Web Appendix to

"Minimax Regret Treatment Choice with Finite Samples"

Jörg Stoye

New York University

January 22, 2009

This appendix uses finite sample results from Stoye (2009) to compute exact minimax regret values for different decision situations and sample designs. The goal is to improve on several quantitative analyses in Manski (2004).

Performance of Empirical Success Rules

The empirical success rule, $\delta^{ES} \equiv \mathbb{I}\{\overline{y}_1 \geq \overline{y}_0\}$, is of special interest for at least three reasons. First, it is probably the most obvious decision rule to employ. Second, it is the one used by Manski (2004). Third, we know from Hirano and Porter (2008) that it is asymptotically minimax regret efficient. Yet at the same time, it was seen to significantly differ from the minimax regret rule for small samples.

How much do these differences matter in terms of regret? By searching over a restricted state space, one can numerically bound from below $\sup_{s \in \mathcal{S}} R(\delta^{ES}, s)$. In table 1, the resulting bound, labelled $\underline{R}$, is contrasted with R_1^* (from corollary 1 in Stoye 2009) for different sample sizes.[1] δ^{ES} incurs significant excess regret when samples are very small, incurring at least double the true minimax regret for $N = 2$. But the table also suggests convergence: For $N = 40$, the demonstrable inefficiency of δ^{ES} drops to 21% of true minimax regret, and this percentage becomes very small for large samples.

A similar exercise can be performed for the case of testing an innovation. Tables 2 through 4 compare the empirical success rule, here $\delta^{ES} = \mathbb{I}(\overline{y}_1 \geq \mu_0)$, to $\widetilde{\delta}_3^*$, using the same three values of μ_0 as Manski (2004).[2] The relative underperformance, in terms of regret, of δ^{ES} is again quite large for

[1] Although I continue to define δ^{ES} as in Manski (2004), results are not driven by suboptimal tie-breaking. The source of suboptimality of δ^{ES} is its discontinuity at $\overline{y}_1 = \overline{y}_0$. Conversely, what sets $\widetilde{\delta}_1^*$ apart is the smoothing (by means of randomization) around $\overline{y}_1 = \overline{y}_0$.

[2] In these tables, the lower bound $\underline{R}$ sometimes displays erratic behavior, e.g. it increases in N. This may occasionally reflect movement in the true regret (minimax regret must weakly decrease in N, but the maximal regret incurred by a specific decision rule need not do so), but probably also attests to the difficulty of finding $\sup_{s \in \mathcal{S}} R(\delta^{ES}, s)$.

N	**2**	**4**	**6**	**8**	**10**	**12**	**14**	**16**	**18**	**20**
R_1^*	.1250	.0870	.0706	.0609	.0543	.0495	.0458	.0428	.0403	.0382
$\underline{R}$	.25	.1481	.1066	.0883	.0765	.0681	.0617	.0568	.0528	.0495
N	**22**	**24**	**30**	**40**	**50**	**100**	**200**	**400**	**1000**	**2000**
R_1^*	.0364	.0348	.0311	.0269	.0241	.0170	.0120	.0085	.0054	.0038
$\underline{R}$	.0467	.0442	.0387	.0326	.0286	.0193	.0132	.0091	.0056	.0039

Table 1: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (two unknown treatments).

N	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R_3^*	.0408	.0358	.0315	.0279	.0248	.0221	.0197	.0176	.0158	.0143
$\underline{R}$	.2256	.1270	.0859	.0634	.0492	.0396	.0326	.0273	.0232	.0199
N	**11**	**12**	**15**	**20**	**25**	**50**	**100**	**200**	**500**	**1000**
R_3^*	.0130	.0120	.0096	.0086	.0080	.0054	.0037	.0026	.0017	.0012
$\underline{R}$	.0193	.0195	.0199	.0203	.0136	.0084	.0059	.0037	.0021	.0014

Table 2: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (one unknown treatment, mu0=0.05).

N	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R_3^*	.0900	.0516	.0389	.0380	.0345	.0299	.0268	.0265	.0252	.0232
$\underline{R}$	.1406	.0625	.0705	.0790	.0517	.0431	.0455	.0473	.0353	.0336
N	**11**	**12**	**15**	**20**	**25**	**50**	**100**	**200**	**500**	**1000**
R_3^*	.0217	.0215	.0187	.0166	.0149	.0104	.0074	.0052	.0033	.0023
$\underline{R}$	.0346	.0354	.0285	.0248	.0193	.0132	.0090	.0060	.0036	.0025

Table 3: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (one unknown treatment, mu0=0.25).

N	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
R_3^*	.0625	.0625	.0435	.0435	.0353	.0353	.0304	.0304	.0272	.0272
$\underline{R}$	.1250	.1099	.0531	.0697	.0462	.0533	.0401	.0441	.0356	.0382
N	**11**	**12**	**15**	**20**	**25**	**50**	**100**	**200**	**500**	**1000**
R_3^*	.0247	.0247	.0214	.0191	.0167	.0120	.0085	.0060	.0038	.0027
$\underline{R}$	.0321	.0340	.0273	.0247	.0206	.0143	.0096	.0066	.0040	.0028

Table 4: Selected values of R* and of a lower bound for regret incurred by the empirical success rule (one unknown treatment, mu0=0.5).

small and moderate sample sizes.

Comparing Sample Stratifications

Manski (2004) also compares the minimax regret value of different ways to stratify samples by covariate. I remove two layers of approximation from this analysis. First, computations are based on exact regret and not an upper bound on it; second, they presume that conditional on sample designs, exact minimax regret rules are chosen, whereas Manski restricts attention to δ^{ES}. Table 5 displays the resulting optimal stratifications, as well as the minimax regret value of the optimal stratification, for the case of a binary covariate $X \in \{m, f\}$. (The tools employed here are corollary 1 and example 1 in Stoye (2009); note that the example is exactly the problem analyzed here.) The cells corresponding to given values of N and $\Pr(X = m)$ give the optimal stratification (N_m, N_f) as well as the minimax regret value $R^* \equiv \Pr(X = m)R_1^*(N_m) + \Pr(X = f)R_1^*(N_f)$ of this stratification. The numbers can be compared to Manski's (2004, table II) "quasi-optimal stratifications" as well as his upper bounds on the resulting regret, all of which are reproduced in parentheses.[3]

The entries for R^* reveal that previous bounds on minimax regret had considerable slack. (Its two sources – suboptimal performance of δ^{ES} and slack of large deviations bounds – are not separated here.) Accordingly, optimal stratifications frequently differ from the quasi-optimal ones, although not by very much. Furthermore, this paper's closed-form results allow for very fast computation of the solutions, so that the table can be extended to much larger sample sizes.

Explanation of Tables

I conclude by giving some remarks on the computation of $\underline{R}$. The problem is to bound the maximal regret incurred by δ^{ES},

$$\sup_{P(Y_0,Y_1)\in\Delta[0,1]^2} R(\delta^{ES}, P(Y_0,Y_1)) = \sup_{P(Y_0,Y_1)\in\Delta[0,1]^2} \left\{ (\mu_1 - \mu_0) \left[\Pr\left(\overline{y}_0 > \overline{y}_1\right) + \frac{1}{2}\Pr(\overline{y}_0 = \overline{y}_1) \right] \right\},$$

where I assume that $\mu_1 > \mu_0$; this is w.l.o.g. by arguments in the proof of proposition 1. I here specify the tie-breaking probability as 1/2 to clarify that results are not driven by suboptimal tie-breaking.

A full treatment of this problem is intricate, but its value can be bounded from below by searching over a restriction of $\Delta[0,1]^2$. I consider the following cases:

- **Case 1:** Both treatments induce Bernoulli distributions, i.e. the search space is restricted to $\Delta\{0,1\}^2$; similar to previous arguments, the search parameters are then just (μ_0, μ_1).

[3] Like Manski (2004), I restrict attention to deterministic stratifications. The truly optimal stratifications may be randomized. They could be computed from the extension of corollary 1 to random N.

	$\Pr(\mathbf{X} = \mathbf{m}) = \mathbf{.05}$			$\Pr(\mathbf{X} = \mathbf{m}) = \mathbf{.25}$		
N	$\mathbf{N}_m$	$\mathbf{N}_f$	$\mathbf{R}^*$	$\mathbf{N}_m$	$\mathbf{N}_f$	$\mathbf{R}^*$
4	0(0)	4(2)	.108(.338)	2(2)	2(1)	.125(.423)
8	2(0)	6(4)	.073(.250)	2(2)	6(3)	.084(.293)
12	2(2)	10(5)	.058(.203)	4(4)	8(4)	.067(.234)
16	2(2)	14(7)	.050(.173)	6(6)	10(5)	.058(.205)
20	2(4)	18(8)	.045(.154)	6(6)	14(7)	.052(.182)
24	4(4)	20(10)	.041(.143)	8(8)	16(8)	.047(.162)
28	4(4)	24(12)	.037(.133)	10(10)	18(9)	.044(.153)
32	4(4)	28(14)	.035(.124)	10(12)	22(10)	.041(.144)
36	4(4)	32(16)	.033(.115)	12(12)	24(12)	.038(.137)
40	6(4)	34(18)	.031(.108)	14(14)	26(13)	.037(.129)
44	6(4)	38(20)	.030(.101)	14(16)	30(14)	.035(.122)
48	6(4)	42(22)	.029(.094)	16(16)	32(16)	.033(.116)
52	6(6)	46(23)	.027(.088)	16(16)	36(18)	.032(.110)
60	8	52	.025	20	40	.030
80	10	70	.022	26	54	.026
100	12	88	.020	32	68	.023
200	24	176	.014	64	136	.016
500	62	438	.009	162	338	.010
1000	124	876	.006	326	674	.007

Table 5: Optimal (deterministic) sample stratifications; values from Manski (2004) in parentheses.

- **Case 2:** Treatment 0 induces a degenerate distribution concentrated at some point μ_0, whereas treatment 1 has a Bernoulli distribution. This case has two subcases according as $\mu_1 > [<]\mu_0$.

Both cases incur open set problems. For case 1, if $\overline{y}_0 = \overline{y}_1$, the decision rule will assign the correct treatment half the time. Regret can, therefore, be increased by letting the better of the two treatments be supported (with unchanged probabilities) on $\{0, 1-\varepsilon\}$ rather than $\{0, 1\}$. By letting $\varepsilon \to 0$, one can approximate the effect of a decision rule whose tie-breaking goes in the wrong direction. The problem with case 2 is similar. To generate well-behaved problems, I therefore rig the tie-breaking rule against the decision maker in all cases.

With these remarks in mind, alternative lower bounds on the regret can be computed as follows. Define $B(n, N, \mu) \equiv \binom{N}{n}\mu^n(1-\mu)^{N-n}$ and $F(n, N, \mu) \equiv \sum_{i=0}^{n} B(i, N, \mu)$, then one can consider the following cases:

- **Case 1:** $\max_{\mu_0,\mu_1\in[0,1]} \left\{(\mu_1 - \mu_0) \sum_{n=0}^{N} B(n, N, \mu_0) \cdot F(n, N, \mu_1)\right\}$.
- **Case 2, first subcase:** $\max_{\mu_0,\mu_1\in[0,1]} \{(\mu_1 - \mu_0)\, F(\mu_0 N, N, \mu_1)\}$.
- **Case 2, second subcase:** $\max_{\mu_0,\mu_1\in[0,1]} \{(\mu_0 - \mu_1)\,(1 - F(\mu_0 N - 1, N, \mu_1)\}$.

Table 1 is generated by numerical evaluation of all of these. The two subcases of case 2 turn out to yield identical regrets. Cases 1 and 2 coincide for $N = 1$. Otherwise, case 2 binds, i.e. yields the highest regret, for $N = 2$, and case 1 binds thereafter.

Tables 2-4 only use case 1, but with the variation that the support of Y_1 is generalized to $\{0, x\}$. The idea here is the following: If one evaluates case 1 only, then $\underline{R}$ sometimes sharply increases in N; for example, this occurs if $\mu_0 = 0.25$ and N moves from 3 to 4. In the specific example, this happens because for $N = 4$, treatment 1 will be rejected even if one success is recorded. For any given Bernoulli distributed treatment, the probability of it being adopted therefore jumps downward, which means that for any such treatment with parameter exceeding μ_0 – i.e. one that should, in fact, be adopted –, regret increases.

This observation spawns an intuition: Perhaps for $N = 4$ and $\mu_0 = 0.25$, one might want to consider distributions Y_1 supported on $\{0, 3/4\}$ because in this case, one success in 4 trials will still lead to rejection. Indeed, this is how the according cell of table 5 was found. More generally, I search over some salient guesses of the upper support point and also execute an algorithm in which this variable is handed down to the maximizer.

References

[1] Hirano, K. and J.R. Porter (2008): "Asymptotics for Statistical Treatment Rules," mimeo, University of Arizona.

[2] Manski, C.F. (2004): "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica* 72: 1221-1246.

[3] Stoye, J. (2009): "Minimax Regret Treatment Choice with Finite Samples," mimeo, New York University.