Large Deviations and Multinomial Probit Choice*

Emin Dokumacı[†] and William H. Sandholm[‡]

July 4, 2011

Abstract

We consider a discrete choice model in which the payoffs to each of an agent's n actions are subjected to the average of m i.i.d. shocks, and use tools from large deviations theory to characterize the rate of decay of the probability of choosing a given suboptimal action as m approaches infinity. Our model includes the multinomial probit model of Myatt and Wallace (2003) as a special case. We show that their formula describing the rates of decay of choice probabilities is incorrect, provide the correct formula, and use our large deviations analysis to provide intuition for the difference between the two.

1. Introduction

In a paper in this journal, Myatt and Wallace (2003) consider a model of stochastic evolution based on the multinomial probit model. Agents in their model optimize after their payoffs are subjected to i.i.d. normal shocks, and their analysis focuses on the agents' long run behavior as the variance of the shocks is taken to zero. Compared to other models of choice used in stochastic evolutionary game theory, the multinomial probit model introduces a novel feature: the rate of decay in the probability of choosing a suboptimal strategy is neither independent of payoffs, as in the mutation models of Kandori et al. (1993) and Young (1993), nor dependent only on the gap between its payoff and the optimal strategy's payoff, as in the logit model of Blume (1993), but can depend on the gaps between its payoff and those of all better performing strategies.[1] The foundation

*An early version of the analysis presented here was first circulated in a working paper entitled "Stochastic Evolution with Perturbed Payoffs and Rapid Play". We thank Olvi Mangasarian, David Myatt, and Chris Wallace for helpful comments. Financial support under NSF Grants SES-0617753 and SES-0851580 is gratefully acknowledged.

†Wisconsin Center for Education Research, University of Wisconsin, 1025 West Johnson Street, Madison, WI 53706, USA. e-mail: edokumaci@wisc.edu; website: http://www.ssc.wisc.edu/~edokumac.

‡Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu; website: http://www.ssc.wisc.edu/~whs.

[1] Related models also appear in unpublished work of Ui (1998).

of the analysis in Myatt and Wallace (2003) (henceforth MW) is their Proposition 1, which characterizes the rates of decay of multinomial probit choice probabilities as the shock variance approaches zero. Their characterization is based on a direct evaluation of the limit of the relevant multiple integral.

In this note, we introduce a model of choice in which the payoffs to each of an agent's n actions are subject to the average of m i.i.d. shocks. One can interpret this average as representing the net effect of many small payoff disturbances. Our model comes equipped with a natural parameterization of the small noise limit: as the number of shocks grows large, the probability of a suboptimal choice approaches zero. Using techniques from large deviations theory, we derive basic monotonicity and convexity properties of the rates of decay of choice probabilities, and we obtain a simple characterization of the rates themselves.

Since the average of independent normal random variables is itself normally distributed, MW's model of choice can be obtained as a special case of ours. Our analysis reveals that MW's formula for the rate of decay of multinomial probit choice probabilities is incorrect. We derive the correct formula for the rate of decay, and we offer an intuitive explanation for the difference between the formulas using the language of large deviations theory.

2. Analysis

2.1 Large Deviations and Cramér's Theorem

Let $\{Z^l\}_{l=1}^{\infty}$ be an i.i.d. sequence of random vectors taking values in $\mathbb{R}^n$. Each random vector Z^l is continuous with convex support, with a moment generating function that exists in a neighborhood of the origin.

Let $\bar{Z}^m = \frac{1}{m}\sum_{l=1}^{m} Z^l$ denote the mth sample mean of the sequence $\{Z^l\}_{l=1}^{\infty}$. The weak law of large numbers tells us that $\bar{Z}^m$ converges in probability to its mean vector $\mu \equiv \mathbb{E}Z^l \in \mathbb{R}^n$. We now explain how methods from large deviations theory can be used to describe the rate of decay of the probability that $\bar{Z}^m$ lies in a given set $U \subset \mathbb{R}^n$ not containing μ.

The *Cramér transform* of Z^l, denoted $R\colon \mathbb{R}^n \to [0, \infty]$, is defined by

$$R(z) = \sup_{\lambda \in \mathbb{R}^n} (\lambda' z - \Lambda(\lambda)), \quad \text{where} \quad \Lambda(\lambda) = \log \mathbb{E} \exp(\lambda' Z^l).$$

Put differently, R is the convex conjugate of the logarithmic moment generating function of Z^l. It can be shown that R is a convex, lower semicontinuous, nonnegative function that satisfies $R(\mu) = 0$. Moreover, R is finite, strictly convex, and continuously differentiable on the interior of the support of Z^l, and is infinite outside the support of Z^l.[2]

For simplicity, we henceforth assume that the components of the random vector $Z^l = (Z_1^l, \ldots, Z_n^l)$ are independent. It is easy to verify that in this case, the Cramér transform of Z^l is $R(z) = \sum_{k=1}^n r_k(z_k)$, where $r_k \colon \mathbb{R} \to [0, \infty]$ is the Cramér transform of component Z_k^l, and so satisfies $r_k(\mu_k) = 0$.

Example 2.1. Suppose that Z^l has a multivariate normal distribution with mean vector $\mathbf{0}$ and covariance matrix $\sigma^2 I$. Then a direct calculation shows that the Cramér transform of component Z_k^l is $r_k(z_k) = \frac{(z_k)^2}{2\sigma^2}$, implying that the Cramér transform of Z^l itself is $R(z) = \sum_{k=1}^n \frac{(z_k)^2}{2\sigma^2}$. ◆

Example 2.2. Suppose that the components of Z^l are independent, each with an *exponential*(λ) distribution. Then $r_k(z_k) = \lambda z_k - 1 - \log \lambda z_k$ when $z_k > 0$ and $r_k(z_k) = \infty$ otherwise, implying that $R(z) = \sum_{k=1}^n (\lambda z_k - 1 - \log \lambda z_k)$ when $z \in \mathbb{R}^n_{++}$ and that $R(z) = \infty$ otherwise. ◆

Cramér's Theorem states that

$$(1) \qquad -\lim_{m \to \infty} \tfrac{1}{m} \log \mathbb{P}(\bar{Z}^m \in U) = \inf_{z \in U} R(z)$$

whenever $U \subseteq \mathbb{R}^d$ is a *continuity set* of R, meaning that $\inf_{z \in \mathrm{int}(U)} R(z) = \inf_{z \in \mathrm{cl}(U)} R(z)$. Roughly speaking, equation (1) says that the probability that $\bar{Z}^m$ takes a value in U is of order $\exp(-m\, R(z^*))$ (that is, that the exponential rate of decay of $\mathbb{P}(\bar{Z}^m \in U)$ is $R(z^*)$), where z^* minimizes the rate function R on the set U. If after a large number of trials the realization of $\bar{Z}^m$ is in U, it is overwhelmingly likely that this realization is one that achieves as small a value of R as possible given this constraint; thus, the rate of decay is determined by this smallest value.

2.2 Discrete Choice and Unlikelihood Functions

Consider an agent who must choose among a set of n actions. The payoff to action i is the sum of the fixed base payoff π_i and the random shock $\bar{Z}_i^m$, which is itself the average of the m random variables $\{Z_i^l\}_{l=1}^m$. The agent chooses the action that is optimal ex post.

[2] These properties of the Cramér transform and Cramér's Theorem can be found in Section 2.2 of Dembo and Zeitouni (1998). In particular, the finiteness, strict convexity, and smoothness of R on the interior of its domain follow from the assumptions that Z^l has convex support and that its moment generating function exists—see Exercises 2.2.24 and 2.2.39 in Dembo and Zeitouni (1998).

The probabilities with which the agent chooses each action are described by the *choice probability function* $C^m \colon \mathbb{R}^n \to \mathbb{R}^n_+$, defined by

$$C_i^m(\pi) = \mathbb{P}\left(\bigcap_{j=1}^{n} \left\{\pi_i + \bar{Z}_i^m \geq \pi_j + \bar{Z}_j^m\right\}\right) = \mathbb{P}\left(D^i(\pi + \bar{Z}^m) \geq \mathbf{0}\right).$$

In the last expression, $D^i \in \mathbb{R}^{n\times n}$ is the matrix $\mathbf{1}e_i' - I$, where e_i is the ith standard basis vector and $\mathbf{1}$ the vector of ones, so that $(D^i\pi)_j = \pi_i - \pi_j$.

Define the *unlikelihood function* $\Upsilon \colon \mathbb{R}^n \to \mathbb{R}^n_+$ by

$$\Upsilon_i(\pi) = -\lim_{m\to\infty} \tfrac{1}{m} \log C_i^m(\pi). \tag{2}$$

In rough terms, equation (2) says that $C_i^m(\pi)$ is of order $\exp(-m\Upsilon_i(\pi))$. Thus, $\Upsilon_i(\pi)$ is the exponential rate of decay of the choice probability $C_i^m(\pi)$ as m grows large.

By Cramér's Theorem, the unlikelihood $\Upsilon_i(\pi)$ can be computed as

$$\Upsilon_i(\pi) = \min \sum_{k=1}^{n} r_k(z_k) \quad \text{subject to } \; D^i(\pi + z) \geq \mathbf{0}. \tag{3}$$

Proposition 2.3 uses (3) to derive some basic qualitative properties of the unlikelihood function, and Proposition 2.4 provides a tractable characterization.

Proposition 2.3. (*i*) $\Upsilon_i(\pi) = 0$ *if and only if* $\pi_i + \mu_i \geq \pi_j + \mu_j$ *for all* $j \neq i$.
(*ii*) $\Upsilon_i(\pi)$ *is nonincreasing in* π_i *and is nondecreasing in* π_j *for* $j \neq i$.
(*iii*) $\Upsilon_i(\pi)$ *is convex in* π.

Proof. Parts (i) and (ii) are immediate. Since the objective function in program (3) is convex, and since the function defining the program's constraints is linear in the vector $(z, \pi) \in \mathbb{R}^{2n}$, part (iii) follows from Mangasarian and Rosen (1964, Lemma 1). ■

Proposition 2.4. *Suppose that* $C_i^1(\pi) > 0$. *Then the unlikelihood function* Υ *satisfies*

$$\Upsilon_i(\pi) = \sum_{k=1}^{n} r_k(z_k^*), \tag{4}$$

where

$$z_j^* = \zeta_j(z_i^*) \equiv (z_i^* + \pi_i - \pi_j) \wedge \mu_j \;\; \textit{for } j \neq i, \tag{5}$$

and where z_i^* is the unique solution to

$$(6) \qquad r_i'(z_i^*) + \sum_{j \neq i} r_j'\left(\zeta_j(z_i^*)\right) = 0,$$

Proof. In the Appendix.

In Proposition 2.4, the vector z^* represents the realization of the average shock vector $\bar{Z}^m$ that is "least unlikely" among those that make action i optimal. To explain the form that z^* takes, it is convenient to focus on the case in which each component Z_k^l of the shock vector Z^l has mean $\mu_k = 0$.[3] In this case, the proposition implies that if action i was not optimal ex ante, then the shock z_i^* must be positive, the shocks to worse-performing actions must be zero, and the shock to each better-performing action j may be either zero or negative, according to whether or not z_i^* is large enough to compensate for the base payoff deficit $\pi_j - \pi_i$. If it is not, the negative payoff shock z_j^* ensures that i and j have the same ex post payoff. Finally, the positive shock value z_i^* is chosen so that the marginal reduction in unlikelihood that would result from lowering z_i^* is exactly offset by the marginal increases in unlikelihood that would result from lowering the negative values of z_j^* by the same amount.

2.3 The Multinomial Probit Model

Because the average of m independent normal random variables is itself normally distributed, our discrete choice model includes MW's multinomial probit model as a special case. Indeed, because the Cramér transform for a $N(0, \sigma^2)$ random variable is the quadratic function $r_k(z_k) = \frac{(z_k)^2}{2\sigma^2}$, the vector z^* from Proposition 2.4 takes a particularly simple form: since $r_k'(z_k) = \frac{z_k}{\sigma^2}$, the first-order condition (6) requires the components of the shock vector z^* to have arithmetic mean zero.

This fact and the considerations described after Proposition 2.4 lead to a simple characterization of the unlikelihood function of the multinomial probit model. To present it most concisely we introduce a new definition: for any set $K \subseteq S$ of cardinality n^K, we let

$$\bar{\pi}^K = \frac{1}{n^K} \sum_{k \in K} \pi_k$$

denote the average payoff of the actions in K.

[3]This is without loss of generality, since one can always eliminate a nonzero mean μ_k by replacing component Z_k^l with it with its demeaned version $Z_k^l - \mu_k$ and replacing the base payoff π_k with $\pi_k + \mu_k$.

Proposition 2.5. *Suppose that each random vector $Z^l \sim N(\mathbf{0}, \sigma^2 I)$ is multivariate normal with i.i.d. components, so that the random vector $\bar{Z}^m \sim N(\mathbf{0}, \frac{\sigma^2}{m} I)$ is multivariate normal with i.i.d. components as well. Then the unlikelihood function Υ is given by*

$$\Upsilon_i(\pi) = \sum_{k=1}^{n} \frac{(z_k^*)^2}{2\sigma^2}, \text{ where} \tag{7}$$

$$z_j^* = \begin{cases} \bar{\pi}^{J\cup\{i\}} - \pi_j & \text{if } j \in J \cup \{i\}, \\ 0 & \text{otherwise,} \end{cases} \tag{8}$$

with the set $J \subset S - \{i\}$ being uniquely determined by the requirement that

$$j \in J \text{ if and only if } \pi_j > \bar{\pi}^{J\cup\{i\}}. \tag{9}$$

Thus J is the set of actions with $z_j^* < 0$.

Proof. In the Appendix.

MW analyze the rates of decay of choice probabilities in the multinomial probit model by directly evaluating the limit of the relevant multiple integral. Their Proposition 1 states that these rates take the form described in equations (7) and (8) above, but with the set $J \cup \{i\}$ being replaced by the set of all actions whose base payoffs are at least π_i. In contrast, Proposition 2.5 requires J to contain only those actions whose payoffs are sufficiently larger than π_i to make a positive contribution to the average payoff of actions in $J \cup \{i\}$. Among other things, this ensures that equation (8) does not assign any action other than i a positive payoff shock.[4]

We illustrate these points through a simple example.

Example 2.6. Let $n = 3$, suppose that payoff shocks are i.i.d. standard normal, and consider a base payoff vector of $\pi = (\pi_1, \pi_2, \pi_3) = (0, b, c)$ with $b > 0$. If $c \leq 0$, so that only action 2's base payoff is higher than action 1's, then both MW's Proposition 1 and our Proposition 2.5 specify the unlikelihood of choosing action 1 as $\Upsilon_1(\pi) = \frac{b^2}{4}$, obtained from shock vector $z^* = (\frac{b}{2}, -\frac{b}{2}, 0)$. Our large deviations analysis shows that the least unlikely way to satisfy the inequality $\bar{Z}_1^m - \bar{Z}_2^m \geq b$ is to have the shocks to actions 1 and 2 "share the burden equally".

Now suppose instead that $c > 0$, so that actions 2 and 3 both have higher base payoffs than action 1. In this case, MW's Proposition 1 suggests that the unlikelihood of choosing

[4]The error in MW's analysis seems to occur on p. 297. It is claimed that the integral in equation (A.2) vanishes because its integrand vanishes, but the requirements of the dominated convergence theorem are not verified.

action 1 is $\Upsilon_1(\pi) = \frac{1}{3}(b^2 - bc + c^2)$, obtained from shock vector $z = (\frac{b+c}{3}, \frac{c-2b}{3}, \frac{b-2c}{3})$. But if $c < \frac{b}{2}$, then the base payoff deficit of action 1 relative to action 3, $\pi_1 - \pi_3 = -c$, is already fully addressed by the positive shock to action 1, $z_1 = \frac{b+c}{3} > c$. Indeed, the shock to action 3 specified above, $z_3 = \frac{b-2c}{3}$, is positive, which can only be counterproductive.

In fact, when $c < \frac{b}{2}$, Proposition 2.5 tells us that the optimal choice of z is still $z^* = (\frac{b}{2}, -\frac{b}{2}, 0)$, for an unlikelihood of $\Upsilon_1(\pi) = \frac{b^2}{4}$. More generally, Proposition 2.5 shows that when b and c are positive,

$$z^* = \begin{cases} (\frac{b}{2}, -\frac{b}{2}, 0) & \text{if } c < \frac{b}{2}, \\ (\frac{b+c}{3}, \frac{c-2b}{3}, \frac{b-2c}{3}) & \text{if } c \in [\frac{b}{2}, 2b], \\ (\frac{c}{2}, 0, -\frac{c}{2}) & \text{if } c > 2b, \end{cases} \quad \text{and} \quad \Upsilon_1(\pi) = \begin{cases} \frac{b^2}{4} & \text{if } c < \frac{b}{2}, \\ \frac{b^2-bc+c^2}{3} & \text{if } c \in [\frac{b}{2}, 2b], \\ \frac{c^2}{4} & \text{if } c > 2b. \quad \blacklozenge \end{cases}$$

2.4 Exponentially Distributed Payoff Shocks

We have seen that when payoff shocks are normally distributed, the components of the shock vector z^* have arithmetic mean equal to $\mu_k = 0$, so that the positive payoff shock to strategy i is equal in absolute value to the sum of the negative payoff shocks to strategies with sufficiently higher base payoffs. For instance, when $n = 2$ and $\pi_j > \pi_i$, the z^* used to determine $\Upsilon_i(\pi)$ is given by $z_i^* = \frac{\pi_j - \pi_i}{2} = -z_j^*$.

If instead the payoff shocks follow an *exponential*(λ) distribution, then the fact that this distribution has no left tail suggests that the shock vector z^* should take a less symmetric form. Indeed, since the relevant Cramér transform is $r_k(z_k) = \lambda z_k - 1 - \log \lambda z_k$, realizations of $\bar{Z}_i^m$ that are significantly above the mean shock λ^{-1} (which thus have $r_i(z_i) \approx \lambda z_i - 1$) are far less uncommon than realizations of $\bar{Z}_j^m$ that are below λ^{-1} to a similar extent (and which thus have $r_j(z_j) \approx -\log \lambda z_j$); of course, negative realizations of $\bar{Z}_j^m$ are impossible.

Proposition 2.7 shows that the correct asymmetric treatment of above- and below-average shocks can be expressed in a surprisingly simple form: with exponential payoff shocks, the *harmonic mean* of the components of z^*,[5]

$$\mathcal{H}(z^*) = \frac{n}{\sum_{k=1}^{n} \frac{1}{z_k^*}},$$

must be equated to mean payoff shock $\mu_k = \lambda^{-1}$.

Proposition 2.7. *Suppose that components of the random vector Z^l are independent, each with*

[5]For interpretation, recall that if $z_1, z_2, \ldots, z_n$ are viewed as the average speeds at which a fixed distance is traversed during n distinct journeys, then $\mathcal{H}(z)$ represents the overall average speed over the n journeys.

an exponential(λ) distribution. Then the unlikelihood function Υ is given by

$$\Upsilon_i(\pi) = \sum_{k=1}^{n} \left(\lambda z_k^* - 1 - \log \lambda z_k^*\right),$$

where

$$z_j^* = \zeta_j(z_i^*) \equiv (z_i^* + \pi_i - \pi_j) \wedge \lambda^{-1} \;\; \textit{for } j \neq i,$$

and where z_i^* is uniquely defined by the requirement that

$$\mathcal{H}(z_i^*, \ldots, \zeta_j(z_i^*), \ldots) = \lambda^{-1}. \tag{10}$$

Proof. Since $r_k(z_k) = \lambda z_k - 1 - \log \lambda z_k$ by Example 2.2, $r_k'(z_k) = \lambda - \frac{1}{z_k}$. Thus, equations (5) and (6) imply that z^* satisfies $n\lambda = \sum_{k=1}^{n} \frac{1}{z_k^*}$. Rearranging this equation and applying Proposition 2.4 proves the result. ■

Propositions 2.5 and 2.7 reveal that the unlikelihood functions for the probit and exponential noise models differ in two important respects. The discussion above emphasizes the symmetry and asymmetry of the shock vectors z^*. It is at least as important that in the probit case, the Cramér transform $R(z)$ is quadratic in z, while in the exponential case, $R(z)$ grows linearly in the positive components of z. This difference reflects the fact that the right tail of the exponential distribution is fatter than the tails of the normal distribution. It implies that the probabilities of suboptimal choices tend to decay much more slowly in the size of the sample when payoff shocks are exponentially distributed rather than normally distributed.

A. Appendix

Proof of Proposition 2.4.

We begin with a lemma.

Lemma A.1. *The optimal solution to program* (3), $z^* \in \mathbb{R}^n$, *is the unique vector satisfying*

$$\sum_{k=1}^{n} r_k'(z_k^*) = 0, \tag{11}$$

$$\pi_i + z_i^* - (\pi_j + z_j^*) \geq 0 \quad \textit{for all } j \neq i, \tag{12}$$

$$z_j^* - \mu_j \leq 0 \quad \textit{for all } j \neq i, \textit{ and} \tag{13}$$

$$(14) \qquad \left(z_j^* - \mu_j\right)\left(\pi_i + z_i^* - (\pi_j + z_j^*)\right) = 0 \quad \textit{for all } j \neq i.$$

The proof of Proposition 2.4 follows easily from this lemma. Conditions (12) and (13) are equivalent to the requirement that $z_j^* \leq (z_i^* + \pi_i - \pi_j) \wedge \mu_j$ for $j \neq i$, and introducing condition (14) is equivalent to requiring the inequality to always bind, yielding condition (5). Given condition (5), equation (11) is equivalent to equation (6).

Proof of Lemma A.1. Since $C_i^1(\pi) > 0$, program (3) admits a feasible solution on the interior of the support of Z^l. Since the Cramér transform of Z^l is differentiable in this region, we can solve program (3) using the Kuhn-Tucker method.

The Lagrangian for program (3) is

$$\mathcal{L}(z, \nu) = \sum_{k=1}^{n} r_k(z_k) - \sum_{j \neq i} \nu_j \left(\pi_i + z_i - (\pi_j + z_j)\right).$$

Since the objective function is convex, z^* is the minimizer if and only it satisfies the constraints (12) and there exist Lagrange multipliers ν^* such that z^* and ν^* together satisfy

$$(15) \qquad r_i'(z_i^*) = \sum_{j \neq i} \nu_j^*,$$

$$(16) \qquad r_j'(z_j^*) = -\nu_j^* \qquad \text{for all } j \neq i,$$

$$(17) \qquad \nu_j^* \geq 0 \qquad \text{for all } j \neq i, \text{ and}$$

$$(18) \qquad \nu_j^* \left(\pi_i + z_i^* - (\pi_j + z_j^*)\right) = 0 \qquad \text{for all } j \neq i.$$

Conditions (15) and (16) together imply condition (11). Since each r_j is strictly convex on its domain and is minimized at μ_j, r_j' satisfies $\operatorname{sgn}(r_j'(z_j)) = \operatorname{sgn}(z_j - \mu_j)$. Thus, conditions (16) and (17) imply condition (13), and conditions (16) and (18) imply condition (14). ∎

Proof of Proposition 2.5.

We apply Proposition 2.4, using the Cramér transform for $N(0, \sigma^2)$ random variables, $r_k(z_k) = \frac{(z_k)^2}{2\sigma^2}$, introduced in Example 2.1. Evidently, equation (4) becomes equation (7), and since $r_k'(z_k) = \frac{z_k}{\sigma^2}$, equation (6) becomes

$$z_i^* + \sum_{j \neq i} \left((z_i^* + \pi_i - \pi_j) \wedge 0\right) = 0.$$

If we define $J^* = \{j \in S\colon \pi_j > \pi_i + z_i^*\}$ and denote this set's cardinality by n^*, we can rewrite the previous equation as

$$(n^* + 1)\, z_i^* = \sum_{j \in J^*} (\pi_j - \pi_i),$$

and hence as

$$z_i^* = \frac{1}{n^* + 1} \sum_{j \in J^* \cup \{i\}} (\pi_j - \pi_i) = \bar{\pi}^{J^* \cup \{i\}} - \pi_i.$$

Thus $J^* = \{j \in S \colon \pi_j > \bar{\pi}^{J^* \cup \{i\}}\}$, so $J^* = J$ by condition (9). Equation (5) thus becomes

$$z_j^* = (z_i^* + \pi_i - \pi_j) \wedge 0 = (\bar{\pi}^{J \cup \{i\}} - \pi_j) \wedge 0,$$

which is equation (8). This completes the proof of the proposition. ■

References

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424.

Dembo, A. and Zeitouni, O. (1998). *Large Deviations Techniques and Applications*. Springer, New York, second edition.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.

Mangasarian, O. L. and Rosen, J. B. (1964). Inequalities for stochastic nonlinear programming problems. *Operations Research*, 12:143–154.

Myatt, D. P. and Wallace, C. C. (2003). A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113:286–301.

Ui, T. (1998). Robustness of stochastic stability. Unpublished manuscript, Bank of Japan.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.