Subjective Probability and Lower and Upper Prevision: a New Understanding *

G. SHAFER
Rutgers Business School – Newark & New Brunswick, NJ, USA

P. R. GILLETT
Rutgers Business School – Newark & New Brunswick, NJ, USA

R. B. SCHERL
Department of Computer Science, Monmouth University, NJ, USA

Abstract

This article introduces a new way of understanding subjective probability and its generalization to lower and upper prevision. Instead of asking whether a person is willing to pay given prices for given risky payoffs, we ask whether the person believes he can make a lot of money at those prices. If not—if the person is convinced that no strategy for exploiting the prices can make him very rich in the long run—then the prices measure his subjective uncertainty about the events involved.

This new understanding justifies Peter Walley's updating principle, which applies when new information is anticipated exactly. It also justifies a weaker principle that is more useful for planning because it applies even when new information is not anticipated exactly. This weaker principle can serve as a basis for flexible probabilistic planning in event trees.

Keywords

Subjective probability, upper and lower prevision, updating, event trees.

1 Introduction

In the established understanding of subjective probability, set out by Bruno de Finetti [2] and his followers, a person's beliefs are revealed by the bets he is willing to make. The odds at which he is willing to bet define his probabilities.

We develop a somewhat different understanding of subjective probability, using Shafer and Vovk's game-theoretic framework [8]. In this framework, probability is understood in terms of two players: one who offers bets, and one to whom

*Research for this article was supported by NSF grant SES-9819116.

the bets are offered. We call these two players *House* and *Gambler*, respectively. The established understanding seems to be concerned with House's uncertainty, since he is the one stating odds and offering to bet. But following Shafer and Vovk, we take Gambler's point of view. Gambler is trying to beat the odds, and Shafer and Vovk's work suggests that what makes odds expressions of a person's uncertainty is his conviction that he cannot beat them.

We briefly introduce our new understanding of subjective probability in Section §2 and immediately generalize it to lower and upper prevision in Section §3.

2 Subjective Probability

Suppose House states odds $p:(1-p)$ and offers Gambler the opportunity to bet any amount he chooses for or against E at these odds. This means that House offers Gambler the payoff

$$\begin{cases} \alpha(1-p) & \text{if } E \text{ happens} \\ -\alpha p & \text{if } E \text{ fails} \end{cases} \tag{1}$$

for any real number α, which Gambler must choose immediately, before any other information becomes available. The absolute value of α is the total stakes for the bet, and the sign of α indicates which side Gambler is taking:

- If α is positive, then Gambler is betting on E happening. Gambler puts up αp, which he loses to House if E fails, while House puts up $\alpha(1-p)$, which he loses to Gambler if E happens. The total stakes are $\alpha p+\alpha(1-p)$, or α.
- If α is negative, then Gambler is betting against E happening. Gambler puts up $-\alpha(1-p)$, which he loses to House if E happens, while House puts up $-\alpha p$, which he loses to Gambler if E fails. The total stakes are $-\alpha(1-p)-\alpha p$, or $-\alpha$.

No principle of logic requires House to state odds at which Gambler can take either side. But mathematical probability has earned our attention by its practical successes over several centuries, and if we follow de Finetti in rejecting as defective all past attempts to provide objective interpretations of probability, then we seem to be left with (1) as the only viable way of interpreting this successful mathematical theory.

De Finetti developed this interpretation from the viewpoint of the player we are calling House. The principle that House should avoid sure loss to Gambler was fundamental to this development. If we agree that House should offer Gambler (1) for some p, then the principle that House should avoid sure loss leads immediately to the conclusion that p should be unique. If House offers (1) for both p_1 and p_2, where $p_1<p_2$, then Gambler can accept the p_1-offer with $\alpha=1$ and the p_2-offer with $\alpha=-1$, and this produces a sure gain of p_2-p_1 for Gambler, no matter whether E happens or fails.

2.1 Protocols

From a thoroughly game-theoretic point of view, the game between House and Gambler also involves a third player, who decides the outcomes on which they are betting. Calling this third player *Reality*, we can lay out an explicit protocol for the game in the style of Shafer and Vovk [8].

PROBABILITY FORECASTING
House announces $p \in [0,1]$.
Gambler announces $\alpha \in \mathbb{R}$.
Reality announces $x \in \{0,1\}$.
$\mathcal{K}_1 := \mathcal{K}_0 + \alpha(x - p)$.

This is a perfect-information protocol; the players move in the order indicated (not simultaneously), and each player sees the other players' moves as they are made. We have written $\mathcal{K}_0$ for Gambler's initial capital and $\mathcal{K}_1$ for his final capital. Reality's announcement indicates the happening or failure of E: $x = 1$ means E happens, and $x = 0$ means E fails. Thus $\alpha(x - p)$ is the same as (1). This is Gambler's net gain, which we can think of as the result of his paying αp for αx; Gambler buys α units of x for p per unit.

We may, for example, present de Finetti's argument for Additivity in this format. Consider the following protocol, where House gives probabilities for the three events E, F, and $E \cup F$, where E and F are disjoint:

MULTIPLE PROBABILITY FORECASTING
House announces $p_E, p_F, p_{E\cup F} \in [0,1]$.
Gambler announces $\alpha_E, \alpha_F, \alpha_{E\cup F} \in \mathbb{R}$.
Reality announces $x_E, x_F, x_{E\cup F} \in \{0,1\}$.
$\mathcal{K}_1 := \mathcal{K}_0 + \alpha_E(x_E - p_E) + \alpha_F(x_F - p_F) + \alpha_{E\cup F}(x_{E\cup F} - p_{E\cup F})$.
Constraint on Reality: Reality must make $x_{E\cup F} = x_E + x_F$ (this expresses the assumptions that E and F are disjoint and that $E \cup F$ is their disjunction).

The constraint on Reality is part of the rules of the game. Like the other rules, it is known to the players at the outset.

To see that House must make $p_{E\cup F} = p_E + p_F$ in order to avoid sure loss in this protocol, set

$$\delta := \begin{cases} 1 & \text{if } p_{E\cup F} > p_E + p_F \\ 0 & \text{if } p_{E\cup F} = p_E + p_F \\ -1 & \text{if } p_{E\cup F} < p_E + p_F \end{cases}$$

and consider the strategy for Gambler in which α_E and α_F are equal to δ and $\alpha_{E\cup F}$ is equal to $-\delta$. Gambler's net gain with this strategy is

$$\delta(x_E - p_E) + \delta(x_F - p_F) - \delta(x_{E\cup F} - p_{E\cup F}) = \delta(p_{E\cup F} - (p_E + p_F)),$$

which is positive unless $p_E + p_F = p_{E \cup F}$.

This argument readily generalizes to an argument for the rule that relates the expected value (or prevision) of a payoff to the probability of the events that determine the payoff.

2.2 Cournot's Principle

The rules of probability can be derived from House's motivation to avoid sure loss. But a clear understanding of how subjective probabilities should be updated over time requires that we shift to Gambler's viewpoint and invoke Cournot's principle. When we assert that certain numbers are valid as objective probabilities, we are asserting that they do not offer anyone any opportunity to get very rich. When we advance them as our subjective probabilities, we are saying something only a little different: we are asserting that they do not offer us, with the knowledge we have, any opportunity to get very rich. When we say this, we put ourselves in the role of Gambler, not in the role of House. The point is not how we got the numbers: the point is what we think we can do with them.

A probability for a single event, if it is not equal to 0 or 1, can hardly be refuted. Even if Gambler chooses the winning side, with stakes high enough to make a lot of money, we will hesitate to conclude that the probability was wrong. Gambler may simply have been lucky. On the other hand, if House announces probabilities for a sequence of events, and Gambler consistently manages to make money, then the validity of the probabilities will be cast in doubt.

Shafer and Vovk [8] have shown that we can make this notion of testing precise within the following protocol, where House announces probabilities $p_1, p_2, \ldots$ for a series of events $E_1, E_2, \ldots$ with indicator variables $x_1, x_2, \ldots$:

SEQUENTIAL PROBABILITY FORECASTING
$\mathcal{K}_0 := 1$.
For $n = 1, 2, \ldots$:
House announces $p_n \in [0,1]$.
Gambler announces $\alpha_n \in \mathbb{R}$.
Reality announces $x_n \in \{0,1\}$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + \alpha_n (x_n - p_n)$.

In this protocol, Gambler can test House's probabilities by trying to get infinitely rich ($\lim_{n\to\infty} \mathcal{K}_n = \infty$) without ever risking bankruptcy (without giving Reality an opportunity to make $\mathcal{K}_n$ negative for any n). If Gambler succeeds in doing this, he has refuted an infinite subset of the set of given probabilities.

Shafer and Vovk use the name *Cournot's principle* for the hypothesis that Reality will not allow Gambler to become infinitely rich without risking bankruptcy. This principle says that no matter what bankruptcy-free strategy for Gambler we specify (in addition to House's and Reality's previous moves, such a strategy may also use other information available to Gambler), we can be confident that Reality will move in such a way that the strategy will not make Gambler infinitely rich. This is an empirical hypothesis—a hypothesis about how Reality will behave, not a rule of the game.

If given probabilities satisfy Cournot's principle for any potential gambler, no matter how much information that gambler has, then we might call them *objective* or *causal* probabilities [5, 6]. On the other hand, if they satisfy Cournot's principle only for gamblers with a certain level of information, then we might call them *subjective* probabilities for that level of information. An individual who believes that the probabilities provided to him by some source or method do not permit any bankruptcy-free strategy to make him very rich might reasonably call them his personal subjective probabilities.

Under this interpretation, a person with subjective probabilities is not merely saying that he does not know how to get very rich betting at these probabilities. He is saying that he is convinced that there is no bankruptcy-free strategy that will make him very rich.

3 Subjective Lower and Upper Prevision

In recent decades, there has been great interest in supplementing subjective probability with more flexible representations of uncertainty. Some of the representations studied emphasize evidence rather than gambling [4, 9, 12]; others use a concept of partial possibility [3]. But many scholars prefer to generalize the story about betting that underlies subjective probability. The first step of such a generalization is obvious. Instead of requiring a person to set odds at which he will take either side of a bet, allow him to set separate odds for the two sides. This leads to lower and upper probabilities and lower and upper previsions rather than additive probabilities and expected values. See the early work of C. A. B. Smith [10, 11] and Peter Williams [17, 18, 19], the influential work of Peter Walley [13, 14, 15, 16], and the recent work of the imprecise probabilities project [1].

In this section, we look at lower and upper previsions from the point of view developed in the preceding section. This leads to a better understanding of how these measures of subjective uncertainty should change with new information, both when the new information is *exact* (i.e., when it is the *only* additional information) and when it is not.

3.1 Pricing Events and Payoffs

Whereas probabilities for events determine expected values for payoffs that depend on those events (see §2.1), lower and upper probabilities are not so informative. The rates at which a person is willing to bet for or against given events do not necessarily determine the prices at which he is willing to buy or sell payoffs depending on those events. We need more than a theory of lower and upper probabilities for events: we need a theory of lower and upper previsions for payoffs.

3.1.1 Lower and Upper Probabilities

Suppose House expresses his uncertainty about E by specifying two numbers, p_1 and p_2. He offers to pay Gambler

$$-\alpha_1(x-p_1) = \begin{cases} -\alpha_1(1-p_1) & \text{if } E \text{ happens} \\ \alpha_1 p_1 & \text{if } E \text{ fails} \end{cases} \tag{2}$$

for any $\alpha_1 \geq 0$, and he also offers to pay Gambler

$$\alpha_2(x-p_2) = \begin{cases} \alpha_2(1-p_2) & \text{if } E \text{ happens} \\ -\alpha_2 p_2 & \text{if } E \text{ fails} \end{cases} \tag{3}$$

for any $\alpha_2 \geq 0$. In (2), Gambler sells α_1 units of x for p_1 per unit, while in (3), he buys α_2 units of x for p_2 per unit. Here is the protocol for this:

FORECASTING WITH LOWER AND UPPER PROBABILITIES

House announces $p_1, p_2 \in [0,1]$.
Gambler announces $\alpha_1, \alpha_2 \in [0,\infty)$.
Reality announces $x \in \{0,1\}$.
$\mathcal{K}_1 := \mathcal{K}_0 - \alpha_1(x-p_1) + \alpha_2(x-p_2)$.

To avoid sure loss, House must make $p_1 \leq p_2$. If $p_1 > p_2$, then Gambler can make money for sure by making α_1 and α_2 strictly positive and equal.

House would presumably be willing to increase his own payoffs by decreasing p_1 in (2) and increasing p_2 in (3). The natural remaining question is how high House will make p_1 and how low he will make p_2. We may call p_1 and p_2 *House's lower and upper probabilities*, respectively, if House will not offer (2) for any value higher than p_1 and will not offer (3) for any value lower than p_2.

When we model our beliefs by putting ourselves in the role of House, we have some flexibility in the meaning we give our refusal to offer higher values of p_1 or lower values of p_2. Perhaps we are certain that we do not want to make additional offers, perhaps we are hesitating, or perhaps we are providing merely an incomplete model of our beliefs (Walley [14], pp. 61–63).

When we instead model our beliefs by putting ourselves in the role of Gambler, the question is what values of p_1 and p_2 we believe will satisfy Cournot's principle. In the context of a sequence of forecasts, we might call p_1 and p_2 *Gambler's lower and upper probabilities* when (1) Gambler believes that no strategy for buying and selling will make him very rich in the long run when he can sell x for p_1 or buy it for p_2 but (2) Gambler is not confident about this in the case where he is allowed to sell x for more than p_1 or buy it for less than p_2.

Clause (2) can be made precise in more than one way. Gambler might be unsure about whether he can get very rich with better values of p_1 or p_2, or he might believe that a strategy available to him would succeed with such values.

3.1.2 Lower and Upper Previsions

The following protocol allows us to price a payoff x that depends on the outcome of more than one event:

FORECASTING WITH LOWER AND UPPER PREVISIONS
House announces $p_1, p_2 \in \mathbb{R}$.
Gambler announces $\alpha_1, \alpha_2 \in [0, \infty)$.
Reality announces $x \in \mathbb{R}$.
$\mathcal{K}_1 := \mathcal{K}_0 - \alpha_1(x - p_1) + \alpha_2(x - p_2)$.

Again, Gambler is allowed to sell x for p_1 and buy it for p_2. If p_1 is the highest price at which Gambler can sell x (either the highest price House will offer or the highest price at which Gambler believes Cournot's principle, depending on our viewpoint), we may call it the *lower prevision* of x. Similarly, if p_2 is the lowest price at which Gambler can buy x, we may call it the *upper prevision* of x.

House may have more to say about x than the lower and upper previsions p_1 and p_2, and even the statement that these are lower and upper previsions is not exactly a statement about the protocol itself. We now turn to a more abstract approach, better suited to general discussion.

3.2 Forecasting in General

Consider a set $\mathbf{R}$, and consider a set $\mathbf{H}$ of real-valued functions on $\mathbf{R}$. We call $\mathbf{H}$ a *belief cone* on $\mathbf{R}$ if it satisfies these two conditions:

1. If $\mathbf{g}$ is a real-valued function on $\mathbf{R}$ and $\mathbf{g}(\mathbf{r}) \leq 0$ for all $\mathbf{r} \in \mathbf{R}$, then $\mathbf{g}$ is in $\mathbf{H}$.
2. If $\mathbf{g}_1$ and $\mathbf{g}_2$ are in $\mathbf{H}$ and a_1 and a_2 are nonnegative numbers, then $a_1\mathbf{g}_1 + a_2\mathbf{g}_2$ is in $\mathbf{H}$.

We write $\mathcal{C}_\mathbf{R}$ for the set of all belief cones on $\mathbf{R}$.

Intuitively, a belief cone is a set of payoffs that House might offer Gambler. Thus if $(\alpha - \mathbf{g}) \in \mathbf{H}$, House is willing to buy $\mathbf{g}$ for α; and if $(\mathbf{g} - \alpha) \in \mathbf{H}$, House is willing to sell $\mathbf{g}$ for α. Condition 1 says that House will at least offer any contract that does not require him to risk a loss. Condition 2 says House will allow Gambler to combine any two of his offers, in any amounts. These conditions are, of course, closely related to Walley's concept of *desirability*.

The following abstract protocol is adapted from p. 90 of [8].

FORECASTING

Parameters: $\mathbf{R}$ and $\mathcal{C} \subseteq \mathcal{C}_{\mathbf{R}}$
Protocol:
House announces $\mathbf{H} \in \mathcal{C}$.
Gambler announces $\mathbf{g} \in \mathbf{H}$.
Reality announces $\mathbf{r} \in \mathbf{R}$.
$\mathcal{K}_1 := \mathcal{K}_0 + \mathbf{g}(\mathbf{r})$.

We call any protocol obtained by a specific choice of $\mathbf{R}$ and $\mathcal{C}$ a *forecasting protocol*. We call $\mathbf{R}$ the *sample space*.

We call a real-valued function on the sample space $\mathbf{R}$ a *variable*. House's move $\mathbf{H}$, itself a set of variables, determines lower and upper previsions for all variables. The *lower prevision* for a bounded variable x is

$$\underline{\mathbb{E}}_{\mathbf{H}}\, x := \sup\{\alpha \mid (\alpha - x) \in \mathbf{H}\}, \tag{4}$$

and the *upper prevision* is

$$\overline{\mathbb{E}}_{\mathbf{H}}\, x := \inf\{\alpha \mid (x - \alpha) \in \mathbf{H}\}. \tag{5}$$

These definitions are similar to those given by Walley ([14], pp. 64–65), with a difference in sign because Walley considers a collection $\mathcal{D}$ of payoffs that House is willing to accept for himself rather than a collection $\mathbf{H}$ that House offers.

The condition $(\alpha - x) \in \mathbf{H}$ in (4) means that Gambler can sell x for α. So roughly speaking, the lower prevision $\underline{\mathbb{E}}_{\mathbf{H}}\, x$ is the highest price at which Gambler can sell x. We say "roughly speaking" because (4) tells us only that Gambler can obtain $\alpha - x$ for α arbitrarily close to $\underline{\mathbb{E}}_{\mathbf{H}}\, x$, not that he can obtain $(\underline{\mathbb{E}}_{\mathbf{H}}\, x) - x$. Similarly, the upper prevision $\overline{\mathbb{E}}_{\mathbf{H}}\, x$ is roughly the lowest price at which Gambler can buy x.

Once we know lower previsions for all variables, we also know upper previsions for all variables, and vice versa, because

$$\overline{\mathbb{E}}_{\mathbf{H}}\, x = -\underline{\mathbb{E}}_{\mathbf{H}}(-x)$$

for every variable x. For additional general properties of lower and upper previsions, see Chapter 2 of Walley [14] and Chapters 1 and 8 of [8].

3.2.1 Regular Protocols

Given $\mathbf{H} \in \mathcal{C}_{\mathbf{R}}$, set

$$\mathbf{H}^* := \{x : \mathbf{R} \mapsto \mathbb{R} \mid \overline{\mathbb{E}}_{\mathbf{H}}\, x \le 0\}.$$

The following facts can be verified straightforwardly:

- $\mathbf{H}^*$ is also a belief cone ($\mathbf{H}^* \in \mathcal{C}_{\mathbf{R}}$),
- $\mathbf{H} \subseteq \mathbf{H}^*$,

- $\overline{\mathbb{E}}_{\mathbf{H}}\,x = \overline{\mathbb{E}}_{\mathbf{H}^*}\,x$ and $\underline{\mathbb{E}}_{\mathbf{H}}\,x = \underline{\mathbb{E}}_{\mathbf{H}^*}\,x$ for every variable x, and
- $(\mathbf{H}^*)^* = \mathbf{H}^*$.

Intuitively, if House offers Gambler all the payoffs in $\mathbf{H}$, then he might as well also offer the other payoffs in $\mathbf{H}^*$, because for every payoff in $\mathbf{H}^*$, there is one in $\mathbf{H}$ that is arbitrarily close.

We call a forecasting protocol *regular* if $\mathbf{H} = \mathbf{H}^*$ for every $\mathbf{H}$ in $\mathcal{C}$. Because any forecasting protocol can be replaced with a regular one with the same lower and upper previsions (enlarge each $\mathbf{H}$ in $\mathcal{C}$ to $\mathbf{H}^*$), little generality is lost when we assume regularity. This assumption allows us to remove the "roughly speaking" from the statements that the lower prevision of x is the highest price at which Gambler can sell x and the upper prevision the lowest price at which he can buy it. It also allows us to say that $\mathbf{H}$ is completely determined by its upper previsions (and hence also by its lower previsions):

$$x \in \mathbf{H} \text{ if and only if } \overline{\mathbb{E}}_{\mathbf{H}}\,x \leq 0.$$

The condition $x \in \mathbf{H}$ says that House will give x to Gambler. The condition $\overline{\mathbb{E}}_{\mathbf{H}}\,x \leq 0$ says that House will sell x to Gambler for 0 or less.

3.2.2 Interpretation

Both interpretations of lower and upper previsions we discussed in §3.1 generalize to forecasting protocols in general. We can put ourselves in the role of House and say that our beliefs are expressed by the prices we are willing to pay—our lower and upper previsions. Or, as we prefer, we can put ourselves in the role of Gambler and subscribe to these prices in the sense of believing that they will not allow us to become very rich in the long run, no matter what strategy we follow.

The reference to the long run in the second interpretation must be understood in terms of a sequential version of our abstract protocol. If we suppose, for simplicity, that Reality and House have the same choices of belief cones and payoffs on every move, this sequential protocol can be written as follows:

SEQUENTIAL FORECASTING
Parameters: $\mathbf{R}$ and $\mathcal{C} \subseteq \mathcal{C}_{\mathbf{R}}$
Protocol:
$\mathcal{K}_0 := 1$.
For $n = 1, 2, \ldots$:
House announces $\mathbf{H}_n \in \mathcal{C}$.
Gambler announces $\mathbf{g}_n \in \mathbf{H}_n$.
Reality announces $\mathbf{r}_n \in \mathbf{R}$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + \mathbf{g}_n(\mathbf{r}_n)$.

The ambiguities we discussed in §3.1 also arise here. If we take House's point of view, we may or may not be categorical about our unwillingness to offer riskier payoffs than those in $\mathbf{H}_n$. If we take Gambler's point of view, we may be more or less certain about whether larger $\mathbf{H}_n$ would also satisfy Cournot's principle.

3.3 Walley's Updating Principle

We turn now to Peter Walley's updating principle. This principle can be shown to entail the rule of conditional probability when it is applied to subjective probability. Here we apply it to our abstract framework for lower and upper previsions.

TWO-STAGE FORECASTING
Parameters: $\mathbf{R}$, a disjoint partition $\mathbf{B}_1,\ldots,\mathbf{B}_k$ of $\mathbf{R}$, $\mathcal{C} \subseteq \mathcal{C}_{\mathbf{R}}$
Protocol:
At time 0:
House announces $\mathbf{H}_0 \in \mathcal{C}$.
Gambler announces $\mathbf{g}_0 \in \mathbf{H}_0$.
Reality announces $i \in \{1,2,\ldots,k\}$.
At time t:
House announces $\mathbf{H}_t \in \mathcal{C}_{\mathbf{B}_i}$.
Gambler announces $\mathbf{g}_t \in \mathbf{H}_t$.
Reality announces $\mathbf{r}_t \in \mathbf{B}_i$.
$\mathcal{K}_t := \mathcal{K}_0 + \mathbf{g}_0(\mathbf{r}_t) + \mathbf{g}_t(\mathbf{r}_t)$.

Because we are considering here how House should make his second move, we leave this move unconstrained by the protocol. In Sections 3.4 and 3.5 below we consider two specific alternatives for this choice. Here, House can choose any belief cone on the reduced sample space $\mathbf{B}_i$.

Walley's updating principle says that if House knows at time 0 that Reality's announcement of i will be House's only new information when he moves at time t, then at time 0, as he makes his move $\mathbf{H}_0$, House should intend for his move $\mathbf{H}_t$ to be the belief cone $\mathbf{w}_t^i$ on $\mathbf{B}_i$ given by

$$\mathbf{w}_t^i := \{\mathbf{g} : \mathbf{B}_i \mapsto \mathbb{R} \mid \mathbf{g}^{\uparrow} \in \mathbf{H}_0\}, \tag{6}$$

where $\mathbf{g}^{\uparrow}$ is defined by

$$\mathbf{g}^{\uparrow}(\mathbf{r}) := \begin{cases} \mathbf{g}(\mathbf{r}) & \text{if } \mathbf{r} \in \mathbf{B}_i \\ 0 & \text{if } \mathbf{r} \notin \mathbf{B}_i. \end{cases} \tag{7}$$

In words: House should intend to offer a given payoff at the second stage after Reality announces i if and only if he is already offering that payoff at the first stage contingent on that value of i. This produces simple formulae relating the new lower and upper previsions to the old ones:

$$\underline{\mathbb{E}}_{\mathbf{w}_t^i}\, x = \sup\{\alpha \mid \underline{\mathbb{E}}_{\mathbf{H}_0}(x-\alpha)^{\uparrow} \geq 0\} \tag{8}$$

and

$$\overline{\mathbb{E}}_{\mathbf{w}_t^i}\, x = \inf\{\alpha \mid \overline{\mathbb{E}}_{\mathbf{H}_0}(x-\alpha)^{\uparrow} \le 0\} \tag{9}$$

for every variable x on the reduced sample space $\mathbf{B}_i$.

3.4 Announcing Future Beliefs in Advance

We now consider House's second move being constrained by announcing his future beliefs in advance. The rule of conditional probability can be shown to be mandated by the principle of House's avoiding sure loss when he announces future subjective probabilities in advance. What if House announces in advance future beliefs that determine only lower and upper previsions?

ADVANCE FORECASTING
Parameters: $\mathbf{R}$, a disjoint partition $\mathbf{B}_1,\ldots,\mathbf{B}_k$ of $\mathbf{R}$, $\mathcal{C} \subseteq \mathcal{C}_{\mathbf{R}}$.
Protocol:
At time 0:
House announces $\mathbf{H}_0 \in \mathcal{C}$ and $\mathbf{H}_t^j \in \mathcal{C}_{\mathbf{B}_j}$ for $j = 1,\ldots,k$.
Gambler announces $\mathbf{g}_0 \in \mathbf{H}_0$.
Reality announces $i \in \{1,2,\ldots,k\}$.
At time t:
Gambler announces $\mathbf{g}_t \in \mathbf{H}_t^i$.
Reality announces $\mathbf{r}_t \in \mathbf{B}_i$.
$\mathcal{K}_t := \mathcal{K}_0 + \mathbf{g}_0(\mathbf{r}_t) + \mathbf{g}_t(\mathbf{r}_t)$.

Consider House's $\mathbf{H}_0$ and his $\mathbf{H}_t^j$ for some particular j. Suppose the variable $\mathbf{g}$ is in $\mathbf{H}_t^j$, but $\mathbf{g}^{\uparrow}$ is not in $\mathbf{H}_0$. Then it would make no difference in what Gambler can do if House were to enlarge $\mathbf{H}_0$ by adding $\mathbf{g}^{\uparrow}$ to it. He can already get the effect of $\mathbf{g}^{\uparrow}$ at time 0 by planning in advance to announce $\mathbf{g}$ at time t.

So we can assume, without changing what Gambler can accomplish, that if $\mathbf{g} \in \mathbf{H}_t^j$, then $\mathbf{g}^{\uparrow} \in \mathbf{H}_0$. This assumption implies $\mathbf{H}_t^j \subseteq \mathbf{w}_t^j$ by (6) and then

$$\underline{\mathbb{E}}_{\mathbf{H}_t^j} \le \underline{\mathbb{E}}_{\mathbf{w}_t^j} \tag{10}$$

by (4). *The lower prevision at time t that is foreseen and announced at time* 0 *should not exceed the lower prevision given by Walley's updating principle.* Writing simply $\underline{\mathbb{E}}_0\, x$ for $\underline{\mathbb{E}}_{\mathbf{H}_0}\, x$ and $\underline{\mathbb{E}}_t\, x$ for $\underline{\mathbb{E}}_{\mathbf{H}_t^i}\, x$ (the lower previsions that House's time-0 announcements imply for time 0 and t, respectively) and recalling (8), we can write (10) in the form

$$\underline{\mathbb{E}}_t\, x \le \sup\{\alpha \mid \underline{\mathbb{E}}_0(x-\alpha)^{\uparrow} \ge 0\}, \tag{11}$$

where x is a variable on the reduced sample space $\mathbf{B}_i$.

The argument for (11) relies on the new viewpoint developed in this article, according to which a person's uncertainty is measured by prices he believes he cannot beat, not by prices he is disposed to offer. We expect (11) to hold because if it did not, the time 0 lower previsions would need to be increased to reflect stronger betting offers that Gambler cannot beat. Strictly speaking, of course, talk about Gambler not being able to beat given prices is talk about the long run, and so a complete exposition of the argument would involve a sequential protocol. We leave this further elaboration of the argument to the reader.

The argument does *not* rely on any assumption about exact information. Possibly House and Gambler will learn more than $\mathbf{B}_i$ by time t. $\underline{\mathbb{E}}_t\, x$, in (11), is not necessarily the lower prevision at time t. It is merely the lower prevision at time t to which House commits himself at time 0. This commitment does not exclude the possibility that House and Gambler will acquire additional unanticipated information and that House will thus offer Gambler more variables at time t than those to which he committed himself at time 0. In this case, the actual lower prevision for x at time t may come out higher than $\underline{\mathbb{E}}_{\mathbf{H}_t^i}\, x$ and even higher than $\underline{\mathbb{E}}_{\mathbf{w}_t^i}\, x$.

For planning at time 0, we are interested in what we can count on already at time 0. This is why the upper bound in (11) is interesting. When time t comes around, positive unanticipated information may lead us to give x a lower prevision exceeding this upper bound, but there is also the possibility of negative unanticipated information, and the upper bound can be thought of as telling us how conservative we need to be in our advance commitments in order to hedge against the possible negative information.

3.5 Updating with Exact Information

Although the case in Section 3.4 above where commitments are made in advance in the face of possible *unanticipated* new information seems to us to have greater practical importance, it is also of interest to consider the case where new information is *anticipated exactly*. This is where Walley's principle applies.

Extending the protocol of §3.3, we obtain the following sequential protocol:

SEQUENTIAL TWO-STAGE FORECASTING
- $\mathcal{K}_0 := 1$.
- For $n = 1, 2, \dots$
  - At time n:
    - House announces $\mathbf{H}_{n0} \in \mathcal{C}$.
    - Gambler announces $\mathbf{g}_{n0} \in \mathbf{H}_{n0}$.
    - Reality announces $i_n \in \{1, 2, \dots, k\}$.
  - At time $n + 1/2$:
    - House announces $\mathbf{H}_{n1} \in \mathcal{C}_{\mathbf{B}_{i_n}}$.
    - Gambler announces $\mathbf{g}_{n1} \in \mathbf{H}_{n1}$.
    - Reality announces $\mathbf{r_n} \in \mathbf{B}_{i_n}$.
  - $\mathcal{K}_n := \mathcal{K}_{n-1} + \mathbf{g}_{n0}(\mathbf{r_n}) + \mathbf{g}_{n1}(\mathbf{r_n})$.

First, we make the following assumptions:

1. House's $\mathbf{H}_{n0}$ satisfy Cournot's principle.
2. House agrees in advance to follow Walley's updating principle: $\mathbf{H}_{n1} = \mathbf{w}_n^{i_n}$, where $\mathbf{w}_n^j := \{\mathbf{g} : \mathbf{B}_j \mapsto \mathbb{R} \mid \mathbf{g}^{\uparrow} \in \mathbf{H}_{n0}\}$.
3. The only new information Gambler acquires between his move at time n and his move at time $n+1/2$ is Reality's choice of of i_n. (By the preceding assumption, he already knows House's move $\mathbf{H}_{n1}$.)
4. Reality disregards Gambler's moves when she chooses her own moves.

Will all of House's announcements (the $\mathbf{H}_{n0}$ and $\mathbf{H}_{n1}$) satisfy Cournot's principle as a group? It is reasonable to conclude that they will. If they did not, then Gambler would have a bankruptcy-free strategy $\mathcal{S}$ that would make him infinitely rich. This strategy would specify $\mathbf{g}_{n0} \in \mathcal{C}$ for $n = 1, 2, \ldots$ and $\mathbf{g}_{n1}^j \in \mathbf{w}_n^j$ for $n = 1, 2, \ldots$ and $j = 1, \ldots, k$. Because Reality's moves do not depend on what Gambler does (Assumption 4) and House will follow Walley's recommendation for $\mathbf{H}_{n1}$ (Assumption 2), Gambler has a strategy $\mathcal{S}'$ for choosing the $\mathbf{g}_{n0}$ alone that makes his capital grow exactly as $\mathcal{S}$ does: to duplicate the effect of $\mathcal{S}$'s move $\mathbf{g}_{n1}$, he adds $(\mathbf{g}_{n1}^j)^{\uparrow}$ to $\mathcal{S}$'s $\mathbf{g}_{n0}$ for $j = 1, \ldots, k$. This strategy does not require knowledge of i_n, and so Gambler would have the information needed to implement it (Assumption 3). So $\mathcal{S}'$ would also make Gambler infinitely rich, contradicting Assumption 1.

This result is a long-run justification for Walley's updating principle in its full generality.

4 Summary and Prospects

In this article, we set forth a new way of understanding probabilities and previsions in which we considered Gambler's viewpoint, and adopted Cournot's principle, in a series of game-theoretic protocols.

The proper handling of updating depends on whether we can exactly anticipate new information.

- We learned in §3.5 that if we can exactly anticipate new information—i.e., if we have an exhaustive advance list $\mathbf{B}_1, \ldots, \mathbf{B}_k$ of possibilities for exactly what all our new information will be, then we can follow Walley's updating principle, deriving new lower previsions from old ones using the formula

$$\underline{\mathbb{E}}_t\, x = \sup\{\alpha \mid \underline{\mathbb{E}}_0(x-\alpha)^{\uparrow} \geq 0\}. \quad (12)$$

- We learned in §3.4 that if we cannot exactly anticipate new information, but we do know that we will learn which of the mutually exclusive events

$\mathbf{B}_1, \ldots, \mathbf{B}_k$ has happened, and we commit ourselves in advance to lower previsions that depend on which $\mathbf{B}_i$ happens, then these preannounced lower previsions should satisfy the upper bound

$$\underline{\mathbb{E}}_t\, x \leq \sup\{\alpha \mid \underline{\mathbb{E}}_0 (x - \alpha)^{\uparrow} \geq 0\}. \tag{13}$$

The requirement of exact new information is very strong. The inequality (13) depends only on the weaker condition that we learn which of the $\mathbf{B}_1, \ldots, \mathbf{B}_k$ happens. There is no requirement that this be all we learn. On the other hand, the inequality only bounds the new lower prevision that can be guaranteed at the outset, at the planning stage. Unanticipated information may produce a higher lower prevision.

In this article, we have invoked Cournot's principle using a relatively simple protocol, in which Reality has a binary choice at each step. This principle can also be adopted, however, when Reality sometimes has more than two choices, and when the choices available to her may depend on what she has done previously. This brings us to the generality of an event tree [5], offering additional flexibility that is needed in planning. Here it may be convenient to suppress the role of House in favor of a formal rule for determining the probabilities offered to Gambler, and to allow for unanticipated information and the refinement of beliefs. We explore these questions in [7].

References

[1] DE COOMAN, G., AND WALLEY, P. The imprecise probabilities project. `http://ippserv.rug.ac.be/home/ipp.html`.

[2] DE FINETTI, B. *Teoria Delle Probabilità*. Einaudi, Turin, 1970. An English translation, by Antonio Machi and Adrian Smith, was published as *Theory of Probability* by Wiley (London, England) in two volumes in 1974 and 1975.

[3] DUBOIS, D., AND PRADE, H. *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. Plenum Press, New York, 1988.

[4] SHAFER, G. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, New Jersey, 1976.

[5] SHAFER, G. *The Art of Causal Conjecture*. The MIT Press, Cambridge, Massachusetts, 1996.

[6] SHAFER, G. Causality and responsibility. *Cardozo Law Review 22* (2001), 101–123.

[7] SHAFER, G., GILLETT, P. R., AND SCHERL, R. B. A new understanding of subjective probability and its generalization to lower and upper prevision. *International Journal of Approximate Reasoning 33* (2003), 1–49.

[8] SHAFER, G., AND VOVK, V. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001. For more information on this book, see `http://www.cs.rhul.ac.uk/~vovk/book/`.

[9] SMETS, P. Varieties of ignorance. *Information Sciences 57–58* (1991), 135–144.

[10] SMITH, C. A. B. Consistency in statistical inference and decision. *Journal of the Royal Statistical Society, Series B 23* (1961), 1–25.

[11] SMITH, C. A. B. Personal probabilities and statistical analysis. *Journal of the Royal Statistical Society, Series A 128* (1965), 469–499.

[12] SRIVASTAVA, R. P., AND MOCK, T. J., Eds. *Belief Functions in Business Decisions*. Springer Verlag, New York, 2002.

[13] WALLEY, P. Coherent lower (and upper) probabilities. Tech. Rep. 22, Department of Statistics, University of Warwick, Coventry, UK, 1981.

[14] WALLEY, P. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.

[15] WALLEY, P. Inferences from multinomial data: learning about a bag of marbles (with discussion). *Journal of the Royal Statistical Society, Series B 58* (1996), 3–57.

[16] WALLEY, P. Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning 24* (2000), 125–148.

[17] WILLIAMS, P. M. Coherence, strict coherence and zero probabilities. In *Fifth International Congress of Logic, Methodology and Philosophy of Science* (1975), pp. VI, 29–33.

[18] WILLIAMS, P. M. Notes on conditional previsions. Research Report, School of Mathematical and Physical Sciences, University of Sussex, February 1975.

[19] WILLIAMS, P. M. Indeterminate probabilities. In *Formal Methods in the Methodology of Empirical Sciences*, M. Przełęcki, K. Szaniawski, and R. Wojcícký, Eds. Ossolineum & Reidel, 1976, pp. 229–246.

Glenn Shafer is on the faculty of the Rutgers Business School—Newark and New Brunswick. He may be contacted at gshafer@andromeda.rutgers.edu.

Peter R. Gillett is on the faculty of the Rutgers Business School—Newark and New Brunswick. He may be contacted at gillett@business.rutgers.edu.

Richard B. Scherl is on the faculty of Monmouth University's Department of Computer Science. He may be contacted at rscherl@monmouth.edu.