasible procedure to fit models with high-dimensional fixed effects. *Stata Journal*, 10 (4), 628-649.

6.2.17 Hansen (1982)

Hansen, Lars Peter (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50 (4), 1029-1054.

6.2.18 Heiss and Winschel (2008)

Heiss, Florian, and Viktor Winschel (2008). Likelihood approximation by numerical integration on sparse grids. *Journal of Econometrics*, 144 (1), 62-80.

6.2.19 Imbens and Lancaster (1994)

Imbens, Guido W., and Tony Lancaster (1994). Combining micro and macro data in microeconomic models. *Review of Economic Studies*, 61 (4), 655-680.

6.2.20 Judd and Skrainka (2011)

Judd, Kenneth L., and Ben Skrainka (2011). High performance quadrature rules: How numerical integration affects a popular model of product differentiation. CeMMAP working paper CWP03/11.

6.2.21 Knittel and Metaxoglou (2014)

Knittel, Christopher R., and Konstantinos Metaxoglou (2014). Estimation of random-coefficient demand models: Two empiricists' perspective. *Review of Economics and Statistics*, 96 (1), 34-59.

6.2.22 Lovell (1963)

Lovell, Michael C. (1963). Seasonal adjustment of economic time series and multiple regression analysis. *Journal of the American Statistical Association*, 58 (304), 993-1010.

6.2.23 Morrow and Skerlos (2011)

Morrow, W. Ross, and Steven J. Skerlos (2011). Fixed-point approaches to computing Bertrand-Nash equilibrium prices under mixed-logit demand. *Operations Research*, 59 (2), 328-345.

6.2.24 Nevo (1997)

Nevo, Aviv (1997). Mergers with differentiated products: The case of the ready-to-eat cereal industry. Competition Policy Center, Working Paper Series qt1d53t6ts, Competition Policy Center, Institute for Business and Economic Research, UC Berkeley.

6.2.25 Nevo (2000)

Nevo, Aviv (2000). A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy*, 9 (4), 513-548.

6.2.26 Newey and West (1987)

Newey, Whitney K., and Kenneth D. West (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 28 (3), 777-787.

6.2.27 Petrin (2002)

Petrin, Amil (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110 (4), 705-729.

6.2.28 Reynaert and Verboven (2014)

Reynaert, Mathias, and Frank Verboven (2014). Improving the performance of random coefficients demand models: The role of optimal instruments. *Journal of Econometrics*, 179 (1), 83-98.

6.2.29 Reynaerts, Varadhan, and Nash (2012)

Reynaerts, Jo, Ravi Varadhan, and John C. Nash (2012). Enhancing the convergence properties of the BLP (1995) contraction mapping. VIVES discussion paper 35.

6.2.30 Rios-Avila (2015)

Rios-Avila, Fernando (2015). Feasible fitting of linear models with N fixed effects. *Stata Journal*, 15 (3), 881-898.

6.2.31 Skrainka (2012)

Skrainka, Benjamin S. (2012). A large scale study of the small sample performance of random coefficient models of demand.

6.2.32 Somaini and Wolak (2016)

Somaini, Paulo, and Frank A. Wolak (2016). An algorithm to estimate the two-way fixed effects model. *Journal of Econometric Methods*, 5 (1), 143-152.

6.2.33 Varadhan and Roland (2008)

Varadhan, Ravi, and Christophe Roland (2008). Simple and globally convergent methods for accelerating the convergence of any EM algorithm. *Scandinavian Journal of Statistics*, 35 (2), 335-353.

LEGAL

Copyright 2019 Jeff Gortmaker

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Part II

Developer Documentation

CONTRIBUTING

Please use the [GitHub issue tracker](#) to report bugs or to request features. Contributions are welcome. Examples include:

- Code optimizations.
- Documentation improvements.
- Alternate formulations that have been implemented in the literature but not in pyblp.

TESTING

Testing is done with the `tox` automation tool, which runs a `pytest`-backed test suite in the `tests/` directory. This [FAQ](#) contains some useful information about how to use `tox` on Windows.

9.1 Testing Requirements

In addition to the installation requirements for the package itself, running tests and building documentation requires additional packages specified by the `tests` and `docs` extras in `setup.py`, along with any other explicitly specified deps in `tox.ini`.

The full suite of tests also requires installation of the following software:

- [Artleys Knitro](#) version 10.3 or newer: testing optimization routines.
- [MATLAB](#): comparing sparse grids with those created by the function `nwspgr` created by Florian Heiss and Viktor Winschel, which must be included in a directory on the MATLAB path.

If software is not installed, its associated tests will be skipped. Additionally, some tests that require support for extended precision will be skipped if on the platform running the tests, `numpy.longdouble` has the same precision as `numpy.float64`. This tends to be the case on Windows.

9.2 Running Tests

Defined in `tox.ini` are environments that test the package under different python versions, check types, enforce style guidelines, verify the integrity of the documentation, and release the package. The following command can be run in the top-level `pyblp` directory to run all testing environments:

```
tox
```

You can choose to run only one environment, such as the one that builds the documentation, with the `-e` flag:

```
tox -e docs
```

9.3 Test Organization

Fixtures, which are defined in `tests.conftest`, configure the testing environment and simulate problems according to a range of specifications.

Most BLP-specific tests in `tests.test_blp` verify properties about results obtained by solving the simulated problems under various parameterizations. Examples include:

- Reasonable formulations of problems should give rise to estimated parameters that are close to their true values.
- Cosmetic changes such as the number of processes should not change estimates.
- Post-estimation outputs should satisfy certain properties.
- Optimization routines should behave as expected.
- Derivatives computed with finite differences should approach analytic derivatives.

Tests of generic utilities in `tests.test_formulation`, `tests.test_integration`, `tests.test_iteration`, and `tests.test_optimization` verify that matrix formulation, integral approximation, fixed point iteration, and nonlinear optimization all work as expected. Example include:

- Nonlinear formulas give rise to expected matrices and derivatives.
- Gauss-Hermite integrals are better approximated with quadrature based on Gauss-Hermite rules than with Monte Carlo integration.
- To solve a fixed point iteration problem for which it was developed, SQUAREM requires fewer fixed point evaluations than does simple iteration.
- All optimization routines manage to solve a well-known optimization problem under different parameterizations.

VERSION NOTES

These notes will only include major changes.

10.1 0.8

- Micro moments that match product and agent characteristic covariances
- Extended use of pseudo-inverses
- Added more information to error messages
- More flexible simulation interface
- Alternative way to simulate data with specified prices and shares
- Tests of overidentifying and model restrictions
- Report projected gradients and reduced Hessians
- Change objective gradient scaling
- Switch to a lower-triangular covariance matrix to fix a bug with off-diagonal parameters

10.2 0.7

- Support more fixed point and optimization solvers
- Hessian computation with finite differences
- Simplified interface for firm changes
- Construction of differentiation instruments
- Add collinearity checks
- Update notation and explanations

10.3 0.6

- Optimal instrument estimation
- Structured all results as classes
- Additional information in progress reports

- Parametric bootstrapping of post-estimation outputs
- Replaced all examples in the documentation with Jupyter notebooks
- Updated the instruments for the BLP example problem
- Improved support for multiple equation GMM
- Made concentrating out linear parameters optional
- Better support for larger nesting parameters
- Improved robustness to overflow

10.4 0.5

- Estimation of nesting parameters
- Performance improvements for matrix algebra and matrix construction
- Support for Python 3.7
- Computation of reasonable default bounds on nonlinear parameters
- Additional information in progress updates
- Improved error handling and documentation
- Simplified multiprocessing interface
- Cancelled out delta in the nonlinear contraction to improve performance
- Additional example data and improvements to the example problems
- Cleaned up covariance estimation
- Added type annotations and overhauled the testing suite

10.5 0.4

- Estimation of a Logit benchmark model
- Support for fixing of all nonlinear parameters
- More efficient two-way fixed effect absorption
- Clustered standard errors

10.6 0.3

- Patsy- and SymPy-backed R-style formula API
- More informative errors and displays of information
- Absorption of arbitrary fixed effects
- Reduction of memory footprint

10.7 0.2

- Improved support for longdouble precision
- Custom ownership matrices
- New benchmarking statistics
- Supply-side gradient computation
- Improved configuration for the automobile example problem

10.8 0.1

- Initial release

Part III

Indices

PYTHON MODULE INDEX

p

`pyblp.data`, [165](#)

`pyblp.options`, [164](#)

A

AbsorptionConvergenceError (class in *pyblp.exceptions*), 174
 AbsorptionInversionError (class in *pyblp.exceptions*), 176
 agent_data (*pyblp.Simulation* attribute), 151
 agent_formulation (*pyblp.Problem* attribute), 111
 agent_formulation (*pyblp.Simulation* attribute), 150
 Agents (class in *pyblp*), 160
 agents (*pyblp.Problem* attribute), 112
 agents (*pyblp.Simulation* attribute), 151

B

beta (*pyblp.ProblemResults* attribute), 123
 beta (*pyblp.Simulation* attribute), 151
 beta_bounds (*pyblp.ProblemResults* attribute), 124
 beta_se (*pyblp.ProblemResults* attribute), 123
 BLP_AGENTS_LOCATION (in module *pyblp.data*), 166
 BLP_PRODUCTS_LOCATION (in module *pyblp.data*), 166
 bootstrap() (*pyblp.ProblemResults* method), 139
 bootstrapped_beta (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_delta (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_gamma (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_pi (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_prices (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_rho (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_shares (*pyblp.BootstrappedResults* attribute), 142
 bootstrapped_sigma (*pyblp.BootstrappedResults* attribute), 142
 BootstrappedResults (class in *pyblp*), 141
 build_blp_instruments() (in module *pyblp*), 94
 build_differentiation_instruments() (in module *pyblp*), 97

build_id_data() (in module *pyblp*), 102
 build_integration() (in module *pyblp*), 107
 build_matrix() (in module *pyblp*), 90
 build_ownership() (in module *pyblp*), 104

C

clipped_costs (*pyblp.ProblemResults* attribute), 124
 clustering_ids (*pyblp.Products* attribute), 159
 collinear_atol (in module *pyblp.options*), 165
 collinear_rtol (in module *pyblp.options*), 165
 computation_time (*pyblp.BootstrappedResults* attribute), 142
 computation_time (*pyblp.OptimalInstrumentResults* attribute), 144
 computation_time (*pyblp.SimulationResults* attribute), 156
 compute_aggregate_elasticities() (*pyblp.ProblemResults* method), 130
 compute_approximate_prices() (*pyblp.ProblemResults* method), 134
 compute_consumer_surpluses() (*pyblp.ProblemResults* method), 138
 compute_costs() (*pyblp.ProblemResults* method), 134
 compute_diversion_ratios() (*pyblp.ProblemResults* method), 131
 compute_elasticities() (*pyblp.ProblemResults* method), 131
 compute_hhi() (*pyblp.ProblemResults* method), 136
 compute_long_run_diversion_ratios() (*pyblp.ProblemResults* method), 132
 compute_markup() (*pyblp.ProblemResults* method), 137
 compute_micro() (*pyblp.SimulationResults* method), 158
 compute_optimal_instruments() (*pyblp.ProblemResults* method), 140
 compute_prices() (*pyblp.ProblemResults* method), 135
 compute_probabilities() (*pyblp.ProblemResults* method), 132

- compute_profits() (*pyblp.ProblemResults method*), 137
- compute_shares() (*pyblp.ProblemResults method*), 136
- contraction_evaluations (*pyblp.BootstrappedResults attribute*), 142
- contraction_evaluations (*pyblp.OptimalInstrumentResults attribute*), 145
- contraction_evaluations (*pyblp.ProblemResults attribute*), 122
- contraction_evaluations (*pyblp.SimulationResults attribute*), 156
- converged (*pyblp.ProblemResults attribute*), 121
- costs_type (*pyblp.Problem attribute*), 112
- costs_type (*pyblp.Simulation attribute*), 152
- CostsNumericalError (*class in pyblp.exceptions*), 172
- CostsReversionError (*class in pyblp.exceptions*), 175
- cumulative_contraction_evaluations (*pyblp.ProblemResults attribute*), 122
- cumulative_converged (*pyblp.ProblemResults attribute*), 121
- cumulative_fp_converged (*pyblp.ProblemResults attribute*), 122
- cumulative_fp_iterations (*pyblp.ProblemResults attribute*), 122
- cumulative_objective_evaluations (*pyblp.ProblemResults attribute*), 122
- cumulative_optimization_iterations (*pyblp.ProblemResults attribute*), 122
- cumulative_optimization_time (*pyblp.ProblemResults attribute*), 121
- cumulative_total_time (*pyblp.ProblemResults attribute*), 121
- ## D
- D (*pyblp.Problem attribute*), 113
- D (*pyblp.Simulation attribute*), 152
- delta (*pyblp.ProblemResults attribute*), 124
- DeltaConvergenceError (*class in pyblp.exceptions*), 174
- DeltaNumericalError (*class in pyblp.exceptions*), 172
- DeltaReversionError (*class in pyblp.exceptions*), 175
- demand_ids (*pyblp.Products attribute*), 159
- demand_instruments (*pyblp.OptimalInstrumentResults attribute*), 144
- demand_shifter_formulation (*pyblp.OptimalInstrumentResults attribute*), 144
- demographics (*pyblp.Agents attribute*), 160
- digits (*in module pyblp.options*), 164
- draws (*pyblp.BootstrappedResults attribute*), 142
- draws (*pyblp.OptimalInstrumentResults attribute*), 144
- dtype (*in module pyblp.options*), 164
- ## E
- ED (*pyblp.Problem attribute*), 113
- ED (*pyblp.Simulation attribute*), 153
- EquilibriumPricesConvergenceError (*class in pyblp.exceptions*), 175
- EquilibriumRealizationNumericalError (*class in pyblp.exceptions*), 174
- ES (*pyblp.Problem attribute*), 113
- ES (*pyblp.Simulation attribute*), 153
- expected_omega_by_theta_jacobian (*pyblp.OptimalInstrumentResults attribute*), 144
- expected_prices (*pyblp.OptimalInstrumentResults attribute*), 144
- expected_xi_by_theta_jacobian (*pyblp.OptimalInstrumentResults attribute*), 144
- extract_diagonal_means() (*pyblp.ProblemResults method*), 133
- extract_diagonals() (*pyblp.ProblemResults method*), 133
- ## F
- F (*pyblp.Problem attribute*), 112
- F (*pyblp.Simulation attribute*), 152
- firm_ids (*pyblp.Products attribute*), 159
- FirstChoiceCovarianceMoment (*class in pyblp*), 120
- FittedValuesInversionError (*class in pyblp.exceptions*), 176
- Formulation (*class in pyblp*), 77
- fp_converged (*pyblp.BootstrappedResults attribute*), 142
- fp_converged (*pyblp.OptimalInstrumentResults attribute*), 144
- fp_converged (*pyblp.ProblemResults attribute*), 122
- fp_converged (*pyblp.SimulationResults attribute*), 156
- fp_iterations (*pyblp.BootstrappedResults attribute*), 142
- fp_iterations (*pyblp.OptimalInstrumentResults attribute*), 145
- fp_iterations (*pyblp.ProblemResults attribute*), 122
- fp_iterations (*pyblp.SimulationResults attribute*), 156
- ## G
- gamma (*pyblp.ProblemResults attribute*), 123
- gamma (*pyblp.Simulation attribute*), 151

- gamma_bounds (*pyblp.ProblemResults* attribute), 124
 gamma_se (*pyblp.ProblemResults* attribute), 123
 GMMMomentCovariancesInversionError (class in *pyblp.exceptions*), 177
 GMMParameterCovariancesInversionError (class in *pyblp.exceptions*), 177
 gradient (*pyblp.ProblemResults* attribute), 125
 GradientReversionError (class in *pyblp.exceptions*), 175
- ## H
- H (*pyblp.Problem* attribute), 113
 H (*pyblp.Simulation* attribute), 153
 hessian (*pyblp.ProblemResults* attribute), 125
 HessianEigenvaluesError (class in *pyblp.exceptions*), 176
- ## I
- I (*pyblp.Problem* attribute), 112
 I (*pyblp.Simulation* attribute), 152
 Integration (class in *pyblp*), 80
 integration (*pyblp.Simulation* attribute), 151
 IntraFirmJacobianInversionError (class in *pyblp.exceptions*), 176
 InvalidMomentCovariancesError (class in *pyblp.exceptions*), 172
 InvalidParameterCovariancesError (class in *pyblp.exceptions*), 172
 inverse_covariance_matrix (*pyblp.OptimalInstrumentResults* attribute), 144
 Iteration (class in *pyblp*), 82
- ## K
- K1 (*pyblp.Problem* attribute), 112
 K1 (*pyblp.Simulation* attribute), 152
 K2 (*pyblp.Problem* attribute), 112
 K2 (*pyblp.Simulation* attribute), 152
 K3 (*pyblp.Problem* attribute), 112
 K3 (*pyblp.Simulation* attribute), 152
- ## L
- last_results (*pyblp.ProblemResults* attribute), 121
 LinearParameterCovariancesInversionError (class in *pyblp.exceptions*), 177
- ## M
- market_ids (*pyblp.Agents* attribute), 160
 market_ids (*pyblp.Products* attribute), 159
 MD (*pyblp.Problem* attribute), 113
 MD (*pyblp.Simulation* attribute), 152
 micro (*pyblp.ProblemResults* attribute), 124
 micro_by_theta_jacobian (*pyblp.ProblemResults* attribute), 125
 MicroMomentCovariancesNumericalError (class in *pyblp.exceptions*), 173
 MicroMomentsByThetaJacobianNumericalError (class in *pyblp.exceptions*), 173
 MicroMomentsByThetaJacobianReversionError (class in *pyblp.exceptions*), 176
 MicroMomentsNumericalError (class in *pyblp.exceptions*), 172
 MicroMomentsReversionError (class in *pyblp.exceptions*), 175
 MS (*pyblp.Problem* attribute), 113
 MS (*pyblp.Simulation* attribute), 153
 MultipleErrors (class in *pyblp.exceptions*), 171
- ## N
- N (*pyblp.Problem* attribute), 112
 N (*pyblp.Simulation* attribute), 152
 nesting_ids (*pyblp.Products* attribute), 159
 NEVO_AGENTS_LOCATION (in module *pyblp.data*), 166
 NEVO_PRODUCTS_LOCATION (in module *pyblp.data*), 165
 nodes (*pyblp.Agents* attribute), 160
 NonpositiveCostsError (class in *pyblp.exceptions*), 171
 NonpositiveSyntheticCostsError (class in *pyblp.exceptions*), 171
- ## O
- objective (*pyblp.ProblemResults* attribute), 124
 objective_evaluations (*pyblp.ProblemResults* attribute), 122
 ObjectiveReversionError (class in *pyblp.exceptions*), 175
 omega (*pyblp.ProblemResults* attribute), 124
 omega (*pyblp.Simulation* attribute), 152
 omega_by_theta_jacobian (*pyblp.ProblemResults* attribute), 124
 OmegaByThetaJacobianNumericalError (class in *pyblp.exceptions*), 173
 OmegaByThetaJacobianReversionError (class in *pyblp.exceptions*), 176
 OptimalInstrumentProblem (class in *pyblp*), 147
 OptimalInstrumentResults (class in *pyblp*), 143
 Optimization (class in *pyblp*), 86
 optimization_iterations (*pyblp.ProblemResults* attribute), 121
 optimization_time (*pyblp.ProblemResults* attribute), 121
 ownership (*pyblp.Products* attribute), 159
- ## P
- parallel() (in module *pyblp*), 161

- parameter_covariances (*pyblp.ProblemResults attribute*), 122
- parameters (*pyblp.ProblemResults attribute*), 122
- pi (*pyblp.ProblemResults attribute*), 123
- pi (*pyblp.Simulation attribute*), 151
- pi_bounds (*pyblp.ProblemResults attribute*), 123
- pi_se (*pyblp.ProblemResults attribute*), 123
- PostEstimationNumericalError (*class in pyblp.exceptions*), 174
- prices (*pyblp.Products attribute*), 160
- Problem (*class in pyblp*), 109
- problem (*pyblp.ProblemResults attribute*), 121
- problem_results (*pyblp.BootstrappedResults attribute*), 141
- problem_results (*pyblp.OptimalInstrumentResults attribute*), 144
- ProblemResults (*class in pyblp*), 121
- product_data (*pyblp.Simulation attribute*), 151
- product_data (*pyblp.SimulationResults attribute*), 156
- product_formulations (*pyblp.Problem attribute*), 111
- product_formulations (*pyblp.Simulation attribute*), 150
- Products (*class in pyblp*), 159
- products (*pyblp.Problem attribute*), 112
- products (*pyblp.Simulation attribute*), 151
- projected_gradient (*pyblp.ProblemResults attribute*), 125
- projected_gradient_norm (*pyblp.ProblemResults attribute*), 125
- psd_atol (*in module pyblp.options*), 165
- psd_rtol (*in module pyblp.options*), 165
- pseudo_inverses (*in module pyblp.options*), 165
- pyblp.data (*module*), 165
- pyblp.options (*module*), 164
- ## R
- reduced_hessian (*pyblp.ProblemResults attribute*), 125
- reduced_hessian_eigenvalues (*pyblp.ProblemResults attribute*), 125
- replace_endogenous () (*pyblp.Simulation method*), 153
- replace_exogenous () (*pyblp.Simulation method*), 155
- rho (*pyblp.ProblemResults attribute*), 123
- rho (*pyblp.Simulation attribute*), 152
- rho_bounds (*pyblp.ProblemResults attribute*), 123
- rho_se (*pyblp.ProblemResults attribute*), 123
- run_distance_test () (*pyblp.ProblemResults method*), 128
- run_hansen_test () (*pyblp.ProblemResults method*), 127
- run_lm_test () (*pyblp.ProblemResults method*), 129
- run_wald_test () (*pyblp.ProblemResults method*), 129
- ## S
- shares (*pyblp.Products attribute*), 159
- SharesByXiJacobianInversionError (*class in pyblp.exceptions*), 176
- sigma (*pyblp.ProblemResults attribute*), 123
- sigma (*pyblp.Simulation attribute*), 151
- sigma_bounds (*pyblp.ProblemResults attribute*), 123
- sigma_se (*pyblp.ProblemResults attribute*), 123
- Simulation (*class in pyblp*), 147
- simulation (*pyblp.SimulationResults attribute*), 156
- SimulationResults (*class in pyblp*), 156
- solve () (*pyblp.Problem method*), 113
- step (*pyblp.ProblemResults attribute*), 121
- supply_ids (*pyblp.Products attribute*), 159
- supply_instruments (*pyblp.OptimalInstrumentResults attribute*), 144
- supply_shifter_formulation (*pyblp.OptimalInstrumentResults attribute*), 144
- SyntheticCostsNumericalError (*class in pyblp.exceptions*), 174
- SyntheticDeltaConvergenceError (*class in pyblp.exceptions*), 175
- SyntheticDeltaNumericalError (*class in pyblp.exceptions*), 173
- SyntheticMicroMomentsNumericalError (*class in pyblp.exceptions*), 174
- SyntheticPricesConvergenceError (*class in pyblp.exceptions*), 175
- SyntheticPricesNumericalError (*class in pyblp.exceptions*), 173
- SyntheticSharesNumericalError (*class in pyblp.exceptions*), 173
- ## T
- T (*pyblp.Problem attribute*), 112
- T (*pyblp.Simulation attribute*), 152
- theta (*pyblp.ProblemResults attribute*), 123
- ThetaConvergenceError (*class in pyblp.exceptions*), 174
- tilde_costs (*pyblp.ProblemResults attribute*), 124
- to_dict () (*pyblp.BootstrappedResults method*), 143
- to_dict () (*pyblp.OptimalInstrumentResults method*), 145
- to_dict () (*pyblp.ProblemResults method*), 127
- to_dict () (*pyblp.SimulationResults method*), 157
- to_problem () (*pyblp.OptimalInstrumentResults method*), 146
- to_problem () (*pyblp.SimulationResults method*), 157

`total_time` (*pyblp.ProblemResults* attribute), 121

U

`unique_firm_ids` (*pyblp.Problem* attribute), 112

`unique_firm_ids` (*pyblp.Simulation* attribute), 151

`unique_market_ids` (*pyblp.Problem* attribute), 112

`unique_market_ids` (*pyblp.Simulation* attribute),
151

`unique_nesting_ids` (*pyblp.Problem* attribute),
112

`unique_nesting_ids` (*pyblp.Simulation* attribute),
151

`updated_W` (*pyblp.ProblemResults* attribute), 125

V

`verbose` (*in module pyblp.options*), 164

`verbose_output` (*in module pyblp.options*), 164

`verbose_tracebacks` (*in module pyblp.options*),
164

W

`W` (*pyblp.ProblemResults* attribute), 125

`weights` (*pyblp.Agents* attribute), 160

X

`X1` (*pyblp.Products* attribute), 160

`X2` (*pyblp.Products* attribute), 160

`X3` (*pyblp.Products* attribute), 160

`xi` (*pyblp.ProblemResults* attribute), 124

`xi` (*pyblp.Simulation* attribute), 152

`xi_by_theta_jacobian` (*pyblp.ProblemResults* at-
tribute), 124

`XiByThetaJacobianNumericalError` (*class in*
pyblp.exceptions), 172

`XiByThetaJacobianReversionError` (*class in*
pyblp.exceptions), 176

Z

`ZD` (*pyblp.Products* attribute), 160

`ZS` (*pyblp.Products* attribute), 160