A Systematic Economic Approach to Evaluating Public Investment in Observations for Weather Forecasting

REBECCA E. MORSS AND KATHLEEN A. MILLER

National Center for Atmospheric Research, Boulder, Colorado*

MAXINE S. VASIL

University of Colorado, Boulder, Colorado

(Manuscript received 26 March 2004, in final form 15 July 2004)

ABSTRACT

Observations of the current state of the atmosphere are a major input to production of modern weather forecasts. As a result, investments in observations are a major component of public expenditures related to weather forecasting. Consequently, from both a meteorological and societal perspective, it is desirable to select an appropriate level of public investment in observations. Although the meteorological community has discussed optimal investment in observations for more than three decades, it still lacks a practical, systematic framework for analyzing this issue. This paper presents the basic elements of such a framework, using an economic approach. The framework is then demonstrated using an example for radiosonde observations and numerical weather forecasts. In presenting and demonstrating the framework, the paper also identifies gaps in existing knowledge that must be addressed before a more complete economic evaluation of investment in observations can be implemented.

1. Introduction

Atmospheric and related observations are a major component of public expenditures for weather prediction (OFCM 2002; NOAA 2002). These observations are used in real time to produce operational weather forecasts, and for research and development to improve future forecasts. The resulting weather forecasts benefit society in many ways, by helping protect lives and property from hazardous weather-related events, aiding the general public in everyday planning, and facilitating more efficient management of transportation, agriculture, energy, and other economic activities.

Because observations have substantial socioeconomic costs and benefits, meteorologists have long recognized the importance of balancing benefits with costs when allocating public resources to observations. For example, in 1967, during international planning for the First Global Atmospheric Research Program (GARP) Global Experiment (FGGE), a World Meteorological Organization (WMO 1967, p. v) report stated:

> Meteorologists have long been in doubt about the optimum density, in both time and space, of observations of the different meteorological parameters necessary for various meteorological functions . . . On the one hand, a paucity of data would result in unsatisfactory output products while, at the other end of the scale, an over-dense network could well pass the point of diminishing returns where further increases would not add economically to forecast accuracy.

After FGGE, in 1980, a National Research Council (NRC) panel recommended that the National Weather Service (NWS) "project benefits and costs of providing alternative levels of weather and hydrological data and information services as a basis for rational allocation of future NOAA resources" (NRC 1980, p. 4). In 1997, the U.S. Weather Research Program (USWRP) convened a scientific panel specifically "to consider ways to promote research that seeks to determine implementations of observing systems that are optimal for weather prediction . . . An 'optimal' measurement system is considered to be one that maximizes the ratio of societal benefit to overall cost" (Emanuel et al. 1997, p. 2859; Emanuel et al. 1995).[1] And over the next decade, a

* The National Center for Atmospheric Research is sponsored by the National Science Foundation.

Corresponding author address: Dr. Rebecca E. Morss, NCAR MMM/ESIG, P.O. Box 3000, Boulder, CO 80307.
E-mail: morss@ucar.edu

[1] Note that we argue that maximizing net benefit = total benefit − total cost is a more appropriate criterion; see section 2.

primary goal of THORPEX: A Global Atmospheric Research Program, an international weather prediction research and field program being developed as a follow-on to GARP/FGGE, is to evaluate the socioeconomic costs and benefits of modifying the global observing system (Shapiro and Thorpe 2004).

Despite this long-term interest, the meteorological community still lacks a practical, systematic framework for analyzing the most appropriate level and mix of public investment in observations for weather prediction. Without such a framework, meteorologists continue to have difficulty providing useful, credible guidance to policy makers on allocating observational resources, and policy makers continue to have difficulty effectively allocating resources for the benefit of both meteorology and society. As pressure on public funding has grown in the United States and internationally, meteorologists are also facing increasing requests from policy makers to justify government expenditures on observations and other aspects of meteorology (e.g., Mikulski 1994; Doswell and Brooks 1998; Zillman and Freebairn 2001). As a step toward ameliorating these difficulties, this paper presents a basic economic framework for evaluating public investments in weather prediction.

Note that public policy decisions involve many considerations in addition to economics. Consequently, the framework presented here does not, on its own, provide a complete guide for decision-making. Rather, the framework provides a structure for organizing economic considerations in observing network design. The information generated can then serve as useful input to decisions about observing networks and weather prediction, to be considered in conjunction with other relevant factors.

The framework treats observations as a variable input in the real-time production of weather forecasts, comparing economic costs and benefits to derive 1) the economically optimal level of weather forecast production and 2) the economically optimal mix of observational inputs to use in producing that level of forecasts. Much of the framework builds on that presented by Zillman and Freebairn (2001), Freebairn and Zillman (2002), and Zillman (2002) (hereafter referred to collectively as ZF) in their derivation of the optimal level of meteorological services. Here, that framework is reviewed, then extended to analyze the optimal level of inputs to forecast production. For ease of exposition, we consider the publicly produced basic forecast services in the "public goods" category discussed by ZF.

As noted by ZF, we currently lack the information required to apply the framework to the real weather observing and prediction system. Instead, the paper demonstrates the framework using an example for radiosonde observations and numerical weather forecasts in the United States, based on previously published results. Although the example is sufficiently simplified that the detailed results cannot be transferred directly to the real world, the results suggest that increasing the number of observations would likely increase net societal benefit. The example is also used to identify gaps in existing knowledge that must be addressed before a more complete economic analysis can be implemented.

Although introduced here in the context of observing network design, the economic framework can also be applied to other public resource allocation issues in atmospheric and related sciences. It could be used, for example, to analyze the level of public investment in other inputs to the weather prediction process, such as model resolution and ensemble size, or to analyze trade-offs among these inputs. It can also be applied in other earth observing and prediction contexts, such as El Niño–Southern Oscillation or streamflow prediction.

Section 2 presents the framework and underlying economic concepts. While the concepts presented are standard economic theory, they provide the necessary background for the primary focuses of the paper: applying the concepts to the meteorological observing and prediction system, and suggesting future applications. In section 3, the framework is demonstrated using a simplified example of the weather observing and prediction system. Section 4 discusses limitations of the framework and example, including gaps in existing knowledge, and section 5 provides a summary.

2. Conceptual economic framework[2]

Weather forecasts can be considered a product, where increasing forecast skill corresponds to increasing the level of output of the product. Weather forecasts are produced by combining various inputs, including different types of observations, computer hardware and software, human labor, and meteorological knowledge. Forecasts at a particular skill level can be produced with different mixes of these inputs. Conceptually, this raises two questions: 1) What is the optimal level of forecast skill? and 2) What is the optimal mix of inputs to use in producing that level of skill? To address these questions, this section presents a standard economic analysis of the optimal level of output of a good or service, and of the optimal combination of inputs to employ in the production process (see, e.g., Stigler 1969; Alchian and Allen 1972; Gisser 1981; Hirshleifer 1984; Frank 2000).

The economic concepts and framework needed to

[2] For consistency with the economic literature, we use standard economic notation. Note that subscripts do not represent derivatives of the variable subscripted; for example, MC_q is the marginal cost of q, and p_a is the price of a. A glossary is provided in the appendix.

address question 1 are outlined in ZF. Section 2a begins with an overview of this framework (with minor modifications to the presentation in ZF), then reviews the underlying economic concepts. Section 2b extends this framework, drawing more heavily on economic production theory, to address question 2.

Addressing these questions first requires defining "optimal." For public investment, economists define optimal as maximization of societal welfare. Because resources are scarce (limited), we can never fully satisfy all competing human wants. Societal welfare is maximized by allocating resources across activities in a way that produces the largest possible set of net benefits, that is, through the criterion of *economic efficiency*. When applied to a particular arena, such as weather forecasting, economic efficiency requires that the most valuable forecasting activities are undertaken first, and that the resources expended on each activity are such that no reallocation of resources—either among forecasting activities or between forecast production and activities in other sectors—could increase net societal benefit.

For simplicity, the only use of observations formally considered here is their immediate use in producing real-time weather forecasts. Other uses of observations (e.g., for research to improve weather forecasts in the longer-term, for climate studies, etc.) are also important to consider when allocating public observational resources; this is discussed briefly in section 4. In general, including other uses of observations in the analysis would increase the net societal benefit for a given level of observations, increasing the economically optimal level of public investment in observations.

a. *Optimal level of output*

As discussed above, the criterion of economic efficiency is satisfied by maximizing *net societal benefit* (NB). Net benefit is the difference between *total benefit* (TB) and *total cost* (TC):

$$\mathrm{NB} \equiv \mathrm{TB} - \mathrm{TC}. \quad (1)$$

As depicted in Figs. 1a and 1b, the total benefit, total cost, and net benefit to society of producing a product generally vary with the level of output of the product, q.

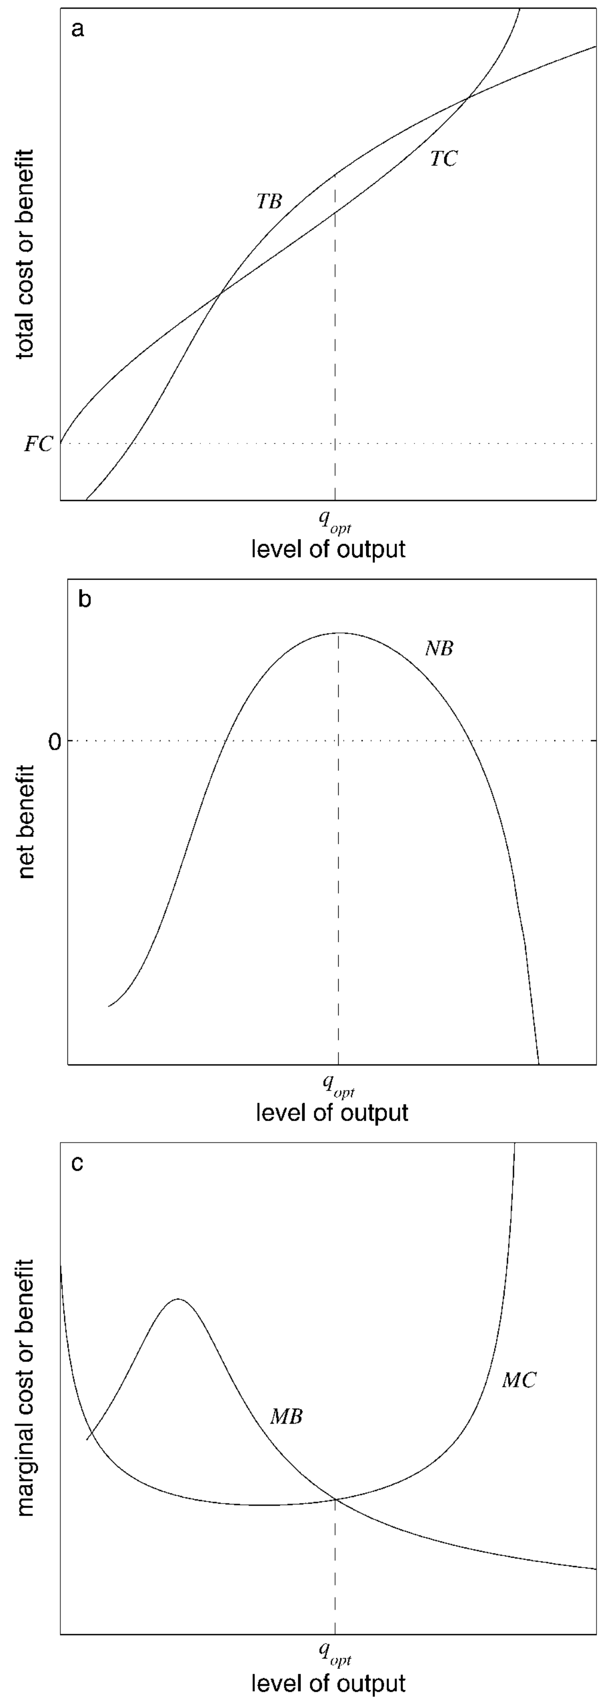

→

FIG. 1. Overview of economic framework, depicting (a) Total benefit (TB) and total cost (TC); (b) net benefit (NB); and (c) marginal benefit (MB) and marginal cost (MC) curves, with shapes typically assumed in economic theory. FC represents the fixed cost of production; TC = FC + VC (variable cost). NB is the difference between the TB and TC curves in (a). MB and MC are the derivatives of the TB and TC curves in (a), respectively. The optimal level of output, q_{opt}, is where NB is maximized or, equivalently, where MB = MC. For details, see section 2a.

The optimal level of output, q_{opt}, is that which generates the maximum net benefit.[3]

Extrema of NB can be found by setting the derivative of Eq. (1) (with respect to q) to zero. The derivative of total benefit is called the *marginal benefit* (MB), and the derivative of total cost is called the *marginal cost* (MC), both depicted in Fig. 1c.[4] Thus, as long as some level of output generates a positive net benefit,[5] the optimal output satisfies the criterion:

$$\mathrm{MB} = \mathrm{MC}. \tag{2}$$

Since more than one level of q can satisfy Eq. (2), we select local maxima of NB by also requiring that the second derivative of Eq. (1) be negative, that is,

$$d\mathrm{MB}/dq < d\mathrm{MC}/dq. \tag{3}$$

This selects the right-hand intersection point of MB and MC in Fig. 1c, the same q_{opt} as in Figs. 1a and 1b.

The curves depicted in Fig. 1 have the general shapes typically assumed in economic theory. Because the optimal output depends on these shapes, let us now take a look at the benefit and cost curves in greater detail.

1) Total and marginal benefit curves

The TB function describes the relationship between the output of a product and the societal benefit derived from that output:

$$\mathrm{TB} \equiv \mathbf{B}(q). \tag{4}$$

Figure 1a depicts a sample total benefit curve, with a shape that displays several typical features. Because some minimal level of output (e.g., minimum level of forecast skill) may be required before society can realize benefits, TB = 0 until q reaches a certain level. As q increases beyond this level, TB may rise at an increasing rate, as synergies allow society to accrue more and more benefits from each increase in output. Eventually, however, the slope of the TB curve will decrease, as further increases in output yield increasingly smaller increments to societal benefit.

The MB is the change in societal benefit derived from producing an additional unit output:

$$\mathrm{MB}_q \equiv d\mathrm{TB}/dq. \tag{5}$$

For a TB curve shaped like that in Fig. 1a, the MB curve will be shaped like that in Fig. 1c: as q increases, MB first increases, then decreases. The decrease in marginal benefit for large q reflects a hypothesized tendency for *diminishing marginal utility*. In other words, our benefit from additional units of any good tends to diminish as we obtain more of that good relative to everything else.

The benefit function is an aggregate over all uses of weather forecasts, including the use of forecast information as a direct consumption good (e.g., by the general public) and as an input to other production processes (e.g., in the private sector or by households). In the case of a service with multiple, diverse uses such as weather forecasts, comprehensively estimating the economic benefits of different levels of production is a challenge, both conceptually and practically. However, as discussed in ZF, basic weather forecast services can often be characterized as *public goods,* in that: (a) many individuals can benefit simultaneously from the forecasts without reducing the benefits derived by anyone else (they are *nonrival*) and (b) it is impractical to exclude nonpaying parties from using the forecasts (they are *nonexcludable*).[6] For public goods, the total benefit can be aggregated by simply summing all individual benefits for any given level of provision of the good. For the more specialized weather forecast services that are *private goods* or *mixed goods*, on the other hand, the rival nature of the use means that estimating total benefit requires accounting for the fact that each individual's consumption of the good diminishes or eliminates its value to others (see, e.g., ZF). Nevertheless, the total benefit derived from basic weather forecasts as a public good is a lower bound on the total benefit from all uses of all weather forecasts.

2) Total and marginal cost curves

The TC function describes the relationship between the level of output of a product and the cost of its production. This can be written as

$$\mathrm{TC} \equiv \mathbf{C}(q) = \mathrm{FC} + \mathrm{VC}(q), \tag{6}$$

where the *fixed cost* (FC) represents the minimum expenditure required to initiate the production process, and the *variable cost* (VC) represents the portion of the cost that varies with the level of production. Fixed costs represent start-up or infrastructure costs, for example, the expenditures required to purchase a supercom-

[3] The framework presented here is similar to standard benefit–cost analysis in that both seek to maximize NB = TB − TC, but is more general. In standard benefit–cost analysis, one or more proposed programs are examined to see whether their NB is greater than zero and/or which has the maximum NB. If each proposed program is represented by a point along the x axis in Figs. 1a and 1c, benefit–cost analysis will select the most economically efficient among them. The public investment in this selected program may, however, be greater or less than the optimal investment derived here. See, for example, Boardman et al. (2001).

[4] In ZF, the MB and MC curves are depicted as straight lines, an approximation when q is close to q_{opt}. Here, we depict MB and MC with nonconstant slopes over a larger range of q because these shapes will be important for determining the optimal levels of inputs (section 2b) and for evaluating public investment when imperfect information is available (section 2c).

[5] If NB < 0 for all q, then $q_{opt} = 0$.

[6] It is these properties of public goods that generally make them most appropriately provided through the public sector, for example, through government intervention, rather than through free competitive markets. See, for example, ZF and Boardman et al. (2001).

puter; build a numerical prediction model; or develop, construct, and launch a satellite. Once the necessary infrastructure is in place, variable costs represent the expenditures required to increase the level of output, for example, to improve a numerical model or obtain and process additional observations from an already-developed observing platform.

Figure 1a depicts a sample total cost curve, with a shape that displays several typical features of production processes. Because a minimum expenditure FC is required to produce any output, TC = FC at $q = 0$ and TC $\geq$ FC for all q. As q increases from 0, initially the slope of the TC curve decreases, as production realizes *economies of scale,* that is, decreasing incremental costs of producing additional units of output. As q increases further, however, TC will tend to rise at an increasing rate, as the principle of *diminishing returns* begins to dominate, that is, as incrementing expenditures by a given amount increases output less and less.

The MC of the product is the change in cost associated with a unit increase in output. Because FC is independent of q,

$$\mathrm{MC}_q \equiv d\mathrm{TC}/dq = d\mathrm{VC}/dq. \qquad (7)$$

Given a total cost curve shaped like that in Fig. 1a, the marginal cost curve will be shaped like that in Fig. 1c: first decreasing with q, then increasing. Note that beyond the first unit of output, marginal cost is independent of fixed cost. This means that once fixed costs have been expended, they drop out of the analysis of the optimal level of output.[7]

b. *Optimal levels (combination) of inputs*

Section 2a derived the optimal level of output of a product, q_{opt}. To determine the optimal level of any *input* to that product, here we examine the relationship between inputs and output that underlies the cost function. That relationship is called a *production function.*

Suppose some inputs (factors of production) have fixed levels, while the levels of other inputs can be varied. The production function can then be written as

$$q \equiv \mathbf{Q}(\bar{f}, a_1, a_2, a_3, \ldots), \qquad (8)$$

where q = the level of output, $\bar{f}$ = the level of the fixed inputs, and $a_1, a_2, a_3, \ldots$ = the levels of the variable inputs. The fixed inputs correspond to the fixed costs described in section 2a(2), and the variable inputs correspond to the variable costs.

1) ONE VARIABLE INPUT

Consider first the case with only one variable input to production, whose level is a. The production function then has only one independent variable, as shown by the *total product* (TP) curve in Fig. 2a. The first derivative of the production function with respect to the variable input is the *marginal product* (MP) of that input,

$$\mathrm{MP}_a = dq/da, \qquad (9)$$

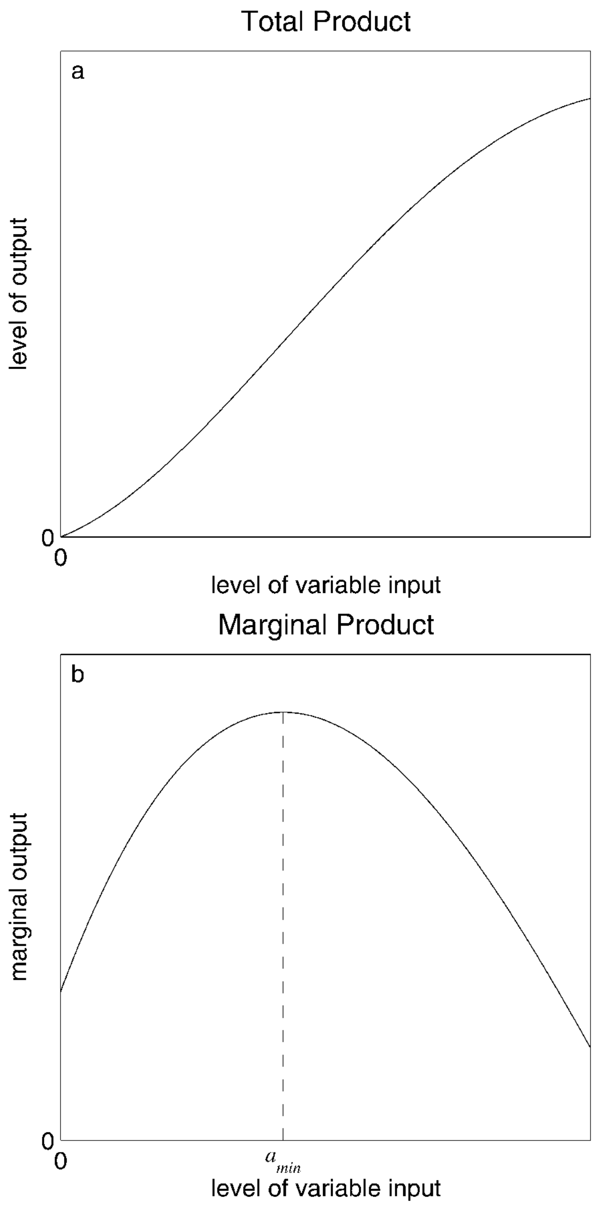

FIG. 2. (a) Total product (TP) curve (plot of production function) for one variable input a, with shape typically assumed in economic theory. (b) Marginal product (MP) curve for one variable input, the derivative of the total product curve in (a). These curves are related to the cost curves in Figs. 1a and 1c as described in section 2b. [Note also that if p_a is constant, the total cost curve in Fig. 1a can be derived from the total product curve in (a) by flipping the x and y axes, multiplying the y axis by p_a, and adding FC.] a_{min} is the minimum efficient level of input to production (assuming reasonable curve shapes), as discussed in section 2c.

[7] In a short-term, one period analysis such as the one presented here. In this sense, the economic analysis is forward-looking.

depicted in Fig. 2b. As discussed in section 2a(2), production processes tend to exhibit economies of scale at some level of production, so that output increases more and more as additional units of input are employed. As the level of input is increased further, however, eventually the principle of diminishing returns dominates, and employing additional units of the input makes progressively smaller contributions to output. Thus, the marginal product depicted in Fig. 2b initially increases, then decreases, corresponding to the initially decreasing, then increasing marginal cost in Fig. 1c. If the price of the input (p_a) is constant, the level of output where marginal product has a maximum in Fig. 2b is the same as that where marginal cost has a minimum in Fig. 1c.

Interpreted in the context of weather forecasting, Figs. 2a and 2b depict how, when observations are available at only a few locations, adding more observations (units of input) would likely improve forecast skill (output) dramatically. Eventually, however, additional observations would tend to become redundant, diminishing their incremental contribution to forecast skill. This redundancy of observations is analogous to the economic principle of diminishing returns. In economics, diminishing returns arise because either 1) the fixed inputs become limiting factors in further production or 2) equal quality units of input have limited availability, so that progressively lower quality units of input must be used to further expand output. The analogous situations in observations for weather prediction are 1) factors that are fixed (in the short term), such as the data assimilation system or numerical model, limit the effects of additional observations on weather forecasts; or 2) observations that provide less and less independent information must be used.

For a production process with one variable input, net benefit is maximized by employing an additional unit of input as long as its price (p_a) is less than the benefit of the additional output it generates ($d\mathbf{B}/dq \times dq/da$ or, equivalently, $\text{MB}_q \times \text{MP}_a$). In other words, as long as some level of output generates a positive net benefit, the optimal level of input, a_{opt}, is that for which the price of an additional unit input just equals the benefit of its incremental increase in output:

$$p_a = \text{MB}_q \times \text{MP}_a \tag{10}$$

or, equivalently,

$$\text{MB}_q = p_a/\text{MP}_a, \text{ at the optimum.} \tag{11}$$

To envision why this criterion yields the optimal level of a, first imagine that a is less than this a_{opt}. The benefit of employing additional units of input to increase output ($\text{MB}_q \times \text{MP}_a$) then exceeds the cost of doing so (p_a); thus, increasing a would increase net benefit. When $a > a_{\text{opt}}$, the cost of employing additional units of input exceeds the benefit of doing so; thus, decreasing a would increase net benefit. Consequently, net benefit is maximized when $a = a_{opt}$. The level of output produced by a_{opt} is q_{opt}, derived from the production function.

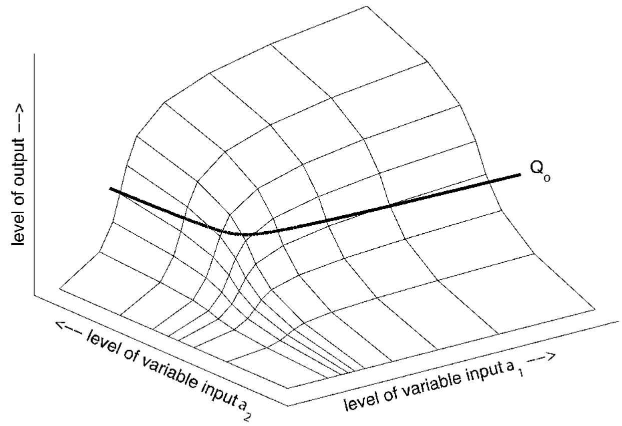

FIG. 3. Sample total product curve for two variable inputs, with levels a_1 and a_2. The isoquant labeled $\mathbf{Q}_o$ represents all combinations of the two inputs that can be used to produce the level of output q_o.

With one variable input, the only variable cost of production is that associated with a. The marginal cost of increasing production is therefore the price of an additional unit input divided by the increment of output realized from employing that additional input. In other words,

$$\text{MC}_q = p_a/\text{MP}_a. \tag{12}$$

Combining Eq. (12) with Eq. (11), we see that the criterion to derive a_{opt} provided by Eq. (11) is equivalent to the criterion to derive q_{opt} provided by Eq. (2) (MB = MC).

2) MULTIPLE VARIABLE INPUTS[8]

Next, consider the more general case with multiple variable inputs, with levels $a_1, a_2, \ldots$. In the context of weather forecasts, the variable inputs can represent types of observations, such as wind, temperature, and water vapor data; observations in different regions, such as continental and oceanic observations; or other uses of resources for weather forecast production, such as model resolution, data assimilation complexity, and ensemble size. Now, we must determine the optimal combination of inputs to use in producing the optimal output.

Figure 3 depicts a sample production function when there are two variable inputs. Note that if either a_1 or a_2 is kept constant in Fig. 3, the two-dimensional curve that remains has a shape similar to that in Fig. 2a. Figure 3 traces out an isoquant, labeled $\mathbf{Q}_o$, that represents all combinations of a_1 and a_2 that produce the level of

[8] Although the example in section 3 has only one variable input, the framework is extended here to multiple variable inputs because that is the primary case of interest for the modern observing and prediction system.

output q_o. Movement along this isoquant represents substituting one variable input for the other in producing q_o. Depending on the prices of the inputs (p_{a1}, p_{a2}, . . .), one point along $\mathbf{Q}_o$ will be the optimal (least cost) way of producing q_o. The total and marginal cost curves are obtained by finding the least-cost combination of inputs that produces each level of output.

The resulting marginal cost curve describes an *efficient expansion path*, along which the levels of the variable inputs increase (or decrease) proportionally such that the following condition applies:

$$\mathrm{MC}_q = p_{a1}/\mathrm{MP}_{a1} = p_{a2}/\mathrm{MP}_{a2} = p_{a3}/\mathrm{MP}_{a3} \ldots \quad (13)$$

(Hirshleifer 1984). Here, the marginal product of each input is a partial derivative, and thus depends on the levels of the other variable inputs. As discussed in section 2a, the optimal level of output is that for which MB = MC [Eq. (2)]. Combining this criterion with Eq. (13), we see that the optimal level of each input is that for which

$$\mathrm{MB}_q = p_{a1}/\mathrm{MP}_{a1} = p_{a2}/\mathrm{MP}_{a2} = p_{a3}/\mathrm{MP}_{a3} \ldots . \quad (14)$$

As one would expect, the optimal combination of inputs depends on their relative prices. For example, if one input becomes more expensive relative to the others, the efficient response is to substitute away from that input, increasing the relative employment of the other inputs (and generally decreasing output in the process).

c. Discussion

Performing the full economic analysis discussed in sections 2a and 2b requires full information about the cost, benefit, and product curves. In a real world problem as complex as weather forecast production, obtaining such information is not feasible. Even with imperfect information, however, the framework can still help narrow the range of the desired level of public investment.

As described in section 2a, typically, above some level of production marginal benefit decreases (due to diminishing marginal utility), and above some (different) level of production marginal cost increases (due to diminishing returns). Figure 1c indicates that the optimal level of production generally occurs on the declining leg of the marginal benefit curve, and on the rising leg of the marginal cost curve. As discussed in section 2b(1), the rising leg of the marginal cost curve corresponds to the declining leg of the marginal product curve. Thus, the maximum net benefit generally occurs on the declining leg of the marginal product curve. This implies that in the case with only one variable input (depicted in Fig. 2b), the optimal level of input will either be zero (if no level of production generates a positive net benefit) or larger than some $a_{\min}$ defined by where the marginal product begins to decline. In other words, assuming reasonable curve shapes, knowing the general shape of the marginal product curve allows one to define an $a_{\min}$ below which it will not be efficient to operate. The case with multiple variable inputs is likely to exhibit a similar minimum efficient scale of production.

Note that as long as some level of production generates a positive net benefit, evaluating the economically optimal level of investment requires only information about marginal cost and benefit—not total cost and benefit. This means that, to apply the framework, one generally does not need to know fixed costs and absolute societal benefits, but rather only how cost and benefit will change with the level of production, and how production will change with the level of the variable inputs.

3. Example: Radiosonde observations and numerical weather forecasts

To demonstrate the framework, this section implements the economic analysis described in section 2 for an idealized example. The example examines synoptic-scale U.S. weather forecast production in a case with only one variable input, the number of vertical observational soundings. Each sounding is similar to a radiosonde profile; no other types of observations are simulated, and all other components of the forecasting system, such as the data assimilation system and forecast model, are kept fixed.[9]

Section 3a uses an idealized set of observing system simulation experiments (OSSEs) to derive a relationship between the number of vertical soundings and forecast skill, that is, a production function for the example. Section 3b combines this production function with a simple estimate of the cost of adding a sounding, deriving a relationship between forecast skill and cost to the United States. Section 3c derives a lower-bound estimate of the marginal benefit to the United States of increasing forecast skill. In section 3d, these estimates are combined to evaluate the level of forecast skill and number of soundings that maximize net economic benefit.

Deriving product, cost, and benefit curves for this example involves several idealizations and approximations. To compensate, we estimate a range of cost and benefit curves, and use these ranges to explore the sensitivity of the results to uncertainties in the analysis. Although the example is presented primarily to illustrate how the framework could be applied, we do briefly discuss potential implications of the results for the real observing system. Missing elements that require further research are discussed in section 4.

a. Production function and marginal product

The production function is derived from a set of idealized OSSEs previously published in Morss et al.

[9] Keeping all inputs but one at a fixed level is analogous to a "short run" economic analysis; see section 4.

(2001, hereafter MES01), with minor modifications. The variable input is the number of vertical soundings, the only type of observation simulated. The level of output is the time- and domain-averaged skill of numerical forecasts, evaluated within the OSSE framework.

1) EXPERIMENTAL SETUP

The experimental method is reviewed briefly here; for additional details, see MES01. All aspects of the experiments are the same as in MES01, except as noted at the end of this section.

The OSSEs are performed using a simple numerical prediction model, a quasigeostrophic (QG) channel model on a midlatitude β plane. The QG model has idealized geometry, including a rigid tropopause lid, a uniform lower boundary, and a circumference that is approximately one-half that of the earth. It is forced by relaxation to a zonal mean state, a Hoskins–West midlatitude jet (Hoskins and West 1979). The model is run at two grid spacings: 250 km (in latitude and longitude) with 5 interior vertical levels ("lower resolution") and 67.5 km with 16 interior vertical levels ("higher resolution"). Further detail on the QG model and its behavior can be found in Rotunno and Bao (1996), MES01, and Snyder et al. (2003).

The simulated observations are vertical wind and temperature profiles, with error characteristics similar to radiosonde soundings, at a set of locations that remains fixed within each experiment. The sounding locations are a randomly selected subset of the lower-resolution model grid points, with a preference toward the middle latitudes of the domain (where the flow is most active). Soundings are spaced a minimum of 350 km apart when 25% or less of the grid points have observations, and 250 km (the gridpoint spacing) otherwise. The observations are incorporated into the QG model using the three-dimensional variational data assimilation (3DVAR) system described in MES01. The 3DVAR system uses the background error statistics developed in Morss (1999) and MES01, optimized for perfect model experiments and 32 simulated sounding locations. Further detail on the observations, observing networks, and data assimilation system can be found in Morss (1999) and MES01.

The OSSEs are run by evolving two parallel trajectories of the QG model: a "true" trajectory that represents the evolution of the atmosphere and a "model" trajectory that simulates a sequence of analyses and forecasts available from a numerical prediction system. Every 12 h, observations are simulated by sampling the true run at the appropriate locations and adding random errors. These observations are then assimilated into the model run using the 3DVAR system. The resulting analysis is used as the initial conditions for numerical forecasts generated with the QG model; the 12-h forecast becomes the background (first guess) field for the assimilation at the next observation time. The true trajectory is evolved using the QG model at higher resolution, and the model trajectory and forecasts are evolved using the QG model at lower resolution. Such experiments are also called "fraternal twin experiments."

As in MES01, a series of OSSEs is performed, varying the number of simulated soundings N. For each N, three experiments are performed, each with different initial states and observing locations. Each experiment is spun up for 30 days, followed by 60 days during which error statistics are accumulated. Analysis and forecast errors are calculated every 12 h with respect to the true state (which is known in this experimental setup).

Forecast skill S is defined as

$$S = \left(1 - \frac{\mathrm{ERR}}{\mathrm{ERR}_0}\right) \times 100, \qquad (15)$$

where ERR is the root-mean-square (rms) domain- and time-averaged error for an experiment at a specific lead time in a specified norm, and ERR_0 is the saturation error for that lead time and norm for the same true trajectory, that is, ERR from the same experiment with no observations. Then S is averaged over the three experiments with N sounding sites. The value of S increases as the analysis/forecast error decreases, that is, with increasing skill, and $S = 0$ when no observations are taken; the maximum possible S is 100, when analyses/forecasts have no error.

Figure 4a depicts how S varies with N for these experiments, for a 1-day forecast lead time and an rms streamfunction error norm. As discussed in MES01, the general shape of this curve is similar for all time-averaged norms tested [including streamfunction, energy (which includes winds and temperature), and enstrophy (squared potential vorticity), vertically averaged and at various model levels]. We chose to use a streamfunction error norm here, even though it is not ideal for relating forecast skill to societal benefit, because this norm describes synoptic-scale errors, and the QG model represents only midlatitude, synoptic-scale dynamics well.

Figure 4a corresponds to the solid line (12-h data assimilation interval results) in MES01's Fig. 2. The results in Fig. 4a are plotted differently than those in MES01; for example, MES01 use the $\mathrm{ERR}/\mathrm{ERR}_0$ term in Eq. (15) as their error norm, reversing the sign convention for increasing skill. The only nonplotting differences between Fig. 4a and MES01 are that here 1) the true trajectory is evolved at higher resolution than the model and forecast trajectories (simulating one type of model error), 2) results are depicted for 1-day forecasts rather than analyses, and 3) the spinup and error accumulation times for each experiment are shorter. The third difference has no noticeable effect; the primary effects of differences 1 and 2 are to decrease the maximum achievable error reduction (decrease S for large N) and decrease the slope of the curve for low–medium N.

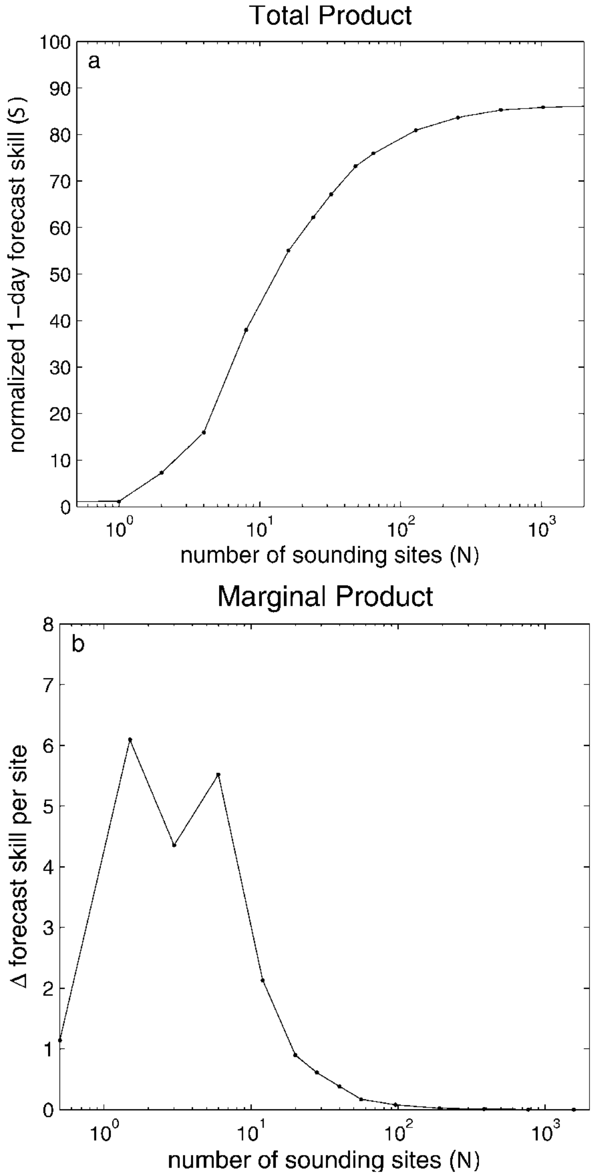

FIG. 4. (a) Production function for the example, generated from the set of idealized OSSEs (fraternal twin experiments) described in section 3a. The x axis is the level of input: the number of vertical soundings N (in the QG model geometry). The y axis is the level of output: the forecast skill S given by Eq. (15), averaged over three experiments with N soundings, with an rms time- and domain-averaged streamfunction error norm. A logarithmic x axis is used to facilitate viewing the results for few observing sites at $N = 0, S = 0$. Note that under the assumptions discussed in section 3b(2), the x axis can be interpreted as one-half the number of Northern Hemisphere, midlatitude vertical sounding sites, and the y axis can be interpreted as average Northern Hemisphere, midlatitude, synoptic-scale forecast skill. (b) Marginal product curve for the idealized example, calculated from the production function in (a) as described in the text. The x axis is the number of vertical soundings N. The y axis is the increment in forecast skill S produced by adding an additional sounding. The dip in marginal product between two and four soundings is due to insufficient sampling (in distributing a small number of observations through the domain) and thus is not significant.

2) PRODUCT CURVES

Figure 4a represents the production function for the example. Figure 4b depicts the corresponding marginal product curve, derived by plotting the slopes of each of the line segments in Fig. 4a at the segment's midpoint. Although the shapes of Figs. 4a and 4b are a product of the statistical and dynamical characteristics of the atmosphere and the prediction system, they are similar to the shapes of the sample production function and marginal product curve depicted in Figs. 2a and 2b. This suggests that the use of observations in weather prediction is analogous to the use of inputs in other types of economic production, in other words, that Figs. 4a and 4b (and corresponding results in MES01) can be interpreted in terms of the economic concepts discussed in section 2. When relatively few observations are employed in numerical prediction, adding observations rapidly improves forecasts, similar to economic production realizing economies of scale. Above a certain number of observations, however, each additional observation improves forecasts less and less, as suggested by the economic principle of diminishing returns.

Numerical prediction models are usually used to generate a suite of forecasts of different lengths from the same initial conditions. In other words, forecasts are generally produced as a "package" of forecasts at different lead times, with related errors. The y axis in Figs. 4a and 4b can therefore be interpreted as the skill of a package of forecasts with different lead times, represented by the skill of 1-day forecasts.[10]

b. *Marginal and variable cost*

To transform the product curves into cost curves, we estimate the price of adding a vertical sounding site (p_a), then use the relationship provided by Eq. (12). Although we assume that p_a is constant, note that the cost of adding a real-world observing site depends on factors such as the observing platform, the number of preexisting observations, and the location. We also employ several approximations to interpret the results from the OSSEs, which were developed using a model with idealized geometry, in terms of real-world weather forecast skill. Costs are estimated in terms of annual costs to the United States in current U.S. dollars.

1) ESTIMATED COST OF ADDING A VERTICAL SOUNDING SITE

The cost of adding a sounding site includes the additional expenditures required to obtain the additional data, communicate and process it, and use it in prediction. Examples of costs to be considered include those

[10] Note that input modifications often affect forecasts at different lead times differently. In general, therefore, the skill of a forecast package is not fully described by a single error norm.

for instrumentation, communication, computer resources, and human labor. Given the other approximations in the example, here we approximate these costs rather than estimate them in detail.

The simulated vertical soundings are similar to radiosonde observations. Purchasing a radiosonde instrument currently costs the U.S. government approximately $70 (J. Facundo 2003, personal communication). With soundings taken every 12 h, the annual instrument cost of adding a radiosonde site is approximately $50,000. For the relatively narrow range of N considered here, most other costs associated with adding an observing site (e.g., for communication, data assimilation) are fixed rather than variable costs, that is, are either one-time expenditures or do not depend on the number of observations. Allowing a small amount ($50,000) for noninstrument variable costs, we estimate that p_a has a lower bound of $100,000 per year. As an upper-bound estimate of p_a, we use $500,000 per year. We chose this large range of p_a so that the example would be relevant for sounding-like observations from a variety of platforms in a variety of regions, not just overland radiosonde soundings. Note, however, that the upper-bound estimate of p_a is probably a significant overestimate, given that in fiscal year 2002 the U.S. Congress appropriated $2.8 billion for *all* (military and civilian) meteorological operations and supporting research (OFCM 2002).

2) COST CURVES

Recall that the QG model used to generate the production function represents only synoptic-scale, midlatitude dynamics, and that the model channel's circumference is approximately half that of the real earth. To scale the production function to the real earth, we assume that 1) in a channel with twice the zonal circumference, obtaining comparable forecast skill would require twice as many sounding sites; and 2) at the lead times of interest, most of the data used in midlatitude forecasts come from midlatitude observations. Under these assumptions, we can interpret $N \times 2$ as the number of Northern Hemisphere, midlatitude soundings, and S as the average skill of Northern Hemisphere, midlatitude, synoptic-scale forecasts.

A large portion of the observational data used in U.S. weather prediction, however, is provided by other countries. To compensate for this, we also assume that, on average, the United States bears approximately one-half of the cost of adding observations in Northern Hemisphere midlatitudes. When transforming the number of sounding sites on the x axis in Fig. 4b to costs on the y axis in the marginal cost curve, this factor of 1/2 cancels the factor of 2 just discussed—and no scaling is required to interpret forecast skill from the OSSEs in terms of Northern Hemisphere, midlatitude, synoptic-scale forecast skill. This scaling is approximate but reasonable, given the other idealizations in the example.

Figure 5a depicts the resulting marginal cost curves

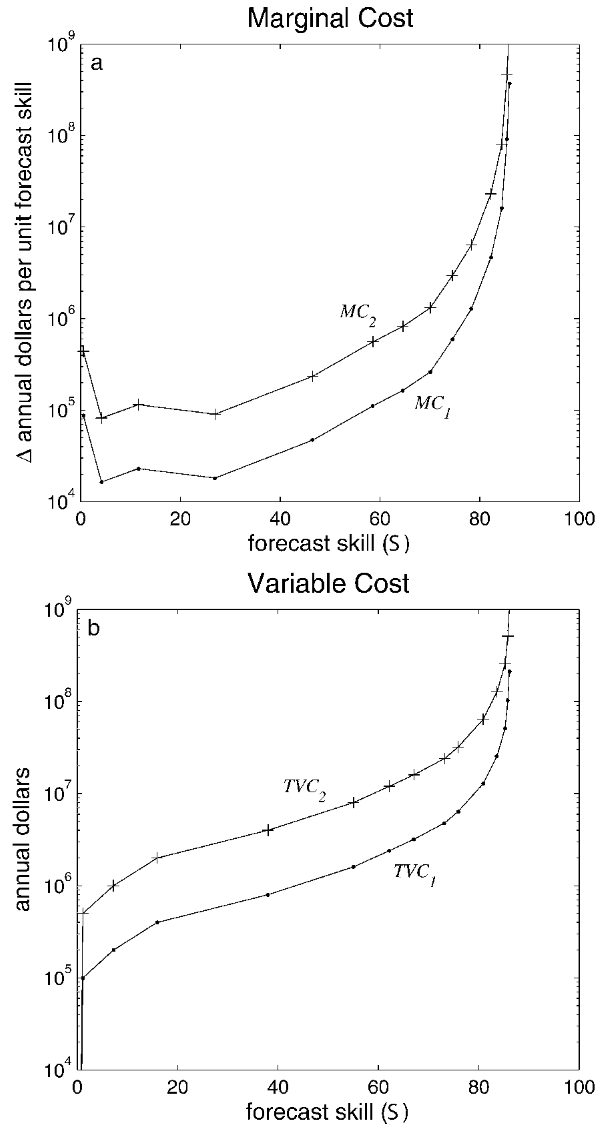

FIG. 5. (a) Marginal cost curve for the idealized example, estimated using the lower- ($100,000 per year) and upper- ($500,000 per year) bound estimates of p_a, labeled MC_1 and MC_2, respectively. The x axis is the forecast skill S (now interpreted as the skill of Northern Hemisphere, midlatitude, synoptic-scale forecasts, as simulated by the OSSEs). The y axis is the approximated additional cost to the United States to obtain an additional unit of forecast skill. A logarithmic y axis is used to facilitate viewing the results for low forecast skill (few observing sites). The curves are derived by combining the marginal product curve from Fig. 4b with the estimates of p_a from section 3b(1). For each level of input in the marginal product curve, Eq. (12) is used to calculate the marginal cost; each value is then plotted at the forecast skill produced by that input, derived from the production function depicted in Fig. 4a. (b) Total variable cost curve for the idealized example, estimated using the lower- and upper-bound estimates of p_a, labeled TVC_1 and TVC_2, respectively. The x axis is the forecast skill S. The y axis is the estimated variable cost (total cost − fixed cost) to the United States of obtaining that level of forecast skill. The variable cost curves are derived by integrating each of the marginal cost curves in (a).

(MC_1, MC_2), derived by using Eq. (12) to combine the marginal product in Fig. 4b with the lower- and upper-bound estimates of p_a from section 3b(1). The corresponding total variable cost curves (TVC_1, TVC_2), each an integral of the corresponding marginal cost curve, are depicted in Fig. 5b. The total cost curves would be the same as the variable cost curves, with a constant (the fixed cost) added to the y axis.

As with the product curves, the marginal and variable cost curves in Figs. 5b and 5a have shapes similar to those described by economic production theory, depicted in Figs. 1a and 1c. When forecast skill is fairly low, increasing forecast skill requires relatively small expenditures (corresponding to relatively few added observations). As forecast skill increases, however, the principle of diminishing returns begins to dominate, and increasing forecast skill further requires progressively larger additional expenditures.

Note that, as discussed in section 2b(1) for a case with constant p_a, the level of output with the minimum marginal cost in Fig. 5a ($S = 5$–30) is the same output (from Fig. 4a) as that produced by the input with the maximum marginal product in Fig. 4b ($N = 2$–8 in the QG model channel).

c. *Estimated marginal benefit*

Limited information is available about the overall societal benefit of weather forecasts, and even less is known about the marginal societal benefit of improving forecasts (ZF; Stratus Consulting 2002, hereafter SC02). As discussed in ZF and section 2a(1), however, basic weather forecast services can often be characterized as public goods. For public good forecasts, the total benefit of any level of forecast skill is simply the sum of all individual benefits. Moreover, the resulting total benefit is a lower bound on the total benefit from all uses of all (basic and specialized) weather forecasts.

We estimate a marginal benefit curve for the example using results from a recent study of the value of daily weather forecasts to the U.S. public as a public good (SC02). The marginal benefit values in SC02 are constants, valid over a limited range of forecast skill. Consequently, the marginal benefit curve derived here is flat (corresponding to a total benefit curve with a constant slope), defined over a limited range of forecast skill. Because SC02 do not include benefits of severe weather forecasts and some sector-specific benefits, their estimates of total benefit can be considered a lower bound, and their estimates of marginal benefit are likely also a lower bound.

SC02 characterize the skill of weather forecasts using four attributes, one of which is accuracy of 1-day forecasts. Using nonmarket valuation methods (structured surveys combined with econometric analysis), SC02 estimate the annual benefit to the average U.S. household of improving forecasts from their current to maximum achievable levels[11] in each of the four attributes. By multiplying the per-household benefit by the number of U.S. households, SC02 then derive the annual incremental benefit to the United States of these maximum achievable forecast improvements, referred to here as ΔB_{MI}.

As with marginal cost, we estimate marginal benefit in terms of annual benefit to the United States, in current U.S. dollars, for each unit improvement in forecast skill. To do so, we divide ΔB_{MI} by $\Delta S_{MI} = S_{max} - S_{current}$, where $S_{current}$ and S_{max} are the current and maximum levels of forecast skill, respectively (in the norm used elsewhere in the example). The resulting estimate of marginal benefit is a constant, valid over the skill range $S_{current}$ to S_{max}.

SC02 estimate that ΔB_{MI} for improving 1-day forecasts is \$1.1 billion, and that ΔB_{MI} for improving the forecast package [see section 3a(2)] consisting of all four attributes is \$1.7 billion. The estimated 95% confidence interval for both values is ±\$300 million. Thus, we estimate ΔB_{MI} = \$1.1–1.7 billion, with a lower bound of \$0.8 billion.

It is not clear how to translate the general forecast skill norm used in SC02 into the rms average streamfunction norm used in the example. We therefore estimate S_{max} and $S_{current}$ using the production function in Fig. 4a, assuming that the skill of numerical weather forecasts is comparable to the skill of weather forecasts received by the public. For S_{max}, we use the S that Fig. 4a asymptotes to for large N (the skill of forecasts from perfect initial conditions): 86.5. For $S_{current}$, we use the S generated by the OSSEs for the approximate number of vertical soundings in the current real-world observing network. The current upper-air observing network used in Northern Hemisphere, midlatitude numerical weather prediction contains approximately 360 twice-daily radiosondes at the observation spacing simulated here.[12] The equivalent N in the QG model channel is 180; from Fig. 4a, S_{max} is then 86.5. With $S_{max} = 86.5$ and $S_{current} = 82$, $\Delta S_{MI} = 4.5$. Note that although we have not accounted for nonradiosonde observations, doing so would increase $S_{current}$, decreasing ΔS_{MI} and increasing MB.

From these estimates of ΔB_{MI} and ΔS_{MI}, Table 1 derives two estimates of MB and a lower bound. Although these values are only approximations, we will

[11] The current and maximum achievable levels of forecast skill were defined in SC02 through consultation with Atmospheric Science Advisors, a meteorology consulting group.

[12] We derived this estimate by taking the set of upper-air soundings decoded operationally at the National Centers for Environmental Prediction from the Global Telecommunications System (obtained online at http://dss.ucar.edu/datasets/ds353.4/data/) for an arbitrary-synoptic time (0000 UTC on 1 October 2003), selecting locations between 20° and 70°N, and counting locations spaced farther apart than approximately 350 km (the observational spacing in the OSSEs).

TABLE 1. Estimates of marginal benefit for the example in section 3; see section 3c.

	ΔB_{MI} (per year)	ΔS_{MI}	MB = $\Delta B_{MI}/\Delta S_{MI}$ (per year per unit S)	Valid range of S
MB_1	\$1.7 billion	4.5	\$380 million	82–86.5
MB_2	\$1.1 billion	4.5	\$240 million	82–86.5
MB_{LB}	\$0.8 billion	4.5	\$180 million	82–86.5

see in the next section that the example's results are relatively insensitive to MB.

d. Analysis

To facilitate depicting the results graphically, we perform the analysis by first deriving the optimal level of forecast skill, using the criterion given by Eq. (2). Figure 6 overlays the marginal cost curves from Fig. 5a with the marginal benefit curves from Table 1; Table 2 summarizes the results. The optimal level of forecast skill (q_{opt}) is where MB and MC intersect: the range S = 84.9–86. According to the production function in Fig. 4a, S = 84.9–86 for approximately 460–1700 soundings in the QG model channel. The optimal number of observations (a_{opt}) in this example is therefore 920–3400 twice-daily soundings in Northern Hemisphere midlatitudes. Note that this range of a_{opt} is the same as that given by the criterion in Eq. (11), that is, is where MB = p_a/MP.

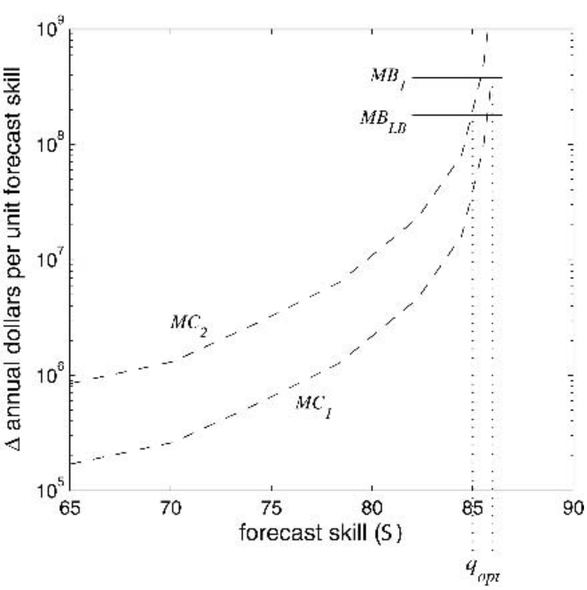

FIG. 6. Overlay of marginal cost (MC) and marginal benefit (MB) curves, for example, depicted for a region of forecast skill (output) near q_{opt} to facilitate viewing the results. The MC_1 and MC_2 curves are from Fig. 5. The MB_1 and MB_{LB} curves are from Table 2; MB_2 (not shown) lies between MB_1 and MB_{LB}. q_{opt} is presented as a range, between the lower-left and upper-right intersection points of MB and MC. The results are summarized in Table 2.

TABLE 2. Results for the example in section 3, for different estimates of MB and MC (p_a). q_{opt} is the optimal level of forecast skill, N is the corresponding optimal number of vertical soundings in the QG model, and a_{opt} is the approximate optimal number of soundings in Northern Hemisphere midlatitudes ($N \times 2$). See section 3d.

	MC_1 (p_a = \$100,000 per year)	MC_2 (p_a = \$500,000 per year)
MB_1 (\$380 million per year per unit S)	q_{opt} = 86.0 N = 1700 a_{opt} = 3400	q_{opt} = 85.4 N = 650 a_{opt} = 1300
MB_2 (\$240 million per year per unit S)	q_{opt} = 85.8 N = 1000 a_{opt} = 2000	q_{opt} = 85.1 N = 490 a_{opt} = 980
MB_{LB} (\$180 million per year per unit S)	q_{opt} = 85.7 N = 920 a_{opt} = 1840	q_{opt} = 84.9 N = 460 a_{opt} = 920

Note that, as discussed in section 2c, a_{opt} and q_{opt} are on the declining leg of the MP curve. In addition, MB and MC intersect in a regime where MC is rising very rapidly, due to rapidly diminishing marginal returns. As Table 2 indicates, in this regime q_{opt} is relatively insensitive to the estimate of MB. However, because MP is very small in this regime, achieving small increments in forecast skill requires large changes in N. Thus, small variations in q_{opt} translate into large variations in a_{opt}.

Although the example contains several idealizations and approximations, limited information is currently available for a more complete analysis. Thus, we briefly discuss what these results might imply for the real observing system. In section 3c, we estimated that the current Northern Hemisphere, midlatitude observing network for numerical weather prediction contains approximately 360 radiosonde locations at the observational spacings simulated here—many fewer than our estimated a_{opt} = 920–3400. Of course, the estimate $a_{current}$ = 360 does not include observations from satellites and other platforms. However, even if the real observing system contains twice as many effective soundings as radiosonde locations (again, at the observational spacings simulated here), $a_{current}$ would still be less than our lower-bound estimate of a_{opt}. Moreover, the lower-bound estimate of a_{opt} results from our upper-bound estimate of p_a (MC) and lower-bound estimate of MB—which are likely over- and underestimates, respectively (see sections 3b and 3b and the introduction to section 2). Thus, the economic analysis suggests that $a_{current} < a_{opt}$, in other words, that increasing the number of observations would increase economic efficiency and net societal benefit.

A more realistic analysis would require improved estimates of marginal cost and marginal benefit, and a production function with multiple variable inputs (representing the different types of observations, e.g., radiosonde, satellite, surface, in the real observing system). These requirements and other aspects of the analysis are discussed next.

4. Discussion

The example analysis in section 3 identifies several areas where improved information is required for a more complete economic evaluation of the real-world observing system. These include the following:

1) improved estimates of marginal product, derived from a production function that includes multiple variable types of observations;
2) improved estimates of p_a, the incremental cost of adding an observation, for various types of observations;
3) improved estimates of the marginal benefit of improving weather forecasts; and
4) forecast verification norms that facilitate connecting the production function with societal benefit, that is, that can be used as both the dependent variable in the production function (independent variable in the cost curves) and the independent variable in the benefit curves.

Obtaining this information requires several types of research and analysis. Item 1 requires meteorological research, to examine how modifying the observing network affects numerical weather forecasts and forecast products. Methodologies for such research include OSSEs and Observing System Experiments (OSEs), data denial experiments, field experiments coupled with comprehensive evaluations of data impact, and studies of the effects of observations on human-generated forecasts. Item 2 requires cost analyses, to estimate the economic cost of obtaining, communicating, processing, and using different types of additional data. Filling these knowledge gaps is within the current capabilities of the meteorological community, and simply requires performing the necessary research and analysis.

Item 3 is a topic for economic and other social science research, and item 4 requires interdisciplinary (joint meteorological–social science) research. Addressing these knowledge gaps therefore requires entraining nonmeteorological expertise. As noted earlier, comprehensively estimating the marginal benefits of weather forecasts is a challenge, both conceptually and practically. Previous studies in this area exist, however, along with several methodologies for performing this research (see, e.g., Katz and Murphy 1997; ZF; SC02; Stratus Consulting 2003).

A major limitation of the example is that it considers only one type of observation, while the real-world observing system contains a mix of observations with different meteorological and cost characteristics. With a multi-input production function and a nonconstant p_a, one could use the framework in section 2b(2) to evaluate the economically optimal combination of observational inputs, and, more generally, to consider trade-offs among different types of observations and observing platforms. A key component of performing such an analysis would be to conduct experiments similar to those presented in section 3a with a more realistic (i.e., primitive equation) model, adding an evaluation of the impact of different types of observations (e.g., wind, temperature, and moisture data with different error structures) at different vertical levels.

Such an analysis could also be extended to examine economically optimal combinations of (and trade-offs among) other variable inputs to weather forecast production, such as the data assimilation system, numerical model components and resolution, and ensemble size. These inputs could be treated in several ways: as fixed (in the short term), as in the example in section 3; as a discrete variable that can (in the long term) shift a production function consisting of other variable inputs; or as a continuously variable input similar to a_1 and a_2 in Fig. 3.

One limitation of the framework presented here is that it assumes that a static benefit curve can be defined, while in reality, societal benefit from weather forecasts is continually changing, as society evolves and different forecasts and technology become available. In addition, the only use of observations considered is real-time production of weather forecasts. Meteorological observations are also used in weather prediction research and development, and in areas such as climate and environmental monitoring, risk analysis, and research. The first leads to additional forecast improvements over the longer term, increasing long-term societal benefit from a given level of investment (by shifting the production function to the right); the second provides additional societal benefits on many time scales. In a general sense, including these uses of observations in an analysis would tend to enhance societal benefit, leading to the same or a greater optimal level of public investment in observations. Evaluating the mix of observations in the multipurpose observing network, however, would require incorporating these uses more systematically, including temporally discounting benefits.

The framework is also limited in that it seeks to optimize net societal benefit, from an economic perspective. In doing so, it may not fully account for other potentially important goals, such as distribution of costs and benefits among different groups (equity) and political feasibility. Implementing the framework also requires deciding whose outcomes will be considered in net societal benefit (who has "standing") and translating outcomes into monetary values, both of which can be controversial and require careful consideration (see, e.g., Boardman et al. 2001). For these reasons, economic analyses such as those presented here are not guides to observing network policy decisions, but rather tools that, if implemented well, can provide useful input to such decisions.

More generally, the framework and economic concepts introduced can help meteorologists and policy makers consider trade-offs among types of observations (and other inputs to weather forecast production), based on their different costs and their different contributions to forecast skill and societal benefit. Such information can contribute to the design of more cost-effective observing networks and forecasting systems, benefiting weather forecasts, the meteorological community, and society.

5. Summary

Designing cost-effective publicly funded meteorological observing networks has been discussed in the meteorological community for decades (e.g., WMO 1967; NRC 1969, 1980) and continues to be of interest to meteorologists and policy makers (e.g., Emanuel et al. 1995, 1997; NRC 1998; Lautenbacher 2003; Shapiro and Thorpe 2004). Yet, despite this long-term interest, the meteorological community lacks a practical, systematic framework for analyzing public investment in observations for weather forecasting. This paper presents the basic elements of such a framework, using an economic approach.

The framework analyzes observations as a variable input in the real-time production of weather forecasts. Using standard economic concepts similar to those presented in Zillman and Freebairn (2001), Freebairn and Zillman (2002), and Zillman (2002), the framework first compares information about economic costs and benefits to derive the optimal level of weather forecast production. The framework is then extended, drawing on economic production theory, to analyze the optimal levels and mix of observational inputs to use in producing the optimal level of forecasts. Although the framework is presented in the context of observations for weather forecast production, it can also be used to analyze mixes of other inputs to the weather forecasting process, as well as the production of other types of forecasts (e.g., climate forecasts) and other uses of earth observations.

The economic framework is demonstrated by constructing and analyzing an example for vertical (radiosonde-like) soundings and numerical weather forecasts in the United States, based on results from previous research. The shapes of the functions derived in the example are similar to those suggested by economic theory, suggesting that the results (and more generally, the use of observations in weather forecast production) can be interpreted using the economic concepts in the framework. Because the example contains several idealizations and approximations, the detailed results cannot be applied directly to the real observing system. However, the results do suggest that increasing the current level of observations is likely to be economically efficient, in other words, to enhance net societal benefit.

The example also identifies several knowledge gaps that currently prevent implementation of a more detailed, more realistic economic analysis. These include 1) weather forecast production functions (which describe the relationship between input levels and forecast output), 2) incremental economic costs of adding different types of observations, 3) marginal societal benefits of improving weather forecasts, and 4) forecast verification norms that are both meteorologically and societally relevant. Addressing these knowledge gaps is a prerequisite to understanding trade-offs among different types of observations—and, more generally, among various inputs to the weather prediction process. Understanding trade-offs among investments in different components of the observing system is a major topic of current interest, in the United States and internationally (e.g., NRC 1998; Shapiro and Thorpe 2004). Thus, even outside the context of a formal economic analysis, performing the research necessary to address these gaps can help meteorologists and policy makers design more cost-effective observing networks and forecasting systems, improve weather forecasts, and benefit society.

Acknowledgments. The authors thank Chris Snyder, Eugene Wahl, Jeff Lazo, and three anonymous reviewers for their comments on the manuscript.

APPENDIX

Glossary of Variables Introduced in Section 2

Variable	Equivalence	Definition
q		Level of product (output)
q_{opt}		Optimal level of product
TB	$= \mathbf{B}(q)$	Total benefit, benefit function
MB, MB_q	$= d\mathrm{TB}/dq = d\mathbf{B}/dq$	Marginal benefit (derivative of total benefit with respect to q)
TC	$= \mathbf{C}(q) = \mathrm{FC} + \mathrm{VC}$	Total cost, cost function
MC, MC_q	$= d\mathrm{TC}/dq = d\mathbf{C}/dq$	Marginal cost (derivative of total cost with respect to q)
NB	$= \mathrm{TB} - \mathrm{TC}$	Net benefit
FC		Fixed cost (constant)

Variable	Equivalence	Definition
VC, TVC		Variable cost (a function of q)
a		Level of variable input (for one variable input)
a_{opt}		Optimal level of variable input
$a_1, a_2, a_3, \ldots$		Levels of variable inputs (for multiple variable inputs)
q	$= \mathbf{Q}(\bar{f}, a_1, a_2, a_3, \ldots)$	Production function with multiple variable inputs ($\bar{f}$ represents the fixed inputs)
MP, MP_a	$= dq/da$	Marginal product for one variable input (derivative of product with respect to a)
$MP_{a1}, MP_{a2}, MP_{a3}, \ldots$	$= \partial q/\partial a_1, \partial q/\partial a_2, \partial q/\partial a_3, \ldots$	Marginal products for multiple variable inputs
p_a		Price of variable input a
$p_{a1}, p_{a2}, p_{a3}, \ldots$		Prices of variable inputs a_1, a_2, a_3, . . .

REFERENCES

Alchian, A. A., and W. R. Allen, 1972: *University Economics: Elements of Inquiry.* 3d ed. Wadsworth, 857 pp.

Boardman, A. E., D. H. Greenberg, A. R. Vining, and D. L. Weimer, 2001: *Cost-Benefit Analysis: Concepts and Practice.* 2d ed. Prentice Hall, 526 pp.

Doswell, C. A., III, and H. E. Brooks, 1998: Budget cutting and the value of weather services. *Wea. Forecasting,* **13,** 206–212.

Emanuel, K., and Coauthors, 1995: Report of the first prospectus development team of the U.S. Weather Research Program to NOAA and NSF. *Bull. Amer. Meteor. Soc.,* **76,** 1194–1208.

——, and Coauthors, 1997: Observations in aid of weather prediction: Report of prospectus development team seven. *Bull. Amer. Meteor. Soc.,* **78,** 2859–2868.

Frank, R. H., 2000: *Microeconomics and Behavior.* 4th ed. McGraw-Hill, 672 pp.

Freebairn, J. W., and J. W. Zillman, 2002: Economic benefits of meteorological services. *Meteor. Appl.,* **9,** 33–44.

Gisser, M., 1981: *Intermediate Price Theory: Analysis, Issues and Applications.* McGraw-Hill, 653 pp.

Hirshleifer, J., 1984: *Price Theory and Applications.* 3d ed. Prentice Hall, 574 pp.

Hoskins, B. J., and N. V. West, 1979: Baroclinic waves and frontogenesis. Part II: Uniform potential vorticity jet flows—Cold and warm fronts. *J. Atmos. Sci.,* **36,** 1663–1680.

Katz, R. W., and A. H. Murphy, Eds., 1997: *Economic Value of Weather and Climate Forecasts.* Cambridge University Press, 240 pp.

Lautenbacher, C. C., Jr., cited 2003: Three imperatives for an integrated earth observation system. Address to the 2003 WMO Congress. [Available online at http://earthobservationsummit.gov/speeches_wmo.html and http://www.noaa.gov/lautenbacher/wmo2003.htm.]

Mikulski, B. A., 1994: Science in the national interest. *Science,* **264,** 221–222.

Morss, R. E., 1999: Adaptive observations: Idealized sampling strategies for improving numerical weather prediction. Ph.D. dissertation, Massachusetts Institute of Technology, 225 pp. [Available from UMI Dissertation Services, P.O. Box 1346, Ann Arbor, MI 48106-1346.]

——, K. A. Emanuel, and C. Snyder, 2001: Idealized adaptive observation strategies for improving numerical weather prediction. *J. Atmos. Sci.,* **58,** 210–232.

NRC, 1969: *Plan for U.S. Participation in the Global Atmospheric Research Program.* National Academy Press, 79 pp.

——, 1980: *Technological and Scientific Opportunities for Improved Weather and Hydrological Services in the Coming Decade.* National Academy Press, 87 pp.

——, 1998: *The Atmospheric Sciences: Entering the Twenty-first Century.* National Academy Press, 384 pp.

NOAA, 2002: Summary of the president's budget: Fiscal year 2003. Department of Commerce, National Oceanic and Atmospheric Administration, 232 pp. [Available online at http://www.publicaffairs.noaa.gov/budget2003/noaabluebook.pdf.]

OFCM, 2002: The federal plan for meteorological services and supporting research: Fiscal year 2003. Department of Commerce FCM-P1-2002, Office of the Federal Coordinator for Meteorology, 229 pp.

Rotunno, R., and J.-W. Bao, 1996: A case study of cyclogenesis using a model hierarchy. *Mon. Wea. Rev.,* **124,** 1051–1066.

Shapiro, M. A., and A. J. Thorpe, cited 2004: THORPEX International Science Plan. [Available online at http://www.wmo.int/thorpex/publications.html.]

Snyder, C., T. M. Hamill, and S. B. Trier, 2003: Linear evolution of error covariances in a quasigeostrophic model. *Mon. Wea. Rev.,* **131,** 189–205.

Stigler, G. J., 1969: *The Theory of Price.* 3d ed. Macmillan, 355 pp.

Stratus Consulting, 2002: Economic value of current and improved weather forecasts in the U.S. household sector. Stratus Consulting SC10050, 211 pp

——, 2003: Benefit analysis for National Oceanic and Atmospheric Administration high-performance computing system for research applications. Stratus Consulting SC10367, 84 pp.

WMO, 1967: Design of optimum networks for aerological observing stations. World Weather Watch Planning Rep. 21, World Meteorological Organization, 58 pp.

Zillman, J. W., 2002: Benefits and costs of weather and climate services. Preprints, *Third Symp. on Environmental Applications,* Orlando, FL, Amer. Meteor. Soc., CD-ROM, J4.0A.

——, and J. W. Freebairn, 2001: Economic framework for the provision of meteorological services. *WMO Bull.,* **50,** 206–215.