Inverse Decision Theory: Characterizing Losses for a Decision Rule With Applications in Cervical Cancer Screening

Richard J. SWARTZ, Dennis D. COX, Scott B. CANTOR, Kalatu DAVIES, and Michele FOLLEN

Identifying an optimal decision rule using Bayesian decision theory requires priors, likelihoods, and losses. In many medical settings, we can develop priors and likelihoods, but specifying losses can be difficult, especially when considering both patient outcomes and economic costs. If there is a widely accepted treatment strategy, then we can consider the inverse problem and find a region in the space of losses where the procedure is optimal. We call this approach inverse decision theory (IDT). We apply IDT to the standard of care for diagnosis and treatment of precancerous lesions of the cervix, and consider an alternative procedure that has been proposed. We use a Bayesian approach to estimate the probabilities associated with the diagnostic tests and make inferences about the region in loss space where these medical procedures are optimal. In particular, we find evidence supporting the current standard of care.

KEY WORDS: Bayesian decision theory; Cervical intraepithelial neoplasia; Cost-benefit ratio; Medical decision making; Squamous intraepithelial neoplasia.

1. INTRODUCTION

The field of medical decision making (MDM) uses formal, quantitative methods, including decision theory, to analyze treatment strategies. Decision theory requires specifying a loss function and expressing all losses on a common scale. However, in medical settings, it is difficult to estimate economic losses and even more difficult to quantify patient outcomes such as pain or death. There is a considerable literature on quantifying patient outcomes in terms of quality-adjusted life years (QALYs) (Weinstein and Fineberg 1980), but this is controversial and leaves them on a different scale from economic costs (Harris 1987). There is also extensive MDM literature discussing various other methods for determining costs and losses, but these are also controversial (Sox, Blatt, Higgins, and Marton 1988).

Given the difficulty in quantifying losses, we could consider the problem from a different perspective. Often there is a widely accepted treatment strategy known as a *standard of care* (SOC), which is the recommendation of a panel of physicians who have arrived at a consensus (Jordan 1985). This suggests investigating the inverse of a typical decision theory problem: Given a decision rule, can we characterize the loss function under which it would be optimal? We call this *inverse decision theory* (IDT). It is similar to "revealed preferences" in economics (Ben-Akiva and Lerman 1985) and inverse optimization (Ahuja and Orlin 2001). To our knowledge, the results presented here represent the first development of IDT.

IDT has two main complications. The conditions which prescribe optimality of a given decision rule involve quantities that may not be known precisely and must be estimated from data. Also, IDT may give not a single loss, but rather a set of losses. Nevertheless, we may still be able to draw some conclusions about a decision rule, as we will discuss in the example considered here.

In the next section we describe the theory underlying IDT and give a simple example. Then we apply IDT to the SOC for detection and treatment of cervical intraepithelial neoplasia (CIN), precancerous lesions of the uterine cervix. We also consider a competing strategy that has been proposed for the detection and treatment of CIN and compare it with the SOC.

2. INVERSE DECISION THEORY

We assume a finite action space and a finite parameter space. Let θ denote the disease state of the patient, $\mathbf{x}$ be the observed outcome of a sequence of tests (we suppose that the data are collected sequentially), $f(\mathbf{x}|\theta)$ be the likelihood, and $p(\theta)$ be the prior. Consider a loss function, $\Lambda(\theta, a)$, where a is an allowable action (i.e., a treatment for the patient). Let c_i be the cost for stage i in the sequence of diagnostic tests with n total tests. For a decision rule, $\delta(\mathbf{x})$, the Bayes risk, $R(\delta)$, is given by

$$R(\delta) = \sum_{\theta} \sum_{x} \Lambda(\theta, \delta(\mathbf{x})) f(\mathbf{x}|\theta) p(\theta) + \sum_{i \le n} c_i \Pr(\text{test } i \text{ is performed}). \quad (1)$$

Note that $R(\delta)$ will be linear in the costs and losses because we consider only finite, discrete loss functions.

Consider a given decision rule, δ^*. This δ^* is optimal under Bayesian decision theory if and only if $R(\delta^*) \le R(\delta)\ \forall \delta \in D$, where D is the set of allowable decision rules. These resulting linear inequalities determine a subset of costs and losses where δ^* is optimal. We call this subset the *region of optimality* (RO) for δ^*, denoted by $\text{RO}(\delta^*)$. Determining $\text{RO}(\delta^*)$ is the problem of IDT.

An IDT analysis can be used in two ways. First, if we believe that δ^* is optimal, then $\text{RO}(\delta^*)$ will characterize the costs and losses associated with δ^*. In contrast, if we do not know

Richard J. Swartz is Instructor, Department of Biostatistics and Applied Mathematics and Department of Behavioral Science, The University of Texas M. D. Anderson Cancer Center, Houston, TX 77030 (E-mail: *rswartz@mdanderson.org*). Dennis D. Cox is Professor, Department of Statistics, Rice University, Houston, TX 77005 (E-mail: *dcox@stat.rice.edu*). Scott B. Cantor is Associate Professor, Department of Biostatistics and Applied Mathematics, The University of Texas M. D. Anderson Cancer Center, Houston, TX 77030. Kalatu Davies is Vigre Postdoctoral Fellow and Pffeifer Lecturer, Department of Statistics, Rice University, Houston, TX 77005. Michele Follen is Professor, Department of Gynecologic Oncology and Director, Biomedical Engineering Center, The University of Texas M. D. Anderson Cancer Center, Houston, TX 77030, and Professor, Department of Obstetrics and Gynecology, The University of Texas Health Science Center, Houston, TX 77030. This research was supported in part by National Cancer Institute grant 2PO1-CA82710 and in part by a cancer prevention fellowship supported by National Cancer Institute grant R25 CA57730.

Journal of the American Statistical Association
March 2006, Vol. 101, No. 473, Applications and Case Studies
DOI 10.1198/016214505000000998

whether δ^* is optimal, then RO(δ^*) will characterize the costs and losses that would be necessary for δ^* to be the optimal decision rule. These costs and losses can then be evaluated by experts to determine whether they are reasonable operating costs.

Consider a simple example: The patient either has or does not have a disease, denoted by Dis $= 1$ and Dis $= 0$ respectively, and the test yields a single binary observation that is either positive (test $= 1$) or negative (test $= 0$). The allowable actions are to treat (treat $= 1$) or not treat (treat $= 0$). We assume that the treatment is effective if the patient has the disease, and that there are negative consequences for treating patients who do not have the disease,

$$\Lambda(1,1) < \Lambda(1,0), \qquad \Lambda(0,1) > \Lambda(0,0). \tag{2}$$

[Recall that $\Lambda(1,0)$ is the loss for Dis $= 1$ and treat $= 0$.] Without loss of generality, we may consider only the net losses for wrong decisions, denoted by

$$\begin{aligned} \lambda_+ &= \Lambda(1,0) - \Lambda(1,1) \qquad \text{and} \\ \lambda_- &= \Lambda(0,1) - \Lambda(0,0). \end{aligned} \tag{3}$$

We suppose that the test is better than a purely random result (i.e., sensitivity plus specificity is greater than unity),

$$\Pr(\text{test} = 1|\text{Dis} = 1) + \Pr(\text{test} = 0|\text{Dis} = 0) > 1. \tag{4}$$

Let δ^* be the decision rule to treat if test $= 1$ and not treat if test $= 0$. Setting its Bayes risk less than or equal to the Bayes risk of all other decision rules and using the foregoing assumptions yields

$$\frac{\Pr(\text{Dis} = 1|\text{test} = 0)}{\Pr(\text{Dis} = 0|\text{test} = 0)} \le \frac{\lambda_-}{\lambda_+} \le \frac{\Pr(\text{Dis} = 1|\text{test} = 1)}{\Pr(\text{Dis} = 0|\text{test} = 1)}. \tag{5}$$

The expression in the middle is the ratio of the net loss of treating an individual without the disease to the net loss of failing to treat an individual with the disease. In the MDM literature, this quantity is known as the cost–benefit (C/B) ratio and is used to summarize costs and losses (Weinstein and Fineberg 1980; Cantor, Sun, Tortolero-Luna, Richards-Kortum, and Follen 1999). Note that what we call "cost" in this article (the cost of performing a medical test, including time costs, monetary costs, and other costs to the patient) is not the same as the "cost" in the C/B ratio (loss due to overtreating nondiseased individuals). We use the C/B ratio in our later analyses.

By estimating the posterior odds [the first and last expressions in (5)], we have information about the C/B ratio. If we believe that the procedure is optimal, then the bounds will tell us about the true C/B ratio. In contrast, if we are not sure the procedure is optimal, then these bounds might help clinicians decide whether or not the procedure is optimal. For example, if the bounds are very wide, then the procedure might be considered optimal under values of the losses considered reasonable by many people.

One final note: When comparing sequential decision rules, it will be more efficient to work backward from the terminal states of the decision rule, δ^*, and derive inequalities on the conditional risk for decision rules that agree with the δ^* up to that stage, similar to the backwards induction method (Ferguson 1967; DeGroot 1970).

3. A PRACTICAL APPLICATION OF IDT

Here we apply IDT to the diagnosis and treatment of CIN. We first describe two different strategies: the standard of care, referred to as the SOC, and a proposed strategy called "see-and-treat colposcopy," which we call as "see-and-treat." Next we apply IDT to obtain the inequalities that define the RO for both procedures. We use a Bayesian approach to estimate the coefficients in the inequalities.

3.1 Diagnosis and Treatment Strategies for Cervical Intraepithelial Neoplasia

The standard of care (SOC) for diagnosing and treating CIN has been established for many years (Jordan 1985; Follen 2001). The strategy is depicted in Figure 1. The three diagnostic stages are: (1) the Papanicolaou (Pap) smear, (2) colposcopy (visual examination of the cervix under magnification), and (3) biopsy. A diagnosis of normal at any stage results in no treatment. A diagnosis of low-grade squamous intraepithelial lesion (LGSIL) leads to monitoring of the patient with colposcopy every 6 months for 2 years. A diagnosis of high-grade squamous intraepithelial lesion (HGSIL) is treated with an outpatient surgical procedure. A diagnosis of invasive cancer involves additional tests followed by more aggressive treatments.

The second diagnostic strategy, "see-and-treat colposcopy," has been proposed as an alternative to the SOC (Ferris, Hiner, Pfenninger, and Zuber 1996; Cardenas-Turanas, Follen,

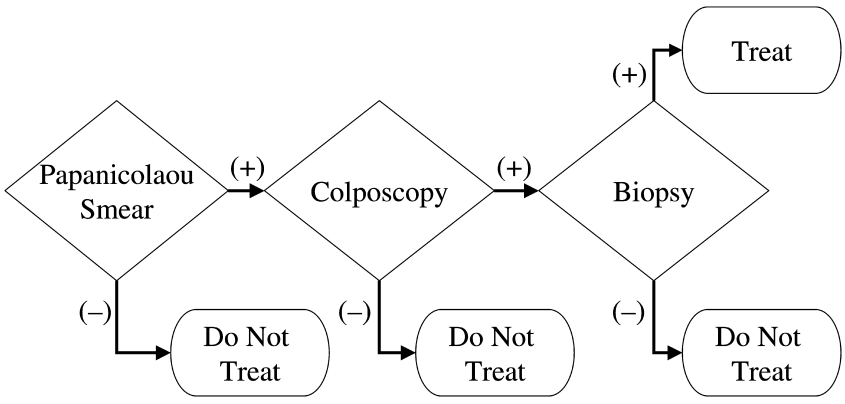

Figure 1. Current Standard of Care for the Screening, Diagnosis, and Treatment of CIN.

Benedet, and Cantor 2005). The "see-and-treat" strategy eliminates the biopsy and treats a patient directly after a positive colposcopic result and is otherwise the same as the SOC.

We make several simplifying assumptions. Because cancer is rare in women who adhere to regular screening with Pap smears, we consider only the population of patients with no worse than HGSIL. We combine a "watchful waiting" clinical management strategy for LGSIL with the nontreatment of women without the disease. In addition, we dichotomize the diagnostic tests according to the positive and negative criteria used by clinicians. The biopsy is assumed to be an error-free test. Furthermore, we consider only decision rules where the sequence of tests is a Pap smear followed by colposcopy followed by biopsy, and the Pap smear is always performed. We discuss these assumptions in Section 4.

3.2 Applying Inverse Decision Theory

Let $\text{Dis} = l$ be disease state l, $\text{Pap} = i$ be Pap smear test result i, $\text{Colpo} = j$ be colposcopy result j, and $\text{Bx} = k$ be biopsy result k, where $i, j, k, l \in \{0 = \text{negative}, 1 = \text{positive}\}$. Let c_c and c_b be the costs of colposcopy and biopsy respectively, λ_+ be the loss from failing to treat a diseased individual, and λ_- be the loss from treating an individual without disease. We assume that these costs and losses are positive and in the same units. We also assume that colposcopy and the Pap smear are better than a random test, as in (4).

First, we investigate $\text{RO}(\delta^{\text{SOC}})$ for the SOC as given in Figure 1. Consider the arm of the decision algorithm where the Pap smear and colposcopy is positive and the patient is now at the final stage: biopsy. Because it is assumed to be error-free, it is optimal to treat patients with positive biopsy and not treat patients with a negative biopsy, so no useful inequalities result.

The conditional expected loss for the SOC at the colposcopy stage is

$$\text{SOC}_1 = c_b \Pr(\text{Colpo} = 1|\text{Pap} = 1) + \lambda_+ \Pr(\text{Dis} = 1, \text{Colpo} = 0|\text{Pap} = 1) + c_c.$$

There are three admissible alternatives to the SOC at this stage that give rise to the following inequalities when we assume that the SOC is optimal:

$$\text{SOC}_1 \le c_b + c_c, \tag{6}$$

$$\text{SOC}_1 \le \lambda_- \Pr(\text{Dis} = 0, \text{Colpo} = 1|\text{Pap} = 1) + \lambda_+ \Pr(\text{Dis} = 1, \text{Colpo} = 0|\text{Pap} = 1) + c_c, \tag{7}$$

and

$$\text{SOC}_1 \le \lambda_- \Pr(\text{Dis} = 0, \text{Colpo} = 1|\text{Pap} = 1) + c_b \Pr(\text{Colpo} = 0|\text{Pap} = 1) + c_c. \tag{8}$$

The right sides of these inequalities correspond to performing biopsies on all patients (6), "see-and-treat colposcopy" (7), and treating after positive colposcopy and performing a biopsy after negative colposcopy (8). There are also three options that are inadmissible. For instance, performing colposcopy but then either treating or not treating everyone are inadmissible rules, because we could have just made the decision after the Pap smear.

Now move back one stage to a positive Pap smear. The options are treat, not treat, or proceed to colposcopy. The latter option has already been considered, and the others yield

$$\text{SOC}_1 \le \lambda_- \Pr(\text{Dis} = 0|\text{Pap} = 1) \tag{9}$$

and

$$\text{SOC}_1 \le \lambda_+ \Pr(\text{Dis} = 1|\text{Pap} = 1). \tag{10}$$

We have now considered all rules conditional on a positive Pap smear.

The SOC for a negative Pap smear is to not treat, and there are five admissible alternatives to the SOC given a negative Pap smear: (1) treat everyone, (2) use the analog of the SOC for a positive Pap smear but with a negative Pap smear, (3) continue to colposcopy and then always to biopsy, (4) continue to the analog of "see-and-treat colposcopy" for a positive Pap smear but with a negative Pap smear, and (5) continue to colposcopy, and from a negative colposcopy to biopsy and a positive colposcopy to treatment. The inequalities can be derived straightforwardly and are not explicitly stated.

The foregoing 10 inequalities along with nonnegativity constraints define the RO associated with the SOC: $\text{RO}(\delta^{\text{SOC}})$. To see this, consider any rule, condition on the Pap smear outcome, and proceed to the last stage where that rule agrees with the SOC. Then compare the conditional risk from after that stage for the SOC and the given rule. This procedure of working backward through the decision tree has led to fewer inequalities than if we had compared the risks with all possible decision rules as in (1). For this setting, there are 6 admissible rules conditional on Pap smear, or 36 rules (of which 10 are inadmissible) unconditional on Pap smear. With our approach, we have only 10 inequalities rather than 35.

Similarly, we can derive inequalities for δ^{st}, "see-and-treat colposcopy," which characterize the region $\text{RO}(\delta^{\text{st}})$ where it is optimal. The right side of (7) is the Bayes risk associated with "see-and-treat" given a positive Pap smear, so we simply set this expression less than or equal to all the other risks conditional on a positive Pap smear [found in (6)–(10)], including the risk for the SOC. Conditional on a negative Pap smear, "see-and-treat" is the same as the SOC, so the same five inequalities apply. We do not explicitly give the linear constraints that define $\text{RO}(\delta^{\text{st}})$.

Note that the coefficients of the costs and losses are functions of the sensitivities and specificities of the tests and the prevalences, which are not precisely known. Also note that because we consider only rules starting with the Pap smear, its cost does not appear in any of the inequalities. It is convenient to set the cost of colposcopy equal to 1,

$$c_c = 1. \tag{11}$$

Thus all costs and losses are in "colposcopy" units. Even though this will not get us absolute costs—to do that, we still would need to express the cost of colposcopy in, for example, dollars—it will allow us to make comparisons and make some inferences on the C/B ratio ($\frac{\lambda_-}{\lambda_+} = \frac{\lambda_-^*}{\lambda_+^*}$); see Section 3.4. With this convention, $\text{RO}(\delta^{\text{SOC}})$ and $\text{RO}(\delta^{\text{st}})$ are subsets of the three-dimensional space of $(\lambda_+^*, \lambda_-^*, c_b^*)$. The superscript "$*$" indicates that these are values relative to the cost of colposcopy.

3.3 Data and Model for Estimating the Coefficients

We used a Bayesian approach to estimate the coefficients appearing in the inequalities discussed in Section 3.2. This allowed us to incorporate available prior information and to propagate uncertainty in the coefficients into uncertainty about $\mathrm{RO}(\delta^{\mathrm{SOC}})$ and $\mathrm{RO}(\delta^{\mathrm{st}})$. The data used in our analysis were collected as part of an ongoing program project funded by the National Cancer Institute, "Optical Technologies for Cervical Neoplasia." It includes subjects from a screening population (patients being seen for a Pap smear) and a diagnostic population (patients who have had a positive Pap smear and have been referred for colposcopy). The data are summarized in Table 1. An observation is a triplet in the form (biopsy, Pap smear, colposcopy), where the true state of disease is determined by the biopsy result because it is a gold-standard test. Negative and positive results are denoted by 0 and 1. The screening or diagnostic population is denoted by Popn = 0 or 1 respectively.

We assume independent random samples from each population, so the likelihood is a product of multinomials, and each sample can be considered a random draw from a multinomial distribution. Our parameters of interest are the various probabilities for all possible combinations of the outcomes of each of the three tests from each of the two populations: $\pi_{p.ijk} = \Pr(\mathrm{Dis} = i, \mathrm{Pap} = j, \mathrm{Colpo} = k|\mathrm{Popn} = p)$ for $p, j, k, l \in \{0, 1\}$, with cell counts denoted by $x_{p.ijk} = \#(\mathrm{Dis} = i, \mathrm{Pap} = j, \mathrm{Colpo} = k|\mathrm{Popn} = p)$. Remember that the disease state is identical to the (dichotomized) biopsy result. The likelihood is the multinomial distribution,

$$f(\mathbf{x}|\pi) = \prod_{p\in\{0,1\}} n_p! \prod_{i,j,k,\in\{0,1\}} \frac{\pi_{p.ijk}^{x_{p.ijk}}}{x_{p.ijk}!}, \tag{12}$$

where π is the column vector of all the $\pi_{p.ijk}$ parameters and $\mathbf{x}$ is the column vector of all the $x_{p.ijk}$ counts. Also, we are subject to the restrictions that $\sum_{i,j,k\in\{0,1\}} \pi_{p.ijk} = 1$.

Because medical tests are characterized in terms of their sensitivity and specificity, it is convenient to parameterize in terms of the prevalence of the disease, sensitivities, and specificities, as

$$q_{p.i} = \Pr(\mathrm{Dis} = i|\mathrm{Popn} = p) = \sum_{j,k\in\{0,1\}} \pi_{p.ijk}, \tag{13}$$

$$q_{p.ij} = \Pr(\mathrm{Pap} = j|\mathrm{Popn} = p, \mathrm{Dis} = i) = \sum_{k\in\{0,1\}} \frac{\pi_{p.ijk}}{q_{p.i}}, \tag{14}$$

Table 1. Cell Counts for Screening and Diagnostic Patients

			Colposcopy	
Population	*Disease state*	*Pap*	*Negative (0)*	*Positive (1)*
Screening (0) [747]	Absent (0) [740]	Negative (0)	640	73
		Positive (1)	15	12
	Present (1) [7]	Negative (0)	1	0
		Positive (1)	4	2
Diagnostic (1) [410]	Absent (0) [287]	Negative (0)	0	0
		Positive (1)	77	210
	Present (1) [123]	Negative (0)	0	0
		Positive (1)	2	121

NOTE: Totals for groups are given in square brackets; index codes, parentheses.

and

$$q_{p.ijk} = \Pr(\mathrm{Colpo} = k|\mathrm{Popn} = p, \mathrm{Dis} = i, \mathrm{Pap} = j) = \frac{\pi_{p.ijk}}{q_{p.ij}q_{p.i}}. \tag{15}$$

Note that $q_{1.i1} = 1$ since the Pap smear is always positive in the diagnostic population. We assume that, conditional on disease and Pap smear, colposcopy is independent of the population, so $q_{1.ijk} = q_{0.ijk} = q_{ijk}$. With these provisions, an independent set of parameters is

$$q_{0.1}, q_{1.1}, q_{0.11}, q_{0.00}, q_{101}, q_{000}, q_{111}, q_{010}, \tag{16}$$

which are the prevalences in the screening and diagnostic populations, the sensitivity and specificity of the Pap smear in the screening population, and the conditional sensitivities and specificities of colposcopy given negative and positive Pap smear results. For convenience, let

$$x_{p.ij} = \sum_k x_{p.ijk} \qquad \text{and} \qquad x_{p.i} = \sum_j x_{p.ij}.$$

Then reexpressing the likelihood in terms of the independent set of q parameters yields

$$\begin{aligned} f(\mathbf{x}|\mathbf{q}) \propto{}& \left[\prod_p q_{p.1}^{x_{p.1}}(1 - q_{p.1})^{x_{p.0}}\right] \\ &\times q_{0.00}^{x_{0.00}}(1 - q_{0.00})^{x_{0.01}} \times q_{0.11}^{x_{0.11}}(1 - q_{0.11})^{x_{0.10}} \\ &\times q_{000}^{x_{p.000}}(1 - q_{000})^{x_{p.001}} \times q_{101}^{x_{p.101}}(1 - q_{101})^{x_{p.100}} \\ &\times q_{010}^{x_{p.010}}(1 - q_{010})^{x_{p.011}} \times q_{111}^{x_{p.111}}(1 - q_{111})^{x_{p.110}}, \end{aligned} \tag{17}$$

where $\mathbf{x}$ and $\mathbf{q}$ denote the vectors of observations and parameters respectively. Note that $q_{1.1}$, the prevalence in a diagnostic population, is not of interest to us, so these factors can be dropped from the likelihood.

We use independent beta priors for each q-parameter, where the beta density is $f(q) = C(a, b)q^a(1 - q)^b$, $0 < q < 1$. The assumption of independence of the priors is reasonable, because they are derived from independent sources, and the beta for the marginal priors is convenient. Denote the prior beta parameters by $a_{p.i}$, $b_{p.i}$, $a_{0.jj}$, $b_{0.jj}$, a_{iji}, and b_{iji}, where the subscripts follow the same convention as for the q's. The posterior distribution is also a product of independent betas, and we denote the posterior parameters by α and β, with corresponding subscripts. They are determined from the data and prior parameters in the same way that the beta parameters are updated when using a beta prior with a binomial sampling distribution, for example,

$$\alpha_{0.00} = x_{0.00} + a_{0.00} \qquad \text{and} \qquad \beta_{0.00} = x_{0.01} + b_{0.00}.$$

We specified our priors in terms of a prior mean, μ, and variance, σ^2. These can be converted to the beta parameters via

$$a = \mu\left[\frac{\mu(1 - \mu)}{\sigma^2} - 1\right] \qquad \text{and} \qquad b = \frac{a(1 - \mu)}{\mu}.$$

We also incorporated constraints on the shape of the beta distribution according to what is reasonable for sensitivities, specificities, and prevalences. An a parameter <1 creates a singularity at 0. This is unreasonable for a prior on a sensitivity or specificity of a diagnostic test, because useful tests have

sensitivities and specificities >0. Therefore, for priors on sensitivities and specificities, we constrain $a \geq 1$. Likewise, a b parameter <1 creates a singularity at 1, which is unreasonable for a prevalence for an uncommon disease, and for prevalences we constrain $b \geq 1$.

$\text{RO}(\delta^{\text{SOC}})$ and $\text{RO}(\delta^{\text{st}})$ given by the linear inequalities discussed in Section 3.2 have coefficients that are functions of these q parameters. Thus an RO is a random set and the posterior distribution for the q parameters can be transformed to a posterior for $\text{RO}(\delta^{\text{SOC}})$ and $\text{RO}(\delta^{\text{st}})$. We performed a Monte Carlo simulation with 10,000 trials to characterize this random set and then computed summary features of this random set, which we describe in the next section.

Once estimated, all of the constraints determining an RO may not be active. A constraint is active when it is satisfied with equality for at least some points in the RO. These active constraints were identified as follows. Consider the linear programming (LP) problem $\mathbf{Ax} \leq \mathbf{b}$, where $\mathbf{A}$ is the matrix of coefficients,

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1^t \\ \mathbf{a}_2^t \\ \vdots \\ \mathbf{a}_k^t \end{bmatrix}.$$

To determine whether the ith constraint is active (satisfied at equality), we use LP to maximize $\mathbf{a}_i^t\mathbf{x}$ subject to the constraints $\mathbf{Ax} \leq \mathbf{b}$. If the solution to this maximization problem, $\mathbf{x}^*$, satisfies $\mathbf{a}_i^t\mathbf{x}^* = b_i$, then $\mathbf{a}_i^t$ is an active constraint.

3.4 Results

After determining the form of the prior distributions and the constraints as discussed in the previous section, we took the prior means and standard deviations from literature values. The prior means and standard deviations—and the sources of those prior means and standard deviations—are summarized in Table 2. Because colposcopy is typically not done in cases of negative Pap smears, there is little information regarding q_{000} and q_{101}. The article by Mitchell, Schottenfeld, Tortolero-Luna, Cantor, and Richards-Kortum (1999) reports on the accuracy of colposcopy as a screening test (i.e., without a previous Pap smear). Note that

$$\Pr(\text{Colpo} = k|\text{Dis} = i) = q_{i1k}q_{0.i1} + q_{i0k}q_{0.i0}.$$

We replaced the $q_{0.ij}$'s and q_{i1k} by their prior means and solved for the prior mean of q_{i0k} to obtain the prior means for q_{101} and q_{000} in Table 2. Because we were uncertain about these prior means, we wanted diffuse priors. Therefore, we chose the highest standard deviations according to the constraints discussed in the previous section; these standard deviations are listed in Table 2.

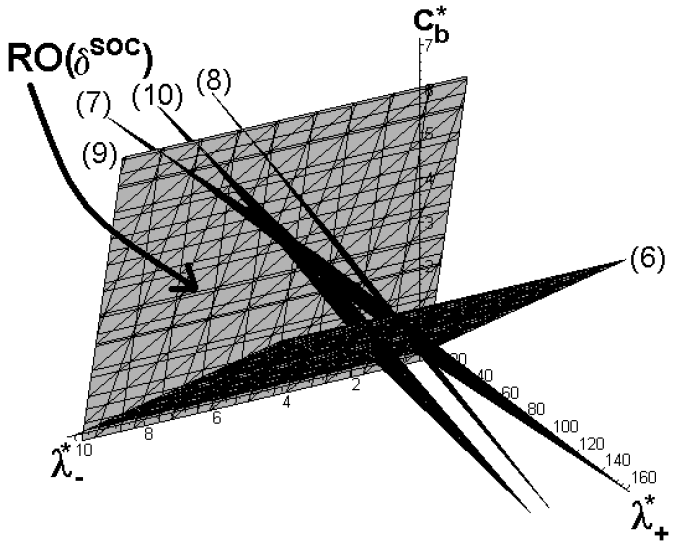

Figure 2. Posterior Mean RO for the SOC: $\text{RO}(\delta^{SOC})$. The five active constraints, (6), (7), (8), (9), and (10). $\text{RO}(\delta^{SOC})$ points to the RO associated with the current standard of care.

After determining the priors, we first computed the posterior means of the coefficients in the inequalities defining $\text{RO}(\delta^{\text{SOC}})$, called the posterior mean $\text{RO}(\delta^{\text{SOC}})$. The active constraints are depicted in Figures 2 and 3, with coefficients given in Table 3. Second, we computed the posterior means of the coefficients defining $\text{RO}(\delta^{\text{st}})$. The active constraints of the posterior mean $\text{RO}(\delta^{\text{st}})$ are depicted in Figures 4 and 5, with coefficients given in Table 4. All active constraints for both $\text{RO}(\delta^{\text{SOC}})$ and $\text{RO}(\delta^{\text{st}})$ involve decisions conditional on a positive Pap smear. In Figure 2 the arrow points to the RO associated with the current SOC. The region contains values for the costs and losses associated with the current SOC being optimal. The boundaries labeled (6)–(10) are the planes associated with (6)–(10) satisfied at equality. The section indicated by the arrow is the RO associated with the current SOC and is the intersection of all of the inequalities. Likewise for Figure 4, the arrow points to the RO associated with "see and treat."

Figures 3 and 5 show slices through the RO associated with the SOC and "see and treat." The slices are cut at various values of c_b^*, as described in the figure legends.

As summary features, we used linear programming to compute the minimum of each of λ_-^*, λ_+^*, and c_b^* for each of the 10,000 Monte Carlo trials. We also computed the minimum and

Table 2. Prior Parameters: Probability Parameters Incorporated in the Prior Distribution

Parameter	*Mean*	*Standard deviation*	*Source*
$q_1 = \Pr(\text{Dis} = 1)$	.03	.1197	Follen (2000)
$q_{11} = \Pr(\text{Pap} = 1\|\text{Dis} = 1)$	.73	.288	Mitchell et al. (1999)
$q_{00} = \Pr(\text{Pap} = 0\|\text{Dis} = 0)$	.63	.300	Mitchell et al. (1999)
$q_{111} = \Pr(\text{Colpo} = 1\|\text{Pap} = 1, \text{Dis} = 1)$	.96	.137	Mitchell et al. (1998)
$q_{010} = \Pr(\text{Colpo} = 0\|\text{Pap} = 1, \text{Dis} = 0)$	.48	.284	Mitchell et al. (1998)
$q_{101} = \Pr(\text{Colpo} = 1\|\text{Pap} = 0, \text{Dis} = 1)$	.59	.2995	Mitchell et al. (1999)*
$q_{000} = \Pr(\text{Colpo} = 0\|\text{Pap} = 0, \text{Dis} = 0)$	.55	.2963	Mitchell et al. (1999)*

*Denotes that the values were computed from given values in the listed sources as described in the text.

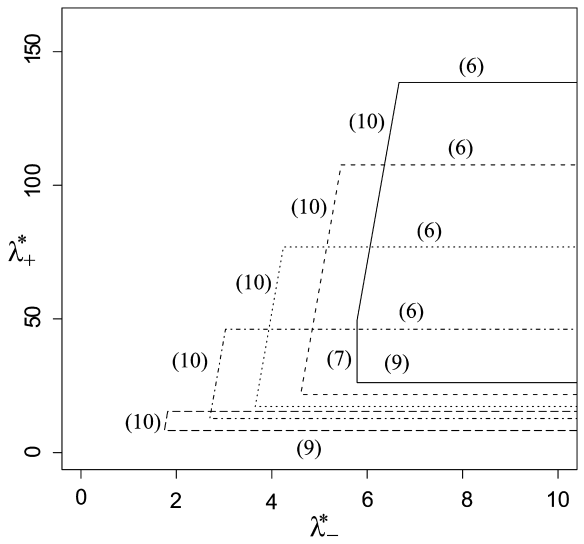

Figure 3. Posterior Mean RO for the SOC: RO(δ^{SOC}) Slice Plot. This figure shows slices through RO(δ^{SOC}) at various values for c_b^* (– – $c_b^* = .5$; ·–· $c_b^* = 1.5$; ····· $c_b^* = 2.5$; - - - $c_b^* = 3.5$; —— $c_b^* = 4.5$). The scale is the same as for Figure 2. The lower boundary on the λ_+ axis is (9) for all slices, but because of the crowded conditions of the plot, only the first and last slices are labeled. Also note that (7) comes into play only for higher values of c_b^*. As a reminder, RO(δ^{SOC}) is the region contained within the bounds and opens to the right (because the region is unbounded as λ_- tends to ∞).

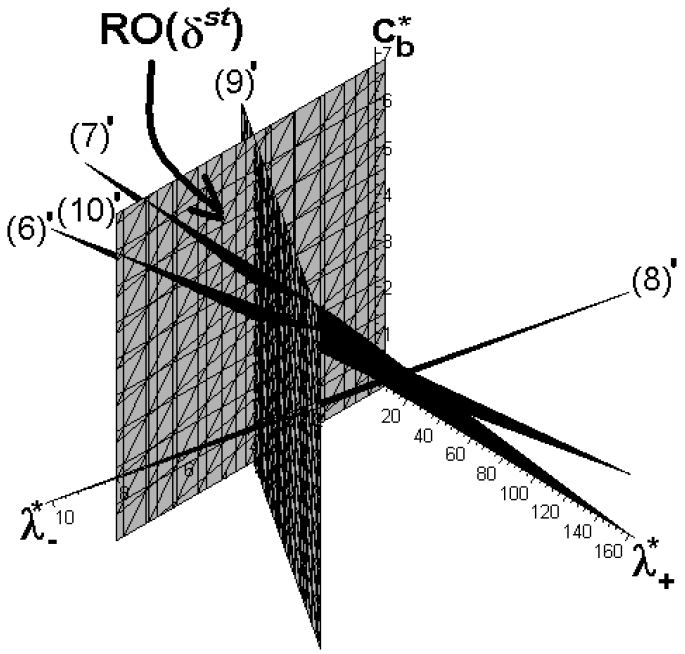

Figure 4. Posterior Mean RO for "See-and-Treat Colposcopy": RO(δ^{st}). The five active constraints, (6)′, (7)′, (8)′, (9)′, and (10)′, for the RO associated with "see-and-treat colposcopy." The "′" mark denotes that left side of the equation is changed to be the Bayes risk associated with "see-and-treat" [the right side of (7)]. Note that (7)′ is the boundary that is between "see-and-treat" and the SOC, and the inequality is simply reversed.

maximum C/B ratio for each of the 10,000 Monte Carlo trials using fractional linear programming (Charnes and Cooper 1973). The posterior means (and standard deviations) of these values are listed in Table 5. Note that the maximum C/B ratio for RO(δ^{SOC}) is infinity. This results from λ_-^* being unbounded in the positive direction because the SOC has no chance of overtreatment.

4. DISCUSSION AND CONCLUSIONS

Recall that there are two main applications of an IDT analysis. First, if we believe that the decision rule is optimal, as we do with the SOC, then we can use the IDT analysis results to characterize the costs and losses associated with that decision rule. Second, if we do not know whether the decision rule under consideration is optimal, as with the "see-and-treat" strategy, then we can use IDT to characterize the costs and losses necessary for the rule to be optimal. These results can then be considered by experts or compared to results from the literature and help assess the optimality of the decision rule.

First, consider the SOC. The fact that the SOC for CIN is widely accepted (the medical community recommends it, insurance companies pay for it, and most patients adhere to it) suggests that many people have faith that it is the optimal procedure. Thus from our IDT analysis of the SOC, we identified minimum values for the costs and losses, which are reported in Table 5.

The boundaries of RO(δ^{SOC}) further restrict the range of the costs and losses. For example, consider the specific bound-

Table 3. Posterior Mean Values (and standard deviations) of Coefficients in Active Constraints for the Posterior Mean RO(δ^{SOC})

Equation	λ_-^*	λ_+^*	c_b^*	*Right side*
(6)	$0_{(0)}$	$.008_{(.005)}$	$-.251_{(.027)}$	0
(7)	$-.582_{(.050)}$	$0_{(0)}$	$.749_{(.027)}$	0
(8)	$-.582_{(.050)}$	$-.008_{(.005)}$	$.499_{(.053)}$	0
(9)	$0_{(0)}$	$-.167_{(.061)}$	$.749_{(.027)}$	−1
(10)	$-.825_{(.064)}$	$.008_{(.005)}$	$.749_{(.027)}$	−1

NOTE: Right side is the right side of the equation. The coefficients correspond to RO(δ^{SOC}) using the vague priors given in Table 2.

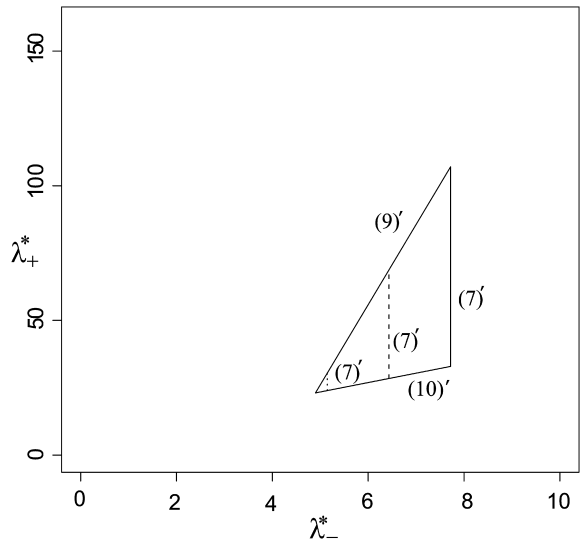

Figure 5. Posterior Mean RO for "See-and-Treat Colposcopy": RO(δ^{st}) Slice Plot. This figure shows slices through RO(δ^{st}) at various values for c_b^* (····· $c_b^* = 4$; - - - $c_b^* = 5$; —— $c_b^* = 6$). The scale is the same as in Figure 4. Note that two of the planes [(9) and (10)] are visible only in the solid line. They are parallel to the c_b^* plane and thus superimpose on each other for the various slices. RO(δ^{st}) is in the interior of the slices.

Table 4. Posterior Mean Values (and standard deviations) of Coefficients in Active Constraints for the Posterior Mean RO(δ^{st})

Equation	λ_-^*	λ_+^*	c_b^*	*Right side*
(6)	$.582_{(.050)}$	$.008_{(.005)}$	$-1_{(0)}$	0
(7)	$.582_{(.050)}$	$0_{(0)}$	$-.749_{(.027)}$	0
(8)	$0_{(0)}$	$.008_{(.005)}$	$-.251_{(.027)}$	0
(9)	$-.243_{(.028)}$	$.008_{(.005)}$	$0_{(0)}$	-1
(10)	$.582_{(.050)}$	$0–.167_{(.061)}$	$0_{(0)}$	-1

NOTE: Right side is the right side of the equation. The coefficients correspond to the region of optimality for "see-and-treat colposcopy," RO(δ^{st}), using the vague priors given in Table 2.

ary in RO(δ^{SOC}) encountered if c_b^* becomes large for a given λ_-^* and λ_+^*. This constraint plane corresponds to (7) (see Fig. 2), and is the boundary between the SOC and "see-and-treat." From Table 3, we deduce that the net cost for treating nondiseased individuals (λ_-) must be larger than $.749/.582 \approx 1.29$ times the cost of biopsy.

In addition, we can use the results of our IDT analysis to confirm other reported results. Cantor et al. (1999) used a qualitative literature review to determine that the C/B ratio for the detection of CIN satisfies $.01 \leq \lambda_-/\lambda_+ \leq .05$. According to our IDT analysis, assuming that the SOC is optimal, the mean minimum C/B ratio is close to the upper end of the interval determined by Cantor et al. In fact, for 33% of the 10,000 Monte Carlo trials, the minimum C/B ratio is $>.05$. This gives evidence that the C/B ratio may not be as low as that reported by Cantor et al.

The second use for IDT involves assessing whether or not a decision rule is optimal. Thus IDT can help experts assess the performance of newly considered decision rules, such as the "see-and-treat" strategy. Consider the cost of biopsy; preliminary data have shown that the incremental monetary cost of a biopsy subsequent to a positive colposcopy result is approximately the same as the monetary cost of colposcopy itself (SBC, personal communication). Furthermore, patients reported similar levels of anxiety for the two procedures, and significantly more pain for the colposcopy than for biopsy (Basen-Engquist et al. 2003). Most other costs involved (e.g., transportation costs, work hours lost) apply only to colposcopy and do not have an incremental cost associated with the biopsy. Considering this information, we conclude that the incremental cost of biopsy should be roughly the same as or less than the cost of colposcopy ($c_b^* \leq 1$).

Table 5 shows that the posterior minimum value for c_b^* found in RO(δ^{st}) is about 3.88, which is much too high given the considerations discussed earlier. In fact, of the 10,000 Monte Carlo trials, the minimum c_b^* value never fell below 3. [In contrast, the posterior minimum for c_b^* in RO(δ^{SOC}) reported in Table 5 is $\leq$1.] This provides substantial evidence that "see-and-treat" could be optimal only when there are substantial contributions to the cost of colposcopy other than those encountered in the patient population considered.

Table 5. Posterior Mean Values (and standard deviations) of Summary Features

	Standard of care	*See-and-treat*
Minimum λ_-^*	$1.52_{(.22)}$	$5.02_{(.82)}$
Minimum λ_+^*	$8.23_{(3.85)}$	$27.43_{(13.64)}$
Minimum c_b^*	$.24_{(.13)}$	$3.88_{(.58)}$
Minimum C/B ratio	$.045_{(.032)}$	$.035_{(.023)}$
Maximum C/B ratio	∞	$.30_{(.14)}$

Several limitations arise from the assumptions that we made to simplify the problem of applying IDT to the detection and treatment of CIN. First, we assumed that the tests resulted in a binary outcome. In fact, the outcomes of the Pap smear, colposcopy, and biopsy compose a set of ordered categories, and the decisions are based on whether or not the result exceeds some threshold. We have considered only decision rules that use the same thresholds as currently implemented in the SOC, which is a subset of all decision rules. This makes our inferences on the lower bound on the C/B ratio conservative, for instance.

Second, we dichotomized the treatment actions into treat and not treat. When assuming that the SOC is optimal, because we combine "watchful waiting" with no treatment, our λ_- is a weighted average of the losses from surgically treating disease-free and LGSIL patients. However, there is some benefit to the surgical treatment for LGSIL patients (because about 20% may progress to HGSIL), so the λ_- values resulting from our IDT analysis are less than the loss from overtreatment of disease-free patients. This assumption does not affect results about λ_+ and has little or no effect on c_b.

Third, the SOC for CIN involves annual screening with a Pap smear, followed by the rest of the sequence if necessary. Our analysis concerns only a single screening event. In particular, λ_+ is the loss from not treating a patient with HGSIL in a single screen. It should not be confused with the loss of not treating a patient before she progresses to cancer, which would definitely be larger. The IDT analysis becomes complicated as one tries to consider all possible screening intervals. This is a subject for future research.

Fourth, we assumed that biopsy was error-free, which is obviously false. The error rate is believed to be very small, and we know of no studies reporting estimates. The fact that there is a nonzero error rate means that a small percentage of individuals are overtreated. This would create a large upper bound on λ_- in RO(δ^{SOC}), but otherwise would have little effect.

Fifth, the unit for the costs is "colposcopy units," because the costs and losses are relative to the cost of colposcopy. When looking for a numerical representation of the costs and losses, one still must perform some sort of analysis to determine the cost of colposcopy that includes the monetary cost of colposcopy and other costs, such as transportation, child care, and pain.

Nevertheless, IDT as presented in this article still provides very useful information about costs and losses associated with decision rules. First, if we believe that a decision rule is optimal, then IDT characterizes the costs and losses associated with that decision rule. In such a case we can use the IDT analysis to confirm other findings regarding the costs and losses. In this study we found evidence suggesting that the C/B ratio for diagnosing and treating CIN is slightly higher than previously reported.

Second, IDT can be used to assess whether or not a decision rule is optimal. We characterized the costs and losses associated with a proposed alternative to the SOC, "see and treat colposcopy," and found that the costs and losses necessary for the competing strategy to be optimal are not practical in the patient population considered. Furthermore, IDT enabled us to

consider a strategy and compare strategies without having to convert patient outcomes, such as pain and anxiety, to monetary amounts.

There are many other potential applications for IDT within medicine, especially where there is an established treatment strategy. Because this is the initial application of IDT, there are several opportunities for further work, including relaxing the assumptions and using priors on the costs and losses. Furthermore, IDT is not limited to the medical setting; it can be used in any decision setting where the standard strategy is believed to be optimal.

[Received October 2003. Revised June 2005.]

REFERENCES

Ahuja, R. K., and Orlin, J. B. (2001), "Inverse Optimization," *Operations Research*, 49, 771–783.

Basen-Engquist, K., Shinn, E. H., Warneke, C., de Moor, C., Le, T., Richards-Kortum, R., and Follen, M. (2003), "Patient Distress and Satisfaction With Optimal Spectroscopy in Cervical Dysplasia Detection," *American Journal of Obstetrics and Gynecology*, 189, 1136–1142.

Ben-Akiva, M., and Lerman, S. R. (1985), *Discrete Choice Analysis: Theory and Application to Travel Demand*, Cambridge, MA: MIT Press.

Cantor, S. B., Sun, C. C., Tortolero-Luna, G., Richards-Kortum, R., and Follen, M. (1999), "A Comparison of C/B Ratios From Studies Using Receiver Operating Characteristic Curve Analysis," *Journal of Clinical Epidemiology*, 52, 885–892.

Cardenas-Turanzas, M. D., Follen, M., Benedet, J. L., and Cantor, S. B. (2005), "See-and-Treat Strategy for Diagnosis and Management of Cervical Squamous Intraepithelial Lesions," *Lancet Oncology*, 6, 43–50.

Charnes, A., and Cooper, W. W. (1973), "An Explicit General Solution in Linear Fractional Programming," *Naval Research Logistics Quarterly*, 20, 449–467.

DeGroot, M. H. (1970), *Optimal Statistical Decisions*, New York: McGraw-Hill.

Ferguson, T. S. (1967), *Mathematical Statistics, A Decision-Theoretic Approach*, New York: Academic Press.

Ferris, D. G., Hiner, B. L., Pfenninger, J. L., and Zuber, T. J. (1996), " 'See and Treat' Electrosurgical Loop Excision of the Cervical Transformation Zone," *Journal of Family Practice*, 42, 253–257.

Follen, M. (2000), "A Clinical Trial Using 4-Hydroxyphenylretinamide for the Treatment of High-Grade Squamous Intraepithelial Lesions of the Cervix: An Interim Analysis," in *Dissertation Abstracts International*, Vol. 61(2B), Ann Arbor, MI: Proquest Information and Learning, Abstract 99-63781, p. 792.

——— (2001), "Preinvasive Squamous Lesions of the Female Lower Genital Tract," in *Operative Gynecology* (2nd ed.), eds. D. M. Gershenson, A. H. DeCherney, S. L. Curry, and L. Brubaker, Philadelphia: Saunders, pp. 273–299.

Harris, J. (1987), "QUALYfying the Value of Life," *Journal of Medical Ethics*, 13, 117–123.

Jordan, J. (1985), "Colposcopy in the Diagnosis of Cervical Cancer and Precancer," *Clinics in Obstetrics & Gynaecology—Cancer of the Cervix: Diagnosis and Treatment*, 12, 67–76.

Mitchell, M. F., Cantor, S. B., Brookner, C., Utzinger, U., Schottenfeld, D., and Richards-Kortum, R. (1999), "Screening for Squamous Intraepithelial Lesions With Fluorescence Spectroscopy," *Obstetrics & Gynecology*, 94, 889–896.

Mitchell, M. F., Schottenfeld, D., Tortolero-Luna, G., Cantor, S. B., and Richards-Kortum, R. (1998), "Colposcopy for the Diagnosis of Squamous Intraepithelial Lesions: A Meta-Analysis," *Obstetrics & Gynecology*, 91, 626–631.

Sox, H. C., Blatt, M. A., Higgins, M. C., and Marton, K. I. (1988), *Medical Decision Making*, Boston: Butterworth.

Weinstein, M. C., and Fineberg, H. V. (1980), *Clinical Decision Analysis*, Philadelphia: Saunders.