A Estimating $D(q)$ Overview

There are several ways to estimate $D_i(q)$, the marginal willingness to pay in time for the q-th extra option. The goal is always to keep time disutility in the supply curve $S_i(q) = c_i(q)$, while $D_i(q)$ captures only the marginal value of variety. This section provides a high-level overview, with details in Sections B and C.

- **Choice-grounded methods (Section B).** Estimate a destination-choice model, compute inclusive-value increments at fixed ranks, and map utils to minutes. This isolates the marginal value of options while keeping time disutility in $S(q)$.
- **Flow-based calibration (Section C).** Estimate a gravity-style model using Poisson Pseudo Maximum Likelihood (PPML). Then, recover the shape of $D(q)$ from inclusive-value increments as options are added by rank, mapping to minutes via the time coefficient.
- **Reduced-form sequences.** Where behavioural data are limited, one can fit a simple monotone decreasing sequence to represent $D(q)$, e.g. $D(q) = \alpha q^{-\beta}$ or $D(q) = \gamma - \delta \ln(1 + q)$. These can be anchored by observed search limits or survey evidence.

In all cases, the aim is for $D(q)$ to reflect diminishing marginal value of variety, while $S(q)$ reflects the disutility of time or cost.

B Recovering $D(q)$ from observed trip times

Let T be observed chosen times for a population with heterogeneous λ and V. For small bins t_k, the survival $1 - F_T(t_k)$ approximates the share of people for whom $D(q) \geq S(q)$ with $q = O(t_k)$. With a parametric $D(q;\theta)$ (e.g., $D = \alpha(1 + q)^{-\beta}$), estimate θ by minimum squared distance between empirical survival and model-implied $\Pr[D \geq S]$; uncertainty by bootstrap.

C Estimating $D(q)$ from flows via PPML

Estimate a flow model that absorbs origin and destination size, leaving the shape in generalised cost:

$$\mathbb{E}[F_{ij}] \;=\; \exp\big(\alpha_i + \gamma_j + g(c_{ij})\big), \tag{9}$$

with PPML and robust errors. Sort destinations by c_{ij} for each origin i. For $q = 1, 2, \ldots$, add the q-th nearest destination to the choice set and compute the inclusive-value increment; convert utils to minutes using the estimated time coefficient to obtain $D_i(q)$. Keep the disutility of time on the supply side $S_i(q) = c_i(q)$, so $D_i(q)$ represents the marginal value of variety by rank. A reduced-form alternative fits a monotone $D_i(q)$, for example $D_i(q) = \alpha_i q^{-\beta}$, anchored by two moments.

D Separating option value from travel-time cost

In standard destination choice, utility typically stacks time cost and attributes, e.g.,

$$U_{ij} = \beta_t c_{ij} + x'_{ij}\beta + \varepsilon_{ij}, \tag{10}$$

and logsum-based welfare mixes both the disutility of time and the value of having more alternatives. Here, we separate the two by design: the *supply* $S_i(q)$ carries the time cost of unlocking the q-th option, while the *demand* $D_i(q)$ captures the *marginal value of variety*, how much additional time a person is willing to pay for one more option, excluding the intrinsic dislike of time.

This separation matters because q and cost often move together; in idealised settings they may be nearly collinear. To avoid double counting and to identify $D_i(q)$ apart from $S_i(q)$, one might use one of the following strategies:

- *Inclusive-value increments at fixed ranks.* For $q = 1, 2, \dots$, add the q-th nearest option to the set and take the change in the inclusive value; convert utils to the *same units as* c (minutes if c is time). This isolates the marginal option value, while the time disutility remains in $S_i(q) = c_{i(q)}$.
- *Within-isochrone variation.* Fix a budget c_0 (time budget t_0 in the examples) and exploit outcome variation within that contour; holding S fixed at c_0 identifies the local $D_i(q)$.
- *Quasi-experimental shocks to availability.* Use openings/closures or capacity changes that alter counts inside a stable travel-time band (network unchanged). These shocks move q holding S roughly fixed, isolating D.
- *Regularised reduced form.* Where q and cost are tightly correlated (e.g., a uniform plane), D is weakly identified from revealed behaviour alone. In such cases, fit a parsimonious decreasing $D(q)$ (e.g., $\alpha q^{-\beta}$ or $\gamma - \delta \ln(1 + q)$), report sensitivity, and anchor with external evidence (surveys, experiments).

In all cases, keep D free of time disutility that is already represented by S. Intuitively, if travel were instantaneous ($S_i \equiv 0$), $D_i(q)$ would still decline as variety saturates; that is the margin D is meant to capture.

E Welfare interpretation

Suppose individual n faces a choice set of opportunities $\mathcal{J}_n$, utility is given by

$$U_{nj} = V_{nj}(x_{nj}) - \lambda_n t_{nj} + \varepsilon_{nj} \tag{11}$$

where t_{nj} is *generalised cost* (time, money, discomfort); in examples we use travel time, and ε_{nj} is i.i.d. type-I extreme value.

The logsum (inclusive value) is:

$$IV_n = \ln \sum_{j \in \mathcal{J}_n} \exp\{V_{nj} - \lambda_n t_{nj}\} \tag{12}$$

Expected money-metric welfare change from a perturbation of t is, as given by Williams (1977); Small and Rosen (1981); Train (2009):

$$\Delta W_n = \frac{1}{\lambda_n} \Delta IV_n \tag{13}$$

For small changes, treat $\mathbb{A}$ as a first-order approximation to welfare (option value) after mapping utils to the same units as t by the relevant coefficient (minutes if t is time).

Ordering opportunities by increasing t induces a continuous index q of marginal opportunities; with $S(q)$ the inverse of cumulative opportunities and $\bar{D}(q)$ the population average marginal WTP for the q-th opportunity, a first-order expansion shows:

$$\Delta W \approx \int_0^{q^\star} \left[\bar{D}(q) - S(q)\right] dq, \tag{14}$$

matching the Access Surplus ($\mathbb{A}$) in the main text.