



Enhancing the investment returns  
of non-profit organizations

# COMMENTARY

A Quarterly Report of TIFF EDUCATION FOUNDATION

JUNE 30, 2004

## BOARD MEMBERS

TIP	<b>Suzanne Brenner</b> Senior Investment Officer Metropolitan Museum of Art New York, NY
TAS TIP	<b>Richard Flannery</b> CEO of TAS President and CEO of TIP Charlottesville, VA
TEF TIP	<b>Harry Hoffman</b> Chief Investment Officer Mayo Foundation Rochester, MN
TEF TIP	<b>Sheryl Johns</b> Chief Financial Officer Houston Endowment Inc. Houston, TX
TAS	<b>Michael McCaffery</b> President and CEO Stanford Management Co. Menlo Park, CA
TAS TIP	<b>William McLean</b> Chief Investment Officer Northwestern University Evanston, IL
TEF TAS	<b>Jane Mendillo</b> Chief Investment Officer Wellesley College Wellesley, MA
TAS	<b>Jack Meyer</b> President and CEO Harvard Management Co. Boston, MA
TAS	<b>Thruston Morton</b> President and CEO Duke Management Co. Durham, NC
TEF TAS TIP	<b>David Salem</b> President and CIO of TAS President of TEF CIO of TIP Charlottesville, VA
TAS	<b>Ellen Shuman</b> Chief Investment Officer Carnegie Corp of New York New York, NY
<i>Legend</i>	
TEF	<b>TIFF Education Foundation</b> a tax-exempt private operating foundation
TAS	<b>TIFF Advisory Services</b> the registered investment advisor for all TIFF vehicles
TIP	<b>TIFF Investment Program</b> a mutual fund family open to charities only

## HIGHLIGHTS

- This quarter's *Commentary*, which marks the start of the second decade of TIFF's education and research program, discusses the program's latest and perhaps most ambitious initiative: to develop a truly integrated approach to investment policy formulation for endowed charities.
- Crafting sensible investment policies for endowed charities is no easy task, for at least two reasons: (1) no one really knows how specific asset classes will behave in the future and (2) it's tough for investment committees to determine *ex ante* the types and degrees of risk they're willing to incur.
- Fiduciaries intent on using quantitative tools to make investment policy choices in a rigorous manner need not despair. As an alternative to mean variance optimization (MVO) — a quantitative tool for making such choices that is as flawed as it is popular — fiduciaries can employ Monte Carlo Simulation (MCS), which is superior to MVO in handling three crucially important aspects of the policy planning challenge: complex return distributions, within-horizon risk, and relative return risk.
- The data germane to truly informed policy choices can be usefully subdivided into two groups: generic and institution-specific. Examples of the former include asset class returns, volatilities, and correlations plus a host of other data that are relevant and reasonably common to most investors (e.g., rebalancing costs, illiquidity restrictions, fees, and the like). The real challenge is to get individual trustee groups to articulate idiosyncratic return goals and risk tolerances in a clear and internally consistent manner.

## ABOUT TIFF

*Origins.* In 1991, a network of foundations founded an investment cooperative whose organizational structure and eligibility criteria have evolved over time but whose core mission has not. Known colloquially as TIFF, the cooperative seeks to improve the investment returns of endowed charities by making available to them a series of multi-manager investment vehicles plus resources aimed at enhancing fiduciaries' knowledge of investing. The cooperative comprises three regulated entities at present: a tax-exempt private operating foundation whose d/b/a (TIFF Education Foundation) is more descriptive of its focus on education and research than its

formal legal name (The Investment Fund for Foundations); the TIFF Investment Program (TIP), a SEC-regulated mutual fund family; and TIFF Advisory Services (TAS), a taxable non-stock corporation and SEC-registered investment advisor that administers all investment vehicles bearing the TIFF name. As noted at left, there is substantial but not complete overlap among these three entities' boards, all of whose members except Richard Flannery and David Salem serve as unpaid volunteers.

*Inquiries.* For more information, please call TIFF at 434-817-8200 or visit [www.tiff.org](http://www.tiff.org).

*Turning Point.* With this report's issuance, TIFF's publications program enters its second decade. For better or worse, the program's primary goal remains unchanged from its initial adoption 10 years ago: to help trustees and officers of endowed charities perform more effectively their investment-related duties. Believers as we are in the principle that no man should serve as judge in his own case, your editor will not opine here on the extent to which our publications program has achieved its primary goal. But we won't shy from claiming that it has achieved more or less fully certain subordinate goals, the most conspicuous being (1) to provoke discussion among trustee groups about endowment return goals and the means most conducive to achievement of these goals and (2) to promote awareness of the unfailingly interesting and (for the last 85 years) unfailingly unsuccessful efforts of your editor's favorite pro sports team to "win it all." We won't mention the team by name here because doing so would violate from the get-go one of several resolutions your editor has adopted as he enters his second decade as TIFF's chief scribe. The first such resolution is to do a better job of achieving with respect to TIFF's publications program a goal pursued by all astute fiduciaries on behalf of the portfolios that they steward: an appropriate degree of diversification. Appropriate in this context means that TIFF's publications henceforth should reflect not a single voice as they have historically but — subject to this editor's acute affinity for arguing about investment policy issues in public and in writing — voices other than that of TIFF's president, if not carrying the main tune then at the least providing counterpoint. Appropriate as used here also means that TIFF's publications henceforth should contain allusions to sports other than the game played in Fenway Park and to teams other than the one which calls that ballpark home.

*Noteworthy Initiative.* Not wishing to delay for a moment implementation of the aforementioned new and more diversified approach to TIFF's publications efforts, your editor invited Jesse Barnes to draft the remainder of this quarter's *TIFF Commentary*. A PhD candidate in financial economics at Harvard University, Jesse has been working closely with TIFF on an initiative that we hope will advance materially our mission of helping endowed charities earn enhanced investment returns. This initiative's aim is to develop a truly integrated approach to investment policy formulation that exploits efficiently the robust and ever-changing set of investment opportunities that endowed charities confront while ensuring that the choices they make among them are sensitive to individual institutions' idiosyncratic goals and risk tolerances. The paragraphs that follow, the drafting of which was done by Jesse, outline how we intend to tackle this concededly daunting task.

*No Easy Task.* As readers of these commentaries know only too well, crafting sensible investment policies for endowed charities is no easy task, for at least two reasons: (1) charlatans aside, no one really knows precisely how specific asset classes or subclasses will behave in the future and (2) as tough as it is for individual investors to determine *ex ante* the types and degrees of risk they're willing to incur, it's tougher by far for groups of human beings (especially groups of lay trustees meeting on a part-time basis) to do so. Of course, absent a clear and common understanding of risk tolerances — which is to say the point past which a trustee group experiencing poor returns is likely to jettison seemingly failed policies — an investment committee cannot determine intelligently what a realistic return goal might be. Nor can it decide among the potentially infinite means of pursuing whatever return goal it ultimately adopts. In light of the difficulties just described, many investment committees make little or no use of quantitative tools — tools that if used properly can help them wring order out of seeming chaos. Moreover, to the extent that fiduciaries utilize such tools — literally computer-based models that exploit the many advances in both modeling techniques and hardware design in recent years — they tend to use them for what might be called defensive purposes, i.e., to produce quantitative justifications of policy choices arrived at by qualitative means. There's nothing illegal about this *modus operandi*, of course, although trustees who employ it must guard vigilantly against adoption of policies that could prove wholly unresponsive to an institution's long-term needs. That said, the dangers of relying too heavily on quantitative techniques when fashioning investment policies arguably exceed the dangers of using them not at all: as was discussed in detail in the *TIFF Commentary* dated March 31, 1999, the quantitative tools we're alluding to here have serious limitations. Yet these tools can be of substantial help in the policy formulation process, especially if used in a manner that is duly sensitive to their inherent limitations.

*State of the Art.* The *TIFF Commentary* just mentioned focuses on the quantitative tool most commonly used by institutional investors when formulating investment policies: Mean Variance Optimization (MVO). Although MVO users can mitigate its most glaring defects — its tendency to assume that future returns will resemble historical averages (don't count on it) and will be "normally" distributed (ditto) — they cannot force MVO models to do something they're inherently unable to do, i.e., help fiduciaries analyze the implications of alternate policies via **intuitively appealing** means. As its name indicates, MVO focuses exclusively on alternate portfolios' mean or average returns and variances. Although these two attributes are hardly irrelevant

to investment policy formulation, they're of limited utility: after all, people have drowned in rivers that are knee-high on average, and we've yet to meet a trustee group that eschews consciously **upward** variance. (Technically, variance and volatility mean different things — variance is actually the standard deviation of a return series multiplied by itself or squared — but for our purposes here we can treat these two phenomena as essentially synonymous.) Indeed, variance *per se* merits little attention when more intuitively appealing measures of risk such as drawdowns and return shortfalls are available. Fortunately, fiduciaries intent on using quantitative tools to make investment policy choices in a more rigorous and unemotional manner than might prevail in the absence of such tools need not despair. In addition to MVO, which for all of its defects can aid the policy formulation process by helping fiduciaries identify alternate policy portfolios worthy of “testing” via other means, fiduciaries can use another highly evolved quantitative technique to aid the policy planning process: Monte Carlo Simulation (MCS).

*A Better Way.* In the context of investment policy formulation, there are at least three respects in which MCS tends to be superior to MVO and related variance-based means of constructing policy portfolios (collectively referred to henceforth as traditional analysis or TA):

- *Complex Return Distributions.* TA begins with the assumption that asset returns are normally distributed. Empirical research continues to confirm what seasoned investors already knew about this assumption: it's about as sound as Jennifer Lopez's marital vows. It's unsound because financial markets display greater kurtosis than normal (in plain English, they display a wider range of outcomes than one finds or rather used to find in leading colleges' grade distributions) and more negative skewness also (i.e., distributions comprising more outcomes far below the mean or average than there are outcomes above the mean, an attribute whose conspicuousness in individual colleges' grading patterns tends to be inversely related to each school's prestige). Moreover, the statistical properties of asset classes and subclasses tend to be highly unstable through time. This is especially true respecting a set of properties that is hugely important in portfolio construction: asset class correlations. Such correlations are highly problematic because history teaches us that they differ hugely depending on whether the market environment being analyzed is “normal” or not. We'll leave for a later publication a detailed discussion of how one determines whether a particular environment is normal or not; the key point is that truly robust approaches to policy formulation should take due note of the distinction between extreme environments and all others and permit fiduciaries to “stress test” alternate

portfolios under both sets of conditions. MCS fits this bill to a far greater extent than MVO. Truly robust approaches to policy formulation should also reflect the tendency of asset returns to exhibit “mean” reversion, i.e., to revert over time to levels compatible with the way capitalism is known to function, despite prolonged and pronounced departures from such levels (*cf.* the Internet “bubble” of the late 1990s). Again, MCS scores much higher than MVO on this crucial metric.

- *Within-Horizon Risk.* As past or present users of TA will recall, its end product is a comparison of the returns that alternate portfolios are expected to produce to their riskiness as measured most typically by standard deviation (SD), with both such attributes — returns and SDs — being expressed over the full time horizons being analyzed (e.g., 20 or 30 years). This is both dangerous and naïve. It's dangerous because for most lay users of MVO a portfolio's SD doesn't illuminate well the extremity of certain potential outcomes. And it's naïve because if and when these extreme events unfold they could cause fiduciaries to “bail,” i.e., to abandon otherwise sensible policies at inopportune times. Differently put, computer-based approaches to policy formulation should consider not only how much wealth a given portfolio is likely to produce at the end of the holding period being considered but also the potential changes in such wealth “along the way.” The longer one extends the holding period or time horizon used to analyze alternate policy choices, the higher the likelihood of experiencing a drawdown that will test and potentially shatter fiduciaries' nerves, because the more time markets will have to misbehave. The opposite is true for end-of-horizon risk, of course, which declines as the time period being analyzed is lengthened, i.e., the longer one's planning horizon is, the safer it is to assume that long-term averages will govern the end result. MCS trumps MVO with respect to within-horizon risk because MCS facilitates infinitely more detailed analyses of within-horizon portfolio changes.

- *Relative Return Risk.* Investment committees may be concerned first and foremost with absolute returns — how much endowment values grow over whatever time horizons such committees employ and how vulnerable such values are to interim losses or drawdowns along the way — but they also are concerned with relative risks, by which we mean the risks of underperforming relevant market indices or peer institutions. Peer group risk can be especially important for publicly supported charities whose potential donors want to see continuing evidence that the organizations they're being asked to help underwrite are well managed in all respects. The most robust asset allocation models (meaning: well-designed MCS packages) take relative risks into account by analyzing them not merely in isolation but in conjunction

with so-called absolute risks, the ultimate aim being to ensure that a given policy portfolio being considered for adoption doesn't entail too high a probability of leaving its stewards "wrong and alone," i.e., with a portfolio that produces returns that are poor when measured both absolutely and relatively. As an aside, our favorite essay on the "wrong and alone" problem was authored by Mark Kritzman of Windham Capital Management, with whom TIFF staffers have had the privilege of spending considerable time in connection with the initiative discussed in this *Commentary*.

*Inner Workings.* As the TIFF team laboring on a fully integrated approach to policy formulation goes about its work, we are taking due care to catalog all data germane to truly informed policy choices. These data can be usefully subdivided into two groups: generic and institution-specific. Examples of the former include asset class returns, volatilities, and correlations plus a host of other data that are relevant and reasonably common to most investors (e.g., rebalancing costs, illiquidity restrictions, fees, and the like). These generic data aren't always readily available, but with adequate Googling skills and a bit of ingenuity enough of them can be obtained to make MCS models worth using. The real challenge in the data collection process is to get individual trustee groups to articulate their return goals and risk tolerances in a clear and internally consistent manner. To be sure, it's not very hard to get trustees or perhaps staff to whom they've delegated such work to codify such straightforward parameters of a well-organized endowment management program as (1) projected external cash flows (i.e., future gifts flowing into the top of the reservoir so to speak as well as future withdrawals flowing out of it), (2) any institution-specific sensitivities that might logically affect investment policy choices (e.g., material linkages between an institution's overall financial condition and the evolving fortunes of particular industries, companies, political jurisdictions, or the like), or (3) the extent to which debt service obligations or covenants circumscribe investment policy options. However, for reasons outlined in an essay on behavioral finance published in the TIFF quarterly report dated March 31, 1996), it tends to be difficult for individuals and **extremely** difficult for groups of individuals to articulate in an internally consistent manner their aversion to various risks.

*To Be Continued.* The difficulties just alluded to bring us back to our starting point, which was to ask whether recent advances in both computer technology and financial economics can be harnessed in a manner that helps endowment fiduciaries wring a semblance of order out of the potentially chaotic environment in which they necessarily go about their work. The work on investment policy formulation being undertaken currently under

TIFF's auspices has the potential to provide an affirmative answer to the question. More particularly, it has the potential to provide investment committees with the tools needed to explore iteratively alternate formulations of both ends and means, by which we mean targeted spending rates, return goals, and risk parameters on the one hand and alternate policy portfolios (i.e., normal asset class weights plus permissible deviations therefrom) on the other. Assessing such choices **iteratively** is not only feasible using relatively low-cost technologies, it's arguably essential due to limitations on how rapidly and accurately even the most acutely gifted human thinkers can compute probabilities. An apparent exception to this generalization was Steve Young, the preternaturally gifted quarterback who brought distinction to the undergraduate *alma mater* that he shares with Jesse Barnes (Brigham Young) and who proved exceptionally gifted at anticipating what opposing defenders would do if he and his teammates moved in certain directions. There's no particular reason for mentioning football here except to underscore the fact that this may be the only TIFF *Commentary* ever published that doesn't mention by name the favorite spectator sport of TIFF's chief scribe! In other words, we've honored in this essay the resolutions set forth in its preface and will honor too our pledge to keep TIFF members apprised of our progress in the policy planning initiative described above. ■

## MEMBERSHIP SUMMARY

	Number of Members
<b>TIFF Membership</b>	<b>444</b>
▪ Private Foundations	183
▪ Community Foundations	38
▪ Educational Institutions	41
▪ Other Endowed Charities	182



THE INVESTMENT FUND FOR FOUNDATIONS  
*Enhancing the investment returns of non-profit organizations*

590 Peter Jefferson Parkway, Suite 250  
Charlottesville, Virginia 22911

Phone: 434-817-8200

Fax: 434-817-8231

Website: [www.tiff.org](http://www.tiff.org)

*Electronic mail inquiries:*

Services offered by TIFF: [info@tiff.org](mailto:info@tiff.org)

Member-specific data: [memberservices@tiff.org](mailto:memberservices@tiff.org)

Manager selection procedures: [managers@tiff.org](mailto:managers@tiff.org)

*For further information about any of TIFF's services, please contact TIFF using the coordinates furnished above.*