Published in final edited form as:

Stroke. 2014 November ; 45(11): 3325–3329. doi:10.1161/STROKEAHA.114.006807.

Picking the Good Apples: Statistics versus Good Judgment in Choosing Stent Operators for a Multicenter Clinical Trial

George Howard, DrPH¹, Jenifer H. Voeks, PhD², James F. Meschia, MD³, Virginia J. Howard, PhD⁴, and Thomas G. Brott, MD³

¹Department of Biostatistics, University of Alabama at Birmingham, School of Public Health, Birmingham, AL

²Department of Neurology, Medical University of South Carolina, Charleston, SC

³Department of Neurology, Mayo Clinic Jacksonville, Jacksonville, FL

⁴Department of Epidemiology, University of Alabama at Birmingham, School of Public Health, Birmingham, AL

Abstract

Background and Purpose—The Carotid Revascularization Endarterectomy versus Stenting Trial (CREST) was completed with a quite low stroke and death rate. A “lead-in” series of patients receiving carotid artery stenting (CAS) was used to select the physician-operators for the study, where performance was evaluated by complication rates and by peer review of cases. Herein, we assess the potential contribution of statistical evaluation of complication rates.

Methods—The ability to discriminate between stent operators who can successfully meet the published guideline of <3% combined rate of stroke and death is calculated under the binomial distribution, based upon a small consecutive case series ($n = 24$ patients).

Results—A criterion of 2 stroke or death events among the 24 patients (<8% event rate) was required of operators. Setting such a high criterion, however, ensures an inability to *exclude* operators who cannot meet the criteria. In fact, if a “good” operator is defined as having a 2% event rate, and a “poor” operator as a 6% event rate, even a series of 240 patients would (on average) still exclude 5.4% of the good operators and include 4.6% of the poor operators.

Conclusions—The low periprocedural event rates in the trial suggest success in separating skillful operators from less skillful. However, it seems unlikely that statistical assessment of event rates in the lead-in contributed to successful selection, but rather successful selection was more likely due to peer review of subjective and other factors including patient volume and technical approaches.

Keywords

Clinical trial; quality control; stroke

Corresponding Author: Thomas G. Brott, MD; Department of Neurology; Mayo Clinic; 4500 San Pablo Road South; Jacksonville, Florida 32224; 904-953-0556; Fax: 904-953-0277; brott.thomas@mayo.edu.

Conflicts of Interest/Disclosures: None.

Introduction

The Carotid Revascularization Endarterectomy versus Stenting Trial (CREST) was completed with periprocedural stroke and death rates in asymptomatic arteries of 2.5% ($\pm$ 0.6%) for those treated with carotid stenting (CAS) and 1.4% ($\pm$ 0.5%) for those treated with carotid endarterectomy (CEA), and corresponding rates in symptomatic arteries of 5.5% ($\pm$ 0.9%) and 3.2% ($\pm$ 0.7%).¹ These rates have been recognized as among the lowest ever reported from randomized clinical trials of carotid surgery and stenting.² Current guidelines specify that the procedural morbidity and mortality rate for asymptomatic patients should be $<3\%$ and for symptomatic patients $<6\%$.²⁻⁴ While the upper limit of the 95% confidence bound on the rates achieved in CREST are above these guidelines, and while stroke and death rates were significantly higher in the CAS than CEA group for the symptomatic patients (and nominally higher for the asymptomatic patients), the observed CAS event rates did meet these guideline criteria.

While CEA was an established procedure, CAS was a newer procedure when CREST was initiated, so the investigators implemented a rigorous assessment of potential interventionists including in a lead-in registry for stent operators.^{5, 6} During this process, the interventional management committee reviewed 10,164 CAS cases from 427 interventionist-applicants (an average of 24 from each operator). The feasible size of the series of patients conducted prior to the admission of the operator to the study is a serious consideration. Even for a very active operator conducting a procedure a week, the stringent requirements of CREST required that approximately six months work need be completed for an operator to be admitted to the study. As such, increases in the size of the qualifying series substantially beyond the neighborhood of 24 are practically infeasible. Accordingly, several criteria were used to approve operators, which have been previously described as: “*a low complication rate, experience with 15 procedures, use of proper standard carotid stenting technique, and avoidance of erroneous techniques (e.g., improper device use, inappropriate balloon sizing, use of 0.0035” wires, use of general anesthesia).*”⁶ These criteria can be classified as evaluating the performance (the low complication rate), establishing the experience (15 procedures), and the reviewing of technique (use of proper technique and avoidance of erroneous techniques). In CREST, of the 427 applicants for participation as an operator, 224 (52%) were approved for randomization; among those 427 applicants, 238 were required to perform cases in the lead-in, and 158 of those applicants were approved for randomization, based upon the lead-in data.⁶ The success of the study in identifying operators with low event rates stands as proof that the process worked, but can this success be attributed to exclusions on the basis performance, experience or technique? Herein, we review the contribution to the statistical assessment of performance in this lead-in registry, and the likelihood that it made a meaningful contribution to selecting the operators who can perform the procedure with a low risk of complication.

The focus of much of this report is on the ability to select operators based on the result of a series of 24 patients, reflecting the average number of patients available for decision making in the CREST trial. However, we also provide estimates of how large an increase in the available series of patients would be required to provide a reasonable basis for the identification of good operators and the exclusion of bad operators. This consideration of the

review process is important more generally, as referring physicians sometimes take a similar approach to ensure the quality of care given to their patients, where requests for performance in recent patients is sometimes used to determine where to make referrals. In addition, patients are advised to seek information on qualifications of potential surgeons or operators prior to scheduling any non-emergency procedures, with lay magazines such as *Consumer Reports* recommending that patients ask four questions in selecting a surgeon prior to a procedure including “What are your success, failure, and complication rates?”.⁷ How, and when, should we look at the “batting averages” of successful procedures to select a surgeon, an operator, or for that matter ... a car mechanic?

Methods

Statistically, assessing the performance of an individual surgeon or operator is “guessing” (statisticians prefer “estimating”) the true rate of poor outcomes following their procedures. Clearly, there are characteristics of the individual patient that could influence the rate of poor outcomes. For example, periprocedural event rates in CAS could be generally higher in the elderly population^{1, 89, 10} or in women.¹¹ However, we assume the goal is to assess the likelihood of stroke or death events for an “average” patient for a specific operator, where a low average risk of events is likely the criteria for approving an individual operator.

One then can assume that there is a “true” event rate for each operator (assumed to not change over time or over patients). For example, the current guidelines suggest CAS and CEA be considered as treatment alternatives for severe asymptomatic carotid artery disease if they can be performed with an event rate <3% in asymptomatic patients.¹² However, if there is variation among operators, how can a study such as CREST systematically choose study operators with event rates below this level? More generally, how can a patient be assured that the “true” event rate for the specific operator is below this level? One might assume that we could look at the observed event rates for a specific operator, and if he/she is a “good apple” their event rate would be <3%, while if they are a “bad apple” their event rate would be >3%. Unfortunately, the situation is more complex than this approach suggests.

For simplicity, assume that the patients treated by a specific operator has a chance of “ p ” of having a periprocedural stroke or death (i.e., poor outcome or complication). Then we are interested in including the physicians with $p \leq 0.03$ (the good apples), and excluding the physicians with $p > 0.03$ (the bad apples). We can estimate an individual operator’s p (we will refer to this estimate as $\hat{p}$) by simply dividing the number of patients with a poor

outcome by the number of patients studied ($\hat{p} = \frac{k}{n}$, where k is the number with poor outcome, and n is the number of patients studied).

There are several challenges to this approach. First, even without changes to the true chance of a poor outcome, chance may imply a different number of complications in two series of patients. For example, a single operator may perform 24 procedures with no complications, and then on the subsequent series of 24 patients the same operator may have 3 poor outcomes. That is, the estimated event rate is only a reflection of the true event rate, and

having an estimated event rate lower than 3% does not assure the reviewer that the true rate is acceptable. The second challenge is inherent in our goal to select operators with a low complication rate of $<3\%$, where in a given 24 procedures, having a single event results in an estimated event rate of 4.2%. As such, a simple rule of having an observed event rate of 3% or less implies that only operators who have no events can participate in the study.

What would be optimal is if we could create a “rule” so that if an operator x has or fewer complications in n patients, then they could be included in the study (for example, we could approve an operator with 0 or 1 complication in 24 procedures). So the goal of this work is to establish whether it is possible to have a rule to identify the good operators and include them in the study (“keep the good apples”) and to identify the bad operators and exclude them from the study (“throw away the bad apples”).

With this as the goal, given that the chance of a poor outcome is constant (p), then the chance of getting x poor outcomes in n patients can be directly calculated from the binomial distribution. We assess if such a rule can be created to keep the good apples and throw away the bad.

We also defined a more general approach to assess the number of patients needed to provide reliable information to include good operators and exclude poor operators. First, for any size of the series of patients, we define a rule to select the threshold of the number of events that a potential operator must meet to be included in the study. Our proposed rule assigns an equal cost of inappropriately excluding a good operator and inappropriately including a poor operator. Specifically, we defined the “error rate” as the sum of the expected percent of good operators (with a “true” 2% event rate) excluded, plus the percent of poor operators (with a true 6%, 8% or 10% event rate) included. For any sample size we selected the threshold to minimize this error rate. For example in Table 1, if the decision was to include operators only if they had zero (0) events, this would result in the exclusion of 38% of the good operators (the 2% column) and the inclusion of 23% of the poor operators (the 6% column), for an error rate of 61%. This error rate is smaller than including operators with 1 or fewer events (an error rate of 66%), or to include operators with 2 or fewer events (an error rate of 84%). Therefore we would declare the “optimal” rule for a series of 24 patients to be to include operators with no events, with an unfortunately high “error rate” of 61%. We then repeated this process for sample size ranging up to 500 and plotted the decline in error rate with the increasing size of the series. We suggest that to make reliable decisions it would be optimal to exclude only about 5% of good operators and include no more than 5% of poor operators, so a goal of an error rate of less than 10% would seem reasonable.

Results

Approaches to keep the good apples

Suppose that a particular operator is a “good apple,” that is his/her true complication rate is 3% or less. However, even among the “good” operators, there are those who are outstanding (true event rate of 1% or $p = 0.01$), those who are very good (true event rate of 2% or $p = 0.02$), and those who just barely meet the criteria (true event rate of 3% or $p = 0.03$). Then suppose we require each operator to perform the procedure on a series of 24 patients. Then

out of 100 potential outstanding operators who each perform 24 procedures, on average how many operators will have no complications, how many 1 complication, how many 2 complications, and so forth? Similarly, what would the outcome be for 100 very good operators or 100 who just meet the criteria? For each of these “good” type of operators, Table 1 shows the expected number of operators who have no events, 1 event, 2 events, 3 events, 4 events, 5 events or 6 or more events. As can be seen, if we set the criteria for inclusion of an operator in the study as having 0 or 1 events, then we will only inappropriately exclude 2 of the outstanding operators; however, we will inappropriately exclude 8 of the very good operators, and inappropriately exclude 16 of the operators who barely meet inclusion criteria. As such, it would seem to meet the needs of the study that operators with 2 or fewer events be included in the study (inappropriately excluding no outstanding operators, a single good operator, and only 3 operators just meeting criteria). There appears to be some success in including operators meeting this criterion; however, to reliably do so we would need to accept operators with event rates of 2 or fewer among 24 patients, an event rate of 8.3% or lower, not our ideal of $< 3\%$.

Approaches to discard the bad apples

From a patient safety perspective, an even more important goal than including the operators meeting criteria (keeping good apples) is excluding the operators who fail to meet criteria (throwing away bad apples). Like there are a spectrum of abilities among the good operators, there are likely operators who barely fail to meet criteria (true event rate of 4%, or $p = 0.04$), those that are moderately unacceptable (say a true event rate of 8% or even 10%, that is $p = 0.08$ or $p = 0.10$), and those that are fundamentally unacceptable (say a true event rate of 20% or $p = 0.20$). Assuming that we have 100 of each of these types of operators, the expected number will have a range of events as shown in Table 1. For those barely not meeting the criteria (4% complication rate), if we set the criteria as above to be having 2 or fewer events, then we would include 94 out of 100 operators that we wish to exclude. Perhaps more disturbing, we would expect to include 70 out of 100 operators with a true 8% event rate (more than twice the practice guideline-defined limit), and 56 of 100 operators with a true 10% event rate (more than three-times the limit). In fact, the exclusion of operators with high event rates can only be achieved if the true event rate is extraordinarily high, such as in the neighborhood of a true 20% complication rate. Even if we set the rate to having no complications, we would still include 23% of operators with a 6% event rate (twice the limit), 14% of operators with an 8% complication rate and 8% of operators with a 10% complication rate. These data suggest that there is little hope of reliably excluding the “bad apples” from the study (unless they are truly unacceptable), and even setting the goal of no events among 24 cases will tend to include as many as 23% of those with twice the acceptable event rate.

Assessment of the Required Size of the Series to Make Reliable Decisions

Figure 1 shows the error rate as a function of the number of patients in the series (using the “optimal” threshold for inclusion of an operator). It would require a series of 240 patients to have an error rate below 10% to distinguish between good operators with a 2% complication rate and poor operators with a 6% complication rate. With 240 patients, the optimal rule would be to include operators with 8 (3.3%) or fewer complications, and at this threshold

5.4% of good operators would be excluded and 4.6% of poor operators would be included (error rate of 10%). If the goal was to distinguish between good operators with a 2% complication rate and poor operators with an 8% complication rate, a series of 120 patients would achieve a 10% error rate. To distinguish between operators with a 2% versus 10% event rate, 83 patients would achieve a 10% error rate.

Discussion

By peer review of a prospective consecutive case series averaging 24 patients per stent operator, CREST was shown to successfully choose operators with very low complication rates. This process ultimately excluded 48% of operators who had applied to be part of the study.⁶ One would presume that the reason for the overall low event rate for the study was that these excluded operators included many “bad apples” that would not truly meet criteria. Likewise, one would presume that the 52% who were accepted for the study included many “good apples” who do meet the criteria. Hence, the randomized trial phase of CREST suggests a proven ability to discriminate between the good and bad apples.

However, looking at complication rates in a series of this size could have provided only modest guidance to select those with good event rates. Recalling that the published criteria for selection were performance (i.e., periprocedural stroke and death rate), volume and peer review of technique, this would suggest that either the assessment of volume or the peer review of technique must have been the primary drivers of the success of the study. In their review of individual cases, the Interventional Management Committee noted a diverse array of “red flags” indicating that specific operators should be excluded such as an administration of 12,000 IU Heparin, needing to use complex sheath access techniques or guiding catheters, utilizing undersized stent diameters or excessive use of long stents, post-dilation with >5 mm balloon or doing multiple inflations, failing to urgently lower blood pressure (if it does not fall) after stent placement, or an protection device time >15 min. Identifying these and other potential red flags requires the judgment of individuals with appropriate high levels of experience and expertise regarding the quality of the operator and the procedure, and these skills not easily quantified. As such, the study is grateful for what must be the skillful subjective review performed by the Interventional Management Committee (members shown in Table 2).

The central (statistical) problem in the assessment of performance is that it is difficult to reliably detect differences in binomial outcomes (in this case, in the neighborhood of $p = 0.03$). We demonstrate that to achieve an acceptably low error rate (sum of exclusion of good operators plus inclusion of poor operators) a series of 240 patients would be required. For a very active operator performing a procedure a week, this would require more than 4.5 years to qualify for the study. Even with the entire CREST experience of 594 CAS procedures in asymptomatic patients, and the lowest reported event stroke and death rate in any randomized trial (2.5%), the 95% confidence intervals for the event rate still extend up to 4.1%. Hence even with a series of 594 patients we cannot definitively state that guideline level of 3.0% was met. This underscores not only challenges of selecting operators for the CREST trial, but also the challenge of following the advice of Consumer Reports in the selection of surgeons to ask “What is your success, failure, and complication rates?”⁷

A further complication is that patients may not be assigned to surgeons/operators in a random manner, but rather the more challenging patients are likely assigned to the “better” surgeon/operators. General approaches for adjustment for case-mix have been suggested;^{13, 14} however, general “rules of thumb” require 10 to 20 events per variable considered for case-mix adjustment. Even with 594 asymptomatic patients in CREST treated with CAS, there were only 15 periprocedural events,¹ making modeling for risk adjustment practically infeasible. Basically, the inability for case-mix modeling arises from the same basic issue as the challenge of statistical assessment of event rates ... there is not a sufficient number of events for reliable statistical analysis.

Finally, the IMC in CREST was charged with selection of outstanding operators to ensure that stenting would be safely done in the randomization phase of the trial. It should be noted that minimization of risk from a procedure can be achieved by means beyond selection of skilled operators. There is growing interest and experience with the use of simulation as a means of reducing procedural risk.¹⁵ Additionally, while attention usually focuses on individual apples, one should not ignore the importance of the apple tree, i.e., the environment in which the procedure is performed. For example, it is now considered standard of care for the cardiac operating room to have ongoing quality improvement projects, checklists, briefings, and formal handoff protocols.¹⁶ Presumably, these elements of care and process would also be of benefit in the setting of CAS.

In conclusion, CREST seems to have selected appropriate and skillful operators by a combination of reviewing performance, volume and a subjective review of technique. Herein, we suggest that while there is some information to be gained in the statistical review of performance, the criteria for selection of the operator is so stringent that the information from a small case series is quite limited. It is uncomfortable for the statistician coauthors to admit that because the CREST review was successful, the greater value in the review would seem to have been from the subjective review of techniques or other sources of information (such as reputation). Statistics is a critically powerful tool to address many issues; however, unless very large case series can be constructed, we suggest that it may not be the best tool for evaluating the ability of individual operators to meet stringent (< 3%) performance rates, and results from small series (e.g. 24) be interpreted with caution. Sometimes you have to pick the good apples and throw away the bad by just looking at them instead of using a statistical assessment.

Acknowledgements

Sources of Funding: CREST was funded by the National Institute of Neurological Disorders and Stroke (US) [R01 NS038384] with additional support from Abbott Vascular, Inc.

References

1. Brott TG, Hobson RW 2nd, Howard G, Roubin GS, Clark WM, Brooks W, et al. Stenting versus endarterectomy for treatment of carotid-artery stenosis. *N Engl J Med*. 2010; 363:11–23. [PubMed: 20505173]
2. Furie KL, Kasner SE, Adams RJ, Albers GW, Bush RL, Fagan SC, et al. Guidelines for the prevention of stroke in patients with stroke or transient ischemic attack: A guideline for healthcare

- professionals from the american heart association/american stroke association. *Stroke*. 2011; 42:227–276. [PubMed: 20966421]
3. Brott TG, Halperin JL, Abbara S, Bacharach JM, Barr JD, Bush RL, et al. 2011 ASA/ACCF/AHA/AANN/AANS/ACR/ASNR/CNS/SAIP/SCAI/SIR/SNIS/SVM/SVS guideline on the management of patients with extracranial carotid and vertebral artery disease: Executive summary: a report of the American College of Cardiology Foundation/American Heart Association Task Force on Practice Guidelines, and the American Stroke Association, American Association of Neuroscience Nurses, American Association of Neurological Surgeons, American College of Radiology, American Society of Neuroradiology, Congress of Neurological Surgeons, Society of Atherosclerosis Imaging and Prevention, Society for Cardiovascular Angiography and Interventions, Society of Interventional Radiology, Society of Neurointerventional Surgery, Society for Vascular Medicine, and Society for Vascular Surgery. *Vasc Med*. 2011; 16:35–77. [PubMed: 21471149]
 4. Goldstein LB, Bushnell CD, Adams RJ, Appel LJ, Braun LT, Chaturvedi S, et al. Guidelines for the primary prevention of stroke: A guideline for healthcare professionals from the American Heart Association/American Stroke Association. *Stroke*. 2011; 42:517–584. [PubMed: 21127304]
 5. Hobson RW 2nd, Howard VJ, Roubin GS, Ferguson RD, Brott TG, Howard G, et al. Credentialing of surgeons as interventionalists for carotid artery stenting: Experience from the lead-in phase of CREST. *J Vasc Surg*. 2004; 40:952–957. [PubMed: 15557910]
 6. Hopkins LN, Roubin GS, Chakhtoura EY, Gray WA, Ferguson RD, Katzen BT, et al. The Carotid Revascularization Endarterectomy versus Stenting Trial: Credentialing of interventionalists and final results of lead-in phase. *J Stroke Cerebrovasc Dis*. 2010; 19:153–162. [PubMed: 20189092]
 7. How to Find the Right Surgeon. [ConsumerReports.org](http://www.consumerreports.org/cro/2012/04/how-to-find-the-right-surgeon/index.htm). 2010 <http://www.consumerreports.org/cro/2012/04/how-to-find-the-right-surgeon/index.htm>.
 8. Hobson RW 2nd, Howard VJ, Roubin GS, Brott TG, Ferguson RD, Popma JJ, et al. Carotid artery stenting is associated with increased complications in octogenarians: 30-day stroke and death rates in the CREST lead-in phase. *J Vasc Surg*. 2004; 40:1106–1111. [PubMed: 15622363]
 9. Voeks JH, Howard G, Roubin GS, Malas MB, Cohen DJ, Sternbergh WC 3rd, et al. Age and outcomes after carotid stenting and endarterectomy: The carotid revascularization endarterectomy versus stenting trial. *Stroke*. 2011; 42:3484–3490. [PubMed: 21980205]
 10. Carotid Stenting Trialists C, Bonati LH, Dobson J, Algra A, Branchereau A, Chatellier G, et al. Short-term outcome after stenting versus endarterectomy for symptomatic carotid stenosis: A preplanned meta-analysis of individual patient data. *Lancet*. 2010; 376:1062–1073. [PubMed: 20832852]
 11. Howard VJ, Voeks JH, Lutsep HL, Mackey A, Milot G, Sam AD 2nd, et al. Does sex matter? Thirty-day stroke and death rates after carotid artery stenting in women versus men: Results from the Carotid Revascularization Endarterectomy versus Stenting Trial (CREST) lead-in phase. *Stroke*. 2009; 40:1140–1147. [PubMed: 19211486]
 12. Brott TG, Halperin JL, Abbara S, Bacharach JM, Barr JD, Bush RL, et al. 2011 ASA/ACCF/AHA/AANN/AANS/ACR/ASNR/CNS/SAIP/SCAI/SIR/SNIS/SVM/SVS guideline on the management of patients with extracranial carotid and vertebral artery disease: Executive summary: A report of the American College of Cardiology Foundation/American Heart Association Task Force on Practice Guidelines, and the American Stroke Association, American Association of Neuroscience Nurses, American Association of Neurological Surgeons, American College of Radiology, American Society of Neuroradiology, Congress of Neurological Surgeons, Society of Atherosclerosis Imaging and Prevention, Society for Cardiovascular Angiography and Interventions, Society of Interventional Radiology, Society of NeuroInterventional Surgery, Society for Vascular Medicine, and Society for Vascular Surgery. Developed in collaboration with the American Academy of Neurology and Society of Cardiovascular Computed Tomography. *Catheter Cardiovasc Interv*. 2013; 81:E76–E123. [PubMed: 23281092]
 13. Krumholz HM, Brindis RG, Brush JE, Cohen DJ, Epstein AJ, Furie K, et al. Standards for statistical models used for public reporting of health outcomes: an American Heart Association Scientific Statement from the Quality of Care and Outcomes Research Interdisciplinary Writing Group: cosponsored by the Council on Epidemiology and Prevention and the Stroke Council. Endorsed by the American College of Cardiology Foundation. *Circulation*. 2006; 113:456–462. [PubMed: 16365198]

14. Katzan IL, Spertus J, Bettger JP, Bravata DM, Reeves MJ, Smith EE, et al. Risk adjustment of ischemic stroke outcomes for comparing hospital performance: A statement for healthcare professionals from the American Heart Association/American Stroke Association. *Stroke*. 2014; 45:918–944. [PubMed: 24457296]
15. Buckley CE, Kavanagh DO, Traynor O, Neary PC. Is the skillset obtained in surgical simulation transferable to the operating theatre? *Am J Surg*. 2014; 207:146–157. [PubMed: 24238602]
16. Wahr JA, Prager RL, Abernathy JH 3rd, Martinez EA, Salas E, Seifert PC, et al. Patient safety in the cardiac operating room: Human factors and teamwork: A scientific statement from the American Heart Association. *Circulation*. 2013; 128:1139–1169. [PubMed: 23918255]

Figure 1.

Minimum sum of the percent of “good operators” (with an event rate of 2%) excluded from participation and the percent of “bad operators” (with event rates of 6%, 8% or 10%) included for participation as a function of the number of patients in the series.

Distribution of operators on the anticipated number of poor outcomes as a function of their true chance of having a poor outcome on a particular patient.

Table 1

# of Poor Outcomes	True Chance of a poor outcome									
	Acceptable Operators ("Good Apples")				Unacceptable Operators ("Bad Apples")					
	1%	2%	3%		4%	6%	8%	10%	20%	
0	79	62	48		38	23	14	8	0	
1	19	30	36		38	35	28	21	3	
2	2	7	13		18	25	28	27	8	
3	0	1	3		5	12	18	22	15	
4	0	0	0		1	4	8	13	20	
5	0	0	0		0	1	3	6	20	
6+	0	0	0		0	0	1	3	34	

Of 100 operators with a specific "true" chance of a bad outcome, the anticipated number with a particular number of bad outcomes. For example, if there were 100 operators with a 3% true chance of a bad outcome, we would anticipate 48 operators to have no bad outcomes, 36 to have 1 bad outcome, 13 to have 2 bad outcomes, and 3 to have 3 bad outcomes.

Table 2

Members of the CREST Interventional Management Committee

Investigator	Title, Specialty
Gary Roubin, MD., Ph.D.	Committee Co-Chair, Interventional Cardiology
Robert Ferguson, M.D.	Committee Co-Chair, Interventional Neuroradiology
Jonathan Goldstein, M.D.	Member, Interventional Cardiology
William Gray, M.D.	Member, Interventional Cardiology
Robert W. Hobson, M.D. (1999 – 2008)	Member, Vascular Surgery
L. Nelson Hopkins, M.D.	Member, Neurosurgery
William Morrish, M.D.	Member, Interventional Neuroradiology
Barry T. Katzen, M.D.	Member, Interventional Radiology
Kenneth Rosenfield, M.D.	Member, Interventional Cardiology
Thomas G. Brott, M.D.	Ex-Officio, Neurology
Elie Chakhtoura, M.D.	Ex-Officio, Interventional Cardiology