Logit Choice and Perturbed Optimization

Minoru Osawa (Created: 2024-12-5; Updated: 2025-01-21)

This note is a brief summary of the well-known connection between the logit choice model and perturbed optimization.

1. Additive random utility models and the logit choice

Consider a decision maker (DM) facing a choice situation. There is a finite set of alternatives, A. The payoff V_a of choosing an alternative $a \in A$ is subject to uncertainty, and may be expressed as a random variable such that

$$V_a = v_a + \epsilon_a, \tag{1}$$

where v_a is known deterministic payoff, and ϵ_a is a random payoff. It is assumed that ϵ_a are i.i.d. across alternatives. It is further assumed that the DM uses randomization, or mixed strategies, so that they choose each alternative a with the probability p_a that a is payoff-maximizing. That is,

$$p_a = \Pr[V_a \geq V_b \ \forall b \in A] = \Pr[v_a + \epsilon_a \geq v_b + \epsilon_b \ \forall b \in A]. \tag{2}$$

This framework is called *additive random utility models* (ARUM).

If ϵ_a is i.i.d. with a differentiable c.d.f. F, we have

$$p_a = \int F'(\epsilon_a) \prod_{b \neq a} F(v_a - v_b + \epsilon_a) \mathrm{d}\epsilon_a. \tag{3}$$

Further suppose that every ϵ_a follows the Gumbel distribution with scale parameter $\eta > 0$ and no location parameter, whose c.d.f. is given as

$$F(\epsilon) \equiv \exp\left(-\exp\left(-\eta^{-1}\epsilon\right)\right) \qquad \epsilon \in (-\infty, \infty). \tag{4}$$

It is known that $\mathbb{E}[\epsilon] = \eta^{-1}\gamma$ with Euler's constant $\gamma \approx 0.5772$, $\mathrm{Var}[\epsilon] = \eta^2\pi^2/6$. The constant η thus represents the magnitude of randomness, and the deterministic payoff v is less (more) relevant for DM's choice when η is large (small).

Under the Gumbel assumption, we obtain the *logit choice rule*:

$$p_a = \frac{\exp(\eta^{-1}v_a)}{\sum_{b \in A} \exp(\eta^{-1}v_b)}. \tag{5}$$

The expected value of the maximized payoff, the *expected maximum utility* (EMU) is

$$\lambda \equiv \mathbb{E}\left[\max_{a \in A} V_a\right] = \eta \log \sum_{a \in A} \exp\left(\eta^{-1}v_a\right) + \eta^{-1}\gamma. \tag{6}$$

It is well known that choice probabilities in the logit model satisfy $\frac{\partial \lambda}{\partial v_a} = p_a$, that is, choice probability vector is the gradient of EMU with respect to the deterministic payoffs. This result also extends to all ARUMs under mild conditions (Williams–Daly–Zachary Theorem) (see Fosgerau et al., 2020).

Note: Computation of p_a and λ

To compute p_a under the Gumbel-distributed ϵ_a, note that, for c.d.f. (4), we have

$$F'(\epsilon) = \rho(\epsilon)F(\epsilon) \quad \text{with} \quad \rho(\epsilon) \equiv \eta^{-1}\exp\left(-\eta^{-1}\epsilon\right) \qquad (= \text{p.d.f.}),$$
$$F(v+\epsilon) = F(\epsilon)^{\exp(-\eta^{-1}v)},$$
$$\{F(\epsilon)^t\}' = tF(\epsilon)^{t-1}F'(\epsilon) = t\rho(\epsilon)F(\epsilon)^t.$$

Then, noting that $\lim_{\epsilon\to 0} F(\epsilon) = 0$ and $\lim_{\epsilon\to\infty} F(\epsilon) = 1$, we see

$$\begin{aligned}
p_a &= \int_{-\infty}^{\infty} F'(\epsilon_a)\prod_{b\neq a} F(v_a - v_b + \epsilon_a)\mathrm{d}\epsilon_a \\
&= \int_{-\infty}^{\infty} \rho(\epsilon_a)F(\epsilon_a)\prod_{b\neq a} F(\epsilon_a)^{\exp(\eta^{-1}(v_b - v_a))}\mathrm{d}\epsilon_a \\
&= \int_{-\infty}^{\infty} \rho(\epsilon_a)F(\epsilon_a)^{1+\sum_{b\neq a}\exp(\eta^{-1}(v_b - v_a))}\mathrm{d}\epsilon_a \\
&= \frac{1}{1+\sum_{b\neq a}\exp(\eta^{-1}(v_b - v_a))}\left[F(\epsilon_a)^{1+\sum_{b\neq a}\exp(\eta^{-1}(v_b - v_a))}\right]_{-\infty}^{\infty} \\
&= \frac{\exp(\eta^{-1}v_a)}{\sum_{b\in A}\exp(\eta^{-1}v_b)}.
\end{aligned}$$

To compute λ, we observe that for $\hat{V} \equiv \max_{a\in A} V_a$,

$$\begin{aligned}
\Pr[\hat{V} \leq x] &= \Pr[\epsilon_a \leq x - v_a \ \forall a \in A] = \prod_{a\in A} F(x - v_a) \\
&= F(x)^{\sum_{a\in A}\exp(\eta^{-1}v_a)} = F(x)^{\exp(\eta^{-1}\lambda_0)} \qquad \text{where} \quad \lambda_0 \equiv \eta \log \sum_{a\in A}\exp(\eta^{-1}v_a) \\
&= F(x - \lambda_0).
\end{aligned}$$

Thus, $\hat{V}$ follows the Gumbel distribution with location parameter λ_0 and scale parameter η, implying $\lambda = \mathbb{E}[\hat{V}] = \lambda_0 + \eta^{-1}\gamma$ as in (6).

2. Mixed-strategy best response and linear optimization problem

Next, consider a simple, deterministic approach. Given alternatives $a \in A$ and payoffs $v = (v_a)$, suppose that the DM's problem is to determine the payoff-maximizing mixed strategy by solving the following linear optimization problem:

$$\max_{y\in\Delta} \quad \langle v, y\rangle \tag{7}$$

where $\Delta \equiv \{y \geq \mathbf{0} \mid \sum_{a\in A} y_a = 1\}$ is the probability simplex and $\langle x, y\rangle$ denotes the inner product of x and y. A solution y^* for this problem should satisfy

$$y_a^* > 0 \Rightarrow a \in \mathrm{br}(v), \tag{8}$$

where $\mathrm{br}(v) \equiv \arg\max_b \{v_b\}_{b\in A}$ is the set of payoff-maximizing alternatives given the payoff vector v. Such y^* form a convex set but uniqueness is not always the case because $\mathrm{br}(v)$ may not be a singleton.

The *dual* problem for (7) is given as

$$\min_{\lambda} \quad \lambda \qquad \text{s.t.} \quad \lambda \geq v_a \quad \forall a \in A. \tag{9}$$

The problem aims to obtain the best (smallest) upper bound for DM's attainable payoff. Evidently, the solution and the optimal value for the problem is $\lambda^* = \max_{a\in A} v_a$ and coincides with the optimal value of (7) (the *strong duality* of linear optimization).

> Note: Derivation of the dual problem
>
> Let λ be the Lagrange multiplier for the constraint $\sum_{a\in A} y_a = 1$. The Lagrangian function is
>
> $$L(y,\lambda) \equiv \langle v, y\rangle - \lambda\left(\langle \mathbf{1}, y\rangle - 1\right) = -\langle \lambda\mathbf{1} - v, y\rangle + \lambda \tag{10}$$
>
> with $y \geq 0$. The Lagrangian dual problem is to minimize the following objective function, implying (9):
>
> $$\omega(\lambda) = \sup_{y\geq 0} L(y,\lambda) = \sup_{y \geq 0} \lambda - \langle \lambda\mathbf{1} - v, y\rangle = \begin{cases} \lambda & \text{if } \lambda \geq v_a \quad \forall a \in A, \\ \infty & \text{otherwise.} \end{cases} \tag{11}$$

3. Perturbed optimization

As seen, the deterministic approach does not provide unique prediction regarding DM's choice. From the mathematical optimization perspective, this stems from the fact that (7) is a linear optimization problem. We can consider adding a regularization term to ensure the uniqueness of the predicted behavior.

Suppose that the DM's problem in (7) is modified as follows:

$$\max_{y\in\Delta} \quad \langle v, y\rangle - H(y) \tag{12}$$

The function $H : \mathrm{int}(\Delta) \to \mathbb{R}$ is assumed to be strictly convex and becomes infinitely steeper as y goes to the boundary of Δ. Since the objective function is strictly concave and the feasible region Δ is convex and compact, the modified problem has unique solution.

Below, as a representative case, suppose that H is the negative entropy

$$H(y) = \eta \sum_{a\in A} y_a \log y_a, \tag{13}$$

where we define $0 \log 0 \equiv 0$. As $\eta \to 0$, the problem (12) recovers the unperturbed problem (7).

The optimal solution y^* is the logit choice rule:

$$y_a^* = p_a = \frac{\exp(\eta^{-1} v_a)}{\sum_{b\in A} \exp(\eta^{-1} v_b)}. \tag{14}$$

The optimal value of the problem (12) is

$$\lambda(v) \equiv \langle v, y^*\rangle - \eta\langle y^*, \log y^*\rangle = \eta \log \sum_{a\in A} \exp(\eta^{-1} v_a). \tag{15}$$

We see that the optimal value can be seen as the expected maximum utility for the logit model. In fact,

$$\frac{\partial \lambda(v)}{\partial v_a} = \frac{\exp(\eta^{-1} v_a)}{\sum_{b\in A} \exp(\eta^{-1} v_b)} = p_a. \tag{16}$$

The optimal value function of (12) is nothing but the convex conjugate (Legendre transform) of H, which also implies the above formula.

The Lagrange dual problem for (12) is

$$\min_{\lambda} \quad \lambda \quad \text{s.t.} \quad \lambda = \eta \log \sum_{a\in A} \exp(\eta^{-1} v_a) \tag{17}$$

whose solution, and hence optimal value, coincides with the optimal value of the primal problem (15) (the strong duality for convex optimization). Observe that $\lambda(v)$ tends to $\lambda^* = \max_{a\in A} v_a$ as $\eta \to 0$. The similarity between the dual problem for the unperturbed case is notable.

Note: Derivations for y_a^* and the Lagrangian dual problem

The Lagrangian function is modified as

$$L(y, v) \equiv -\langle v, y\rangle + \lambda \left(\langle \mathbf{1}, y\rangle - 1\right) + H(y). \tag{18}$$

The optimality condition is

$$y_a \frac{\partial L(y,\lambda)}{\partial y_a} = 0, y_a \geq 0, \frac{\partial L(y,\lambda)}{\partial y_a} = -v_a + \lambda + \eta \log y_a + \eta \geq 0, \tag{19}$$

$$\frac{\partial L(y,\lambda)}{\partial \lambda} = \sum_a y_a - 1 = 0. \tag{20}$$

Since $\frac{\partial L(y,\lambda)}{\partial y_a} \to -\infty$ as $y_a \to 0$, $y_a = 0$ violates (19). Then, $y_a > 0$ and $\frac{\partial L(y,\lambda)}{\partial y_a} = 0$ for all a, implying $y_a = \exp\left(\eta^{-1}(v_a - \lambda) - 1\right)$. Thus, from $\sum_a y_a = 1$, we obtain

$$\lambda = \eta \log \sum_{a\in A} \exp(\eta^{-1} v_a) - \eta. \tag{21}$$

Since $\inf_{y\geq 0} L(y,\lambda) = \lambda + \eta$, the dual problem is equivalent to (17) where we redefine $\lambda \coloneqq \lambda + \eta$.

Observe that when we take the limit $\eta \to 0$, $y_a > 0$ can occur only if $a \in \operatorname{br}(v)$, and $y_a \to 0$ as $\eta \to 0$ if $a \notin \operatorname{br}(v)$, which are consistent with the unperturbed case. To see this, observe

$$y_a = \frac{1}{\sum_{b\in A} \exp(\eta^{-1}(v_b - v_a))}. \tag{22}$$

If $a \notin \operatorname{br}(v)$, $y_a \to 0$ because the denominator goes to infinity as $\eta \to 0$ when $v_b > v_a$ for some b. If $a \in \operatorname{br}(v)$, y_a tends to $1/|\operatorname{br}(v)|$ as $\eta \to 0$, which is slightly different from the unperturbed case where mixed-strategy best response can be nonunique.

Considering a different convex function for H induces a different choice rule. All practically used ARUM have such deterministically perturbed optimization representation but converse is not true.

4. Further readings

- Hofbauer and Sandholm (2002), Theorem 2.1; Hofbauer and Sandholm (2007), Appendix.
- Anderson et al. (1992)

- 土木学会 (1998), Ch.6
- Fudenberg et al. (2015)
- Fosgerau et al. (2020)

References

Anderson, S. P., de Palma, A., and Thisse, J. F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.

Fosgerau, M., Melo, E., De Palma, A., and Shum, M. (2020). Discrete choice and rational inattention: A general equivalence result. *International Economic Review*, 61(4):1569–1589.

Fudenberg, D., Iijima, R., and Strzalecki, T. (2015). Stochastic choice and revealed perturbed utility. *Econometrica*, 83(6):2371–2409.

Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294.

Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132(1):47–69.

土木学会 (1998). 交通ネットワークの均衡分析：最新の理論と解法. 土木学会.