The Annals of Statistics
1992, Vol. 20, No. 4, 2177–2182

INVARIANCE PROPERTIES OF DENSITY RATIO PRIORS[1]

BY LARRY WASSERMAN

Carnegie Mellon University

Density ratio neighborhoods are classes of probabilities that are used in robust Bayesian inference. These classes are invariant under Bayesian updating and marginalization. This makes them computationally convenient in robust Bayesian inference. We show that this is the unique class of probabilities that has these invariance properties. Aside from its theoretical value, this result has computational implications as well.

1. Introduction. Classes of probability measures are used in Robust Bayesian inference [Berger (1984, 1985, 1990)] and in the theory of upper and lower probabilities [Waley (1991)]. Various classes have been considered [see, for example, Berger (1990), Berger and Berliner (1986), Lavine (1991a, b), Walley (1991), Wasserman (1992) and Wasserman and Kadane (1990)]. An intuitively appealing class of priors was proposed by DeRobertis and Hartigan (1981). Berger (1990) called this class the *density ratio class*. This class is invariant under Bayesian updating and marginalization. Thus, if a density ratio neighborhood around a prior π is updated by Bayes' theorem, it is transformed into a density ratio neighborhood around the corresponding posterior. As shown in Wasserman (1992) and Wasserman and Kadane (1992a), these invariance properties simplify many computations. In particular, Wasserman (1992) showed that these properties allow us to apply Gibbs sampling techniques [Gelfand and Smith (1990)] to the problem of finding bounds on posterior expectations. Apparently, this was the first application of Monte Carlo techniques to these types of robust Bayesian problems and it raises the following question: Do other classes possess these invariance properties? In this paper we show that, subject to mild regularity conditions, the density ratio class is the only class that satisfies these invariance properties. In Section 2 we present some preliminary material. In Section 3 we study updating and marginalization invariance and we present the main result (Theorem 3.1). In Section 4 we discuss the implications of the results.

2. Preliminaries. Let Θ be a compact subset of $\mathbb{R}$ and let $\mathscr{B}(\Theta)$ be the Borel subsets of Θ. Without loss of generality, take $\Theta = [0, 1]$. Let $\mathscr{L}$ be the set of all measurable functions ℓ mapping Θ into $[0, \infty)$ and let $\mathscr{F} = \{f \in \mathscr{L}; \int f(\theta)\lambda(d\theta) = 1\}$, where λ is Lebesgue measure. We shall regard every $f \in \mathscr{F}$ as a prior probability density function and every $\ell \in \mathscr{L}$ as a likelihood function. Given $f \in \mathscr{F}$ and $\ell \in \mathscr{L}$, we write $\ell \otimes f$ to mean the posterior

Received July 1991; revised January 1992.
[1]Research supported by NSF Grant DMS-90-05858.
AMS 1980 *subject classifications*. Primary 62F15; secondary 62F35.
Key words and phrases. Density ratio classes, invariance, robust Bayesian inference.

density function $\ell(\theta)f(\theta)/\int \ell(\theta)f(\theta)\lambda(d\theta)$ as long as the denominator is positive and finite. If $\Gamma \subset \mathscr{F}$ is a class of densities, then we define $\ell \otimes \Gamma = \{\ell \otimes f;\ f \in \Gamma\}$. Let μ be the uniform probability density function.

Let Γ_k: $\mathscr{F} \to 2^{\mathscr{F}}$, where $2^{\mathscr{F}}$ is the set of all subsets of $\mathscr{F}$ and k is a nonnegative real number, and assume that, for every $f \in \mathscr{F}$, $\Gamma_0(f) = \{f\}$ and that $k < m$ implies that $\Gamma_k(f)$ is strictly contained in $\Gamma_m(f)$. We also assume that $\Gamma_k(f) = \bigcup_{m<k}\Gamma_m(f) = \bigcap_{m>k}\Gamma_m(f)$. We call $\Gamma_k(f)$ a *neighborhood* of f and we call $\{\Gamma_k(f);\ k \geq 0,\ f \in \mathscr{F}\}$ a *neighborhood system*.

Following Ryff (1965) and Marshall and Olkin (1979), we say that two densities f and g are *equimeasurable* if $\lambda(\{\theta;\ f(\theta) > t\}) = \lambda(\{\theta; g(\theta) > t\})$ for all real t. In this case write $f \sim g$. This is the continuous version of saying that g is a permutation of f. Given f, there is a unique, nonincreasing, right-continuous function f^*, called the *decreasing rearrangement* of f, such that $f \sim f^*$. If $f, g \in \mathscr{F}$, we say that f is *majorized* by g and we write $f \prec g$ if $\int_0^s f^*(\theta)\lambda(d\theta) \leq \int_0^s g^*(\theta)\lambda(d\theta)$ for every s. This is usually taken to mean that f is "more uniform" than g. A fact that we shall use is that if f is constant over a set A and over A^c and if $\int_A f(\theta)\lambda(d\theta) = \int_A g(\theta)\lambda(d\theta)$, then $f \prec g$. We shall make the following regularity assumptions:

ASSUMPTION 1. If $f \in \Gamma_k(\mu)$ and $g \prec f$, then $g \in \Gamma_k(\mu)$.

ASSUMPTION 2. $\Gamma_k(\mu)$ is convex and weakly compact.

Let $S(f)$ be the support of f and define the *density ratio metric* δ on $\mathscr{F}$ by

$$\delta(f, g) = \operatorname*{ess\,sup}_{\theta, \phi \in S(f)} \log \frac{f(\theta)/f(\phi)}{g(\theta)/g(\phi)}$$

if $S(f) = S(g)$ and $\delta(f, g) = \infty$ otherwise. Note that $\delta(\mu, g) = \log(g^*(0)/g^*(1))$. Define the *density ratio neighborhood* by $\overline{\Gamma}_k(f) = \{g;\ \delta(f, g) \leq k\}$. This is a special case of the classes studied in DeRobertis and Hartigan (1981) and has been investigated by Berger (1990), Lavine (1991a, b), Wasserman (1992) and Wasserman and Kadane (1992a, b). For a given neighborhood system Γ_k define

$$\rho(k) = \sup_{f \in \Gamma_k(\mu)} \delta(\mu, f) = \sup_{f \in \Gamma_k(\mu)} \log \frac{f^*(0)}{f^*(1)}.$$

Note that $\rho(0) = 0$ and that $\rho(k)$ is increasing in k. We also make the following assumption about Γ_k:

ASSUMPTION 3. For every $k \geq 0$, $\rho(k) < \infty$.

For every $n \geq 1$, let $\Theta_{(n)} = \{\theta_1, \ldots, \theta_n\}$ represent a finite set with n elements and let $\mathscr{F}_{(n)}$ be the set of all probability mass vectors on $\Theta_{(n)}$. We assume that Γ_k is well defined over $\mathscr{F}_{(n)}$ for all n, that is, if $f = (f_1, \ldots, f_n) \in \mathscr{F}_{(n)}$, then $\Gamma_k(f) \subset \mathscr{F}_{(n)}$. The decreasing rearrangement of $f =$

$(f_1, \dots, f_n)$ is the vector obtained by rearranging the elements of f in decreasing order. Then $f \prec g$ if $\Sigma_1^j f_i^* \le \Sigma_1^j g_i^*$ for $j = 1, \dots, n$. Let $\mu_{(n)} = (1/n, \dots, 1/n)$. For $\mathscr{F}_{(n)}$, the density ratio metric is defined by

$$\delta(f, g) = \sup_{\theta_i, \theta_j \in S(f)} \log \frac{f_i/f_j}{g_i/g_j}$$

if $S(f) = S(g)$ and $\delta(f, g) = \infty$ otherwise. Define

$$\hat{\rho}(k) = \sup_{f \in \mathscr{F}_{(2)}} \delta(\mu_{(2)}, f) = \log \sup_{f \in \mathscr{F}_{(2)}} f_1^*/f_2^*.$$

We assume that Assumptions 1–3 apply to the finite sets as well. It follows that $\hat{\rho}(k)$ is a strictly increasing, continuous function and that $\Gamma_k(\mu_{(2)}) = \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)})$.

3. Invariance. We say that a neighborhood system Γ_k is *update invariant* if $\Gamma_k(\ell \otimes f) = \ell \otimes \Gamma_k(f)$ for every $f \in \mathscr{F}$, $k \ge 0$ and $\ell \in \mathscr{L}$. This means that operations of Bayesian updating and neighborhood formation commute.

LEMMA 3.1. *If Γ_k is update invariant, then $\Gamma_k(f) \subset \overline{\Gamma}_{\rho(k)}(f)$ for every $k \ge 0$ and every $f \in \mathscr{F}$.*

PROOF. Let $f \in \Gamma_k(\mu)$. Then $\delta(\mu, f) \le \rho(k)$ so $f \in \overline{\Gamma}_{\rho(k)}(\mu)$. Hence $\Gamma_k(\mu) \subset \overline{\Gamma}_{\rho(k)}(\mu)$. Now, using update invariance we have $\Gamma_k(f) = \Gamma_k(f \otimes \mu) = f \otimes \Gamma_k(\mu) \subset f \otimes \overline{\Gamma}_{\rho(k)}(\mu) = \overline{\Gamma}_{\rho(k)}(f \otimes \mu) = \overline{\Gamma}_{\rho(k)}(f)$ $\square$

If $X: \Theta \to \mathbb{R}$ is a simple random variable that is constant over the elements of the partition $\pi = \{A_1, \dots, A_n\}$ and $f \in \mathscr{F}$, then the marginal of f induced by X can be regarded as a mass vector $f^{\pi} = (f_1, \dots, f_n) \in \mathscr{F}_{(n)}$, where $f_i = \int_{A_i} f(\theta)\lambda(d\theta)$. Given $\Gamma \subset \mathscr{F}$, define $\Gamma^{\pi} = \{f^{\pi};\ f \in \Gamma\}$. We shall say that Γ_k is *marginalization invariant* if, for all $f \in \mathscr{F}$, for all $k \ge 0$ and all partitions π, $\Gamma_k^{\pi}(f) = \Gamma_k(f^{\pi})$. In words, the marginal of the neighborhood equals the neighborhood around the marginal.

REMARK. At the end of Section 2 we pointed out that $\Gamma_k(\mu_{(2)}) = \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)})$. If Γ_k is update invariant, then for every $f \in \mathscr{F}_{(2)}$, $\Gamma_k(f) = \Gamma_k(f \otimes \mu_{(2)}) = f \otimes \Gamma_k(\mu_{(2)}) = f \otimes \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)}) = \overline{\Gamma}_{\hat{\rho}(k)}(f)$.

To make further progress, we will need the following lemma, which gives an integral representation of $\overline{\Gamma}_k$. For every $t \in (0, 1)$ and $z \ge 1$ define a probability density function $\tau^{z,t}$ by $\tau^{z,t}(\theta) = zc^{z,t}$ if $\theta \le t$ and $\tau^{z,t}(\theta) = c^{z,t}$ if $\theta > t$, where $c^{z,t} = \{zt + (1 - t)\}^{-1}$. For $t = 0$ or 1 define $\tau^{z,t} = \mu$.

LEMMA 3.2. *The following two statements are equivalent*:

(i) $f \in \overline{\Gamma}_k(\mu)$.

(ii) *There exists a probability measure R on $\mathscr{B}(\Theta)$ and a number $z \in [1, \exp(k)]$ such that $f^*(\theta) = \int_0^1 \tau^{z,t}(\theta)R(dt)$ for all $\theta \in \Theta$.*

PROOF. Suppose (i) holds. Then $1 \le z = f^*(0)/f^*(1) \le \exp(k)$. Let $h(\theta) = f^*(\theta)/f^*(1)$. If $z = 1$, take R to be a point mass at $t = 1$ and we are done. So assume $z > 1$. Define a set function V by $V([0,\theta]) = (z - h(\theta))/(z-1)$. Then $V(\Theta) = 1$, and V can be extended to be a probability measure on $\mathscr{B}(\Theta)$. Let $h^{z,t}(\theta) = \tau^{z,t}(\theta)/c^{z,t}$. Then $\int_0^1 h^{z,t}(\theta)V(dt) = V([0,\theta]) + z(1 - V([0,\theta])) = h(\theta)$. Now, $f^*(\theta) = f^*(1)h(\theta) = f^*(1)\int_0^1 h^{z,t}(\theta)V(dt) = \int_0^1 \tau^{z,t}(\theta)R(dt)$, where R is defined by $R([0,t]) = V([0,t])f^*(1)/c^{z,t}$. Now we confirm that R has total mass 1 and so can be extended to be a probability measure. We have

$$1 = \int_0^1 f^*(\theta)\lambda(d\theta) = \int_0^1\int_0^1 \tau^{z,t}(\theta)R(dt)\lambda(d\theta)$$

$$= \int_0^1\int_0^1 \tau^{z,t}(\theta)\lambda(d\theta)R(dt) = \int_0^1 R(dt) = R(\Theta).$$

Hence (ii) holds. Now suppose that (ii) holds. Then $\delta(\mu, f) = \log(f^*(0)/f^*(1)) = \log h(0) = \log z$ for some $z \in [1, \exp(k)]$ since $h(\theta) = \int_0^1 h^{z,t}(\theta)\lambda(d\theta)$. Thus, (i) holds. □

LEMMA 3.3. *If Γ_k is update invariant and marginalization invariant, then $\rho(k) = \hat{\rho}(k)$ for all $k \ge 0$.*

PROOF. Let $z = \exp(\rho(k))$. Choose $f \in \Gamma_k(\mu)$ such that $f^*(0)/f^*(1) = z$. Choose $\delta > 0$. Since f^* is nonincreasing, for sufficiently small, positive ε, $f^*(\theta)/f^*(\phi) > z - \delta$ for all $\theta \in [0,\varepsilon]$ and $\phi \in [1-\varepsilon, 1]$. Hence

$$\frac{\int_0^\varepsilon f^*(\theta)\lambda(d\theta)}{\int_{1-\varepsilon}^1 f^*(\theta)\lambda(d\theta)} \ge \frac{\varepsilon f^*(\varepsilon)}{\varepsilon f^*(1-\varepsilon)} > z - \delta. \tag{1}$$

Let $B = [0,\varepsilon] \cup [1-\varepsilon, 1]$ and let ℓ be the indicator function for B. Then $g = \ell \otimes f^* \in \ell \otimes \Gamma_k(\mu) = \Gamma_k(\ell \otimes \mu) = \Gamma_k(\mu_B)$, where μ_B is the density that is equal to $1/(2\varepsilon)$ on B and 0 otherwise. Let $\pi = \{A, A^c\}$, where $A = [0, 1/2]$. Then $\mu^\pi = \mu_{(2)}$ and $g^\pi = (g_1, g_2)$, where $g_1 = \int_A g(\theta)\lambda(d\theta)$. Appealing to marginalization invariance, we have $g^\pi \in \Gamma_k^\pi(\mu_B) = \Gamma_k(\mu_B^\pi) = \Gamma_k(\mu_{(2)}) = \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)})$. From (i) we conclude that $\delta(\mu_{(2)}, g^\pi) > \log(z - \delta)$. Since $g^\pi \in \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)})$, we conclude that $\delta(\mu_{(2)}, g^\pi) \le \hat{\rho}(k)$. Thus $\log(z-\delta) \le \hat{\rho}(k)$ for every $\delta > 0$ so that $\hat{\rho}(k) \ge \rho(k)$.

Now, with π defined as above, $\Gamma_k^\pi(\mu) = \Gamma_k(\mu^\pi) = \Gamma_k(\mu_{(2)}) = \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)})$. Choose $f \in \overline{\Gamma}_{\hat{\rho}(k)}(\mu_{(2)})$ such that $\delta(\mu_{(2)}, f) = \hat{\rho}(k)$. Then there exists $g \in \Gamma_k(\mu)$ such that $g^\pi = f$. Define the density h by $h(\theta) = \int_A g(\theta)\lambda(d\theta)/\lambda(A)$, if $\theta \in A$, and $h(\theta) = \int_{A^c} g(\theta)\lambda(d\theta)/\lambda(A^c)$, otherwise. Then $h \prec g$ so $h \in \Gamma_k(\mu)$. Note that $\int_A h(\theta)\lambda(d\theta) = \int_A g(\theta)\lambda(d\theta)$. Hence $\delta(\mu, h) = \delta(\mu_{(2)}, f) = \hat{\rho}(k)$. But $h \in \Gamma_k(\mu) \subset \overline{\Gamma}_{\rho(k)}(\mu)$ so $\delta(\mu, h) \le \rho(k)$ and we conclude that $\hat{\rho}(k) \le \rho(k)$. □

REMARK. It follows from Lemma 3.3 and the assumptions about Γ_k in Section 2 that $\rho(0) = 0$ and $\rho(k)$ is a strictly increasing, continuous function of k.

THEOREM 3.1. *If* Γ_k *is update invariant and marginalization invariant, then* $\overline{\Gamma}_{\rho(k)}(f) = \Gamma_k(f)$ *for every* $k \geq 0$ *and every* $f \in \mathscr{F}$. *Hence,* Γ_k *and* $\overline{\Gamma}_k$ *generate the same neighborhoods.*

PROOF. The theorem is clearly true if $k = 0$. Now assume $k > 0$. Let $z = \exp(\rho(k))$. First we show that $\tau^{z,t} \in \Gamma_k(\mu)$. Let $\pi = \{A, A^c\}$, where $A = [0, t]$. By marginalization invariance, $\Gamma_k^{\pi}(\mu) = \Gamma_k(\mu^{\pi}) = \overline{\Gamma}_{\hat{\rho}(k)}(\mu^{\pi}) = \overline{\Gamma}_{\rho(k)}(\mu^{\pi})$. There must exist $f \in \Gamma_k(\mu)$ such that $\delta(\mu^{\pi}, f^{\pi}) = \rho(k)$. Note that $\mu^{\pi} = (t, 1 - t)$ and $f^{\pi} = (f_1, f_2)$, where $f_1 = \int_0^t f(\theta)\lambda(d\theta)$ and $f_2 = 1 - f_1$. Then

$$\delta(\mu^{\pi}, f^{\pi}) = \log \max\left\{ \frac{tf_2}{(1-t)f_1}, \frac{(1-t)f_1}{tf_2} \right\}.$$

Consider the case where $tf_2/((1-t)f_1) \leq ((1-t)f_1)/(tf_2)$—the other case is proved in a similar manner. Then $\rho(k) = \delta(\mu^{\pi}, f^{\pi}) = \log((1-t)f_1/(tf_2))$ so that $f_1 = zt/(1 - t + zt)$. Hence, $\int_A f(\theta)\lambda(d\theta) = \int_A \tau^{z,t}(\theta)\lambda(d\theta)$ and since $\tau^{z,t}$ is constant over A and A^c, $\tau^{z,t} \prec f \in \Gamma_k(\mu)$. By Assumption 1, $\tau^{z,t} \in \Gamma_k(\mu)$. Also, $\tau^{r,t} \in \Gamma_k(\mu)$ for all $r \in [1, z]$. Let $g \in \overline{\Gamma}_{\rho(k)}(\mu)$. By Lemma 3.2, g^* is a mixture of the $\tau^{r,t}$'s for some $r \in [1, z]$ so, by the regularity assumptions, $g \in \Gamma_k(\mu)$. Hence, $\overline{\Gamma}_{\rho(k)}(\mu) \subset \Gamma_k(\mu)$. Now, for any $f \in \mathscr{F}$, $\overline{\Gamma}_{\rho(k)}(f) = f \otimes \overline{\Gamma}_{\rho(k)}(\mu) \subset f \otimes \Gamma_k(\mu) = \Gamma_k(f)$. By Lemma 3.1, $\Gamma_k(f) \subset \overline{\Gamma}_{\rho(k)}(f)$ and the proof is complete. □

4. Discussion. Suppose we use a prior f in a Bayesian analysis of a statistical problem. To evaluate the sensitivity of our inferences to the choice of prior, we may embed f in a convenient class of priors and compute bounds on posterior quantities of interest. If the density ratio class $\overline{\Gamma}_k(f)$ is used, then, by the update invariance property, the set of posteriors after combining the prior with the likelihood ℓ is $\overline{\Gamma}_k(\ell \otimes f)$. In other words, the set of posteriors is again a density ratio class. Further, Wasserman and Kadane (1992a) showed that if it is possible to draw a random sample from f, then bounds on posterior expectations are easily estimated for a density ratio class.

Specifically, suppose we wish to compute the upper bound $\bar{\rho}$ on the posterior expectation of a function $\phi = \phi(\theta)$ of a parameter θ when the prior g varies in the class $\overline{\Gamma}_k(f)$. Then Wasserman and Kadane (1992a) showed the following: Let $\phi^1, \dots, \phi^N$ be a random draw from f_ϕ, the marginal probability distribution of ϕ under f. Let $\phi^{(1)}, \dots, \phi^{(N)}$ be the ϕ^i's ordered from smallest to largest and let $\bar{\phi}$ be the average of the ϕ^i's. As a consequence of marginalization invariance, to find the upper bound of the expectation of ϕ over $\overline{\Gamma}_k(f)$, we need only compute the upper bound over $\{g; \delta(f_\phi, g) \leq k\}$. This maximum is attained at the g that has density proportional to kf_ϕ over a set where ϕ is greater than some constant and is proportional to f_ϕ otherwise. Using this fact, together with standard facts about order statistics, they showed that $\bar{\rho} = \hat{\rho} + O_P(1/\sqrt{N})$, where

$$\hat{\rho} = \max_{1 \leq i \leq N} \{(1 - i/N)(z - 1) + 1\}^{-1} \left\{ (z - 1) \sum_{j=i}^{N} \phi^{(j)}/N + \bar{\phi} \right\},$$

where $z = \exp(k)$. This simplifies the computation of posterior bounds considerably. This simplification is possible because of the invariance properties possessed by the density ratio class. The results in this paper imply that no other class has these properties.

Acknowledgments. The author would like to thank Isabella Verdinelli, the referees and an Associate Editor for helpful comments.

REFERENCES

BERGER, J. and BERLINER, M. (1986). Robust Bayes and empirical Bayes analysis with ε-contaminated priors. *Ann. Statist.* **14** 461–486.

BERGER, J. O. (1984). The robust Bayesian viewpoint (with discussion). In *Robustness in Bayesian Statistics* (J. Kadane, ed.) 63–124. North-Holland, Amsterdam.

BERGER, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer, New York.

BERGER, J. O. (1990). Robust Bayesian analysis: Sensitivity to the prior. *J. Statist. Plann. Inference* **25** 303–328.

DEROBERTIS, L. and HARTIGAN, J. A. (1981). Bayesian inference using intervals of measures. *Ann. Statist.* **9** 235–244.

GELFAND, A. and SMITH, A F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409.

LAVINE, M. (1991a). An approach to robust Bayesian analysis for multidimensional spaces. *J. Amer. Statist. Assoc.* **86** 396–399.

LAVINE, M. (1991b). Sensitivity in Bayesian statistics: The prior and the likelihood. *J. Amer. Statist. Assoc.* **86** 400–403.

MARSHALL, A. W. and OLKIN, I. (1979). *Inequalities: Theory of Majorization and Its Applications.* Academic, New York.

RYFF, J. (1965). Orbits of L^1-functions under doubly stochastic transformations. *Trans. Amer. Math. Soc.* **117** 92–100.

WALLEY, P. (1991). *Statistical Reasoning With Imprecise Probabilities. Monographs Statist. Appl. Probab.* Chapman and Hall, London.

WASSERMAN, L. (1992). Recent methodological advances in robust Bayesian inference (with discussion). In *Bayesian Statistics IV: Fourth Valencia International Meeting on Bayesian Statistics* (J. M. Bernardo, J. O. Berger, A. P. David and A. F. M. Smith, eds.) 483–502. Clarendon, Oxford.

WASSERMAN, L. A. and KADANE, J. B. (1990). Bayes' theorem for Choquet capacities. *Ann. Statist.* **18** 1328–1339.

WASSERMAN, L. and KADANE, J. B. (1992a). Computing bounds on expectations. *J. Amer. Statist. Assoc.* **87** 516–522.

WASSERMAN, L. and KADANE, J. B. (1992b). Symmetric upper probabilities. *Ann. Statist.* **20** 1720–1736.

DEPARTMENT OF STATISTICS
232 BAKER HALL
CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PENNSYLVANIA 15213-3890