

To What Extent Is Foresight Luck?

ABSTRACT

This study assesses the relative plausibility of cognitive, evolutionary, and chance theories of foresight by quantifying the extent to which foresight is luck. Results from a regression-toward-the-mean approach applied to foresight estimates derived from a beta item response theory model using data from 265 forecasters who made 40,514 forecasts on 212 questions over two years in a forecasting tournament suggest that foresight is on average 32% luck. This means that the foresight of an individual forecaster regresses on average by 32% toward the mean foresight of all forecasters over one year. This regression toward the mean is stronger among top forecasters (35%) and forecasters with below average foresight (47%) but weaker for the average forecaster (23%). Despite this regression toward the mean, the top forecasters in year 1 did not display worse foresight than other forecasters in year 2. I discuss the implications of these findings for cognitive, evolutionary, and chance theories of foresight.

Keywords:

foresight; cognition; evolution; luck; item response theory; regression toward the mean

To What Extent Is Foresight Luck?

Theories in strategy differ in the extent to which they consider foresight—“[the] ability to make predictions about strategic outcomes” (Csaszar & Laureiro-Martínez, 2018, p. 513)—to be luck. Cognitive theories suggest that strategists possess “boundedly rational foresight” (Gavetti, Greve, Levinthal, & Ocasio, 2012: 9) that is grounded in their mental representations of the world rendering foresight mainly a skill (Csaszar, 2018; Csaszar & Laureiro-Martínez, 2018; Gavetti & Levinthal, 2000). Evolutionary theories (Levinthal, 2021; Murmann, 2003; Nelson & Winter, 1982) acknowledge that strategists may exhibit myopic foresight grounded in their cognitive preadaptation that enables them to discover opportunities serendipitously (Denrell, Fang, & Winter, 2003; Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016). Chance theories question the existence of managerial foresight by emphasizing the role of random processes rendering accurate foresight mostly luck (Denrell & Fang, 2010; Denrell, Fang, & Liu, 2015; Liu, 2020; Liu & de Rond, 2016).

However, the relative plausibility of cognitive, evolutionary, and chance theories of foresight has received scant empirical attention. One explanation for this may be that unlike research on the relative importance of organizational and industry effects (McGahan & Porter, 1997; Rumelt, 1991; Wang, 2023), it is empirically challenging to directly examine the relative explanatory power of these theories in a single study. Such a study would require the measurement of the content and structure of mental representations of strategists (Csaszar & Laureiro-Martínez, 2018; Csaszar & Ostler, 2020), the cognitive preadaptation of strategists in relation to their environment and changes thereof over time (Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016), and an exhaustive set of variables controlling for alternative accounts of foresight as well as the comparison of alternative random processes to evaluate the role of chance in foresight

(Denrell et al., 2015). Despite these empirical challenges, it is important to assess the relative plausibility of cognitive, evolutionary, and chance theories of foresight because otherwise we may not fully understand to what extent foresight—and by extension the strategist—may be a source of competitive advantage resulting in superior organizational performance.

Therefore, the purpose of this paper is to assess the relative plausibility of cognitive, evolutionary, and chance theories of foresight by quantifying the extent to which foresight is luck. Although cognitive, evolutionary, and chance theories of foresight propose different explanations for foresight rendering a joint examination of these explanations unfeasible, these theories may be compared based on their implications for the extent of luck in foresight. Cognitive theories suggest that foresight is mainly skill and only somewhat luck. Evolutionary theories consider foresight to be somewhat skill and mainly luck. Chance theories consider foresight to be mostly luck. Quantifying the extent to which foresight is luck may, therefore, provide important insights into the relative explanatory power of different theories of foresight. To this end, I apply a regression toward the mean approach (Liu, 2020; Samuels, 1991; Schmittlein, 1989) to foresight estimates derived from a beta item response theory model using data from 265 forecasters who made 40,514 forecasts on 212 questions over two years in a forecasting tournament. The results suggest that foresight is on average 32% luck. This means that the foresight of an individual forecaster regresses on average by 32% toward the mean foresight of all forecasters over one year. This regression toward the mean is stronger among top forecasters (35%) and forecasters with below average foresight (47%) but weaker for the average forecasters (23%). Despite this regression toward the mean, the top forecasters in year 1 did not display worse foresight than other forecasters in year 2.

This study contributes to the strategy literature by assessing the relative plausibility of cognitive, evolutionary, and chance theories of foresight. Although foresight may be somewhat

luck, it appears to be mainly skill. This is consistent with cognitive theories proposing that managers may exhibit boundedly rational foresight grounded in their mental representations of the world (Csaszar, 2018; Csaszar & Laureiro-Martínez, 2018; Gavetti & Levinthal, 2000). Furthermore, the fact that forecasters were able to accurately predict future states of the world three to six months and in some cases up to one year in advance on a wide range of questions suggests that this boundedly rational foresight was less myopic than assumed by evolutionary theories, albeit skillful foresight may still be partially driven by cognitive preadaptation (Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016). At the same time, this result cautions against ascribing omniscience to strategists who displayed foresight historically. Although such strategists may continue to exhibit foresight, their foresight was partly luck (Denrell & Fang, 2010; Denrell et al., 2015; Liu, 2020; Liu & de Rond, 2016). Taken together, this study suggests that the strategist's foresight may be a source of competitive advantage with potentially positive effects on organizational performance.

The remainder of this paper is structured as follows: First, I will briefly outline cognitive, evolutionary, and chance theories of foresight. Second, I will explain the regression-toward-the-mean approach, introduce a beta item response theory model to measure foresight, and describe the data and analytical approach of this study. Third, I will report the results of this paper. Last, I will discuss the theoretical and practical implications of this study, acknowledge its limitations, and articulate opportunities for future research.

THEORETICAL BACKGROUND

In the following, I will briefly outline cognitive, evolutionary, and chance theories of foresight. For each theory, I will first describe the nature of the strategist's foresight and why strategists may or may not display foresight. Second, I will derive to what extent these theories

consider foresight to be luck. Last, I will explain how the relative plausibility of cognitive, evolutionary, and chance theories of foresight can be assessed based on their different implications for the extent to which foresight is luck. Table 1 summarizes cognitive, evolutionary, and chance theories of foresight.

 Insert Table 1 about here

Cognitive Theories of Foresight

Cognitive theories of foresight suggest that strategists possess “boundedly rational foresight” (Gavetti et al., 2012: 9) that is grounded in their mental representations of the world (Csaszar, 2018; Csaszar & Laureiro-Martínez, 2018; Gavetti & Levinthal, 2000). A mental representation is a “model of reality held in the mind of an individual, who can use this representation to generate predictions about reality” (Csaszar & Levinthal, 2016: 2031f.). Based on their predictions, strategists choose strategic decision alternatives they expect to satisfice organizational goals (March & Simon, 1958; Simon, 1947). In other words, mental representations guide strategists in their search of strategic initiatives that improve performance (Gavetti & Levinthal, 2000; Gavetti, Levinthal, & Rivkin, 2005) by enabling them to predict future states of the world. Therefore, accurate mental representations underlying the strategist’s foresight have been theorized to positively affect organizational performance (Csaszar, 2018; Gavetti, 2012).

From a cognitive perspective, foresight is considered to be mainly a skill as mental representations can be learned and adapted. Strategists may learn and use complex representations if they do not yet understand the drivers of organizational performance in an environment, whereas they may learn and use simple representations if they do (Csaszar & Ostler, 2020). Given the importance of foresight in strategic decision making, strategists should actively search for mental representations that enable them to predict future states of the world (Csaszar & Levinthal, 2016).

Strategists may learn accurate mental representations in strategy courses (Heshmati & Csaszar, 2024) that teach frameworks like Porter's Five Forces (Porter, 1979, 2008). By internalizing and adapting such external representations (Csaszar, Hinrichs, & Heshmati, 2024) for their specific organization, strategists may make better strategic decisions. That being said, the strategist's cognitive, informational, and time limitations (March & Simon, 1958; Simon, 1947) as well as the inherent randomness of the world (Packard & Clark, 2020; Schuler, Murmann, Beisemann, & Satopää, 2024) mean that uncertainty about future cannot be fully resolved rendering foresight somewhat luck. Taken together, cognitive theories of foresight suggest that strategists possess boundedly rational foresight grounded in their mental representations of the world rendering foresight mainly a skill but also somewhat luck.

Evolutionary Theories of Foresight

Evolutionary theories (Levinthal, 2021; Murmann, 2003; Nelson & Winter, 1982) suggest that strategists may exhibit myopic foresight because they are cognitively preadapted to a future environment (Denrell et al., 2003; Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016). Strategists are cognitively preadapted to a future environment if they possess deep domain knowledge that enables them to see unique opportunities that other strategists cannot see once the future environment realizes (Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016). Thus, foresight represents a form of myopic, serendipitous exploitation (Busch, 2020, 2024; Denrell et al., 2003) of deep domain knowledge in a changed environment rather than the active anticipation of distant future states of the world (Levinthal, 2017, 2021). While evolutionary theories acknowledge that foresight may contribute to organizational performance serendipitously, organizational routines that are well adapted to the environment are considered to be the main driver of organizational performance (Levinthal, 2021; Murmann, 2003; Nelson & Winter, 1982).

From an evolutionary perspective, foresight is somewhat skill and largely luck. Foresight is somewhat skill because strategists may deepen and broaden their domain knowledge to become cognitively preadapted to an increasing number of potential future environments (Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016). Furthermore, strategists still need to attend to changes in the environment and exert effort to seize the opportunities in the environment that they see based on their cognitive preadaptation (Busch, 2020, 2024; Denrell et al., 2003). At the same time, foresight is mainly luck because it remains a form of myopic, serendipitous exploitation of changing environments that are mostly not under the control of the strategist but determine what deep domain knowledge is valuable (Levinthal, 2017, 2021). Overall, evolutionary theories of foresight suggest that strategists may exhibit myopic foresight because they are cognitively preadapted to a future environment rendering foresight somewhat skill but largely luck.

Chance Theories of Foresight

Chance theories suggest that true foresight may not exist because accurate predictions of future states of the world by strategists are due to chance (Denrell et al., 2015; Liu, 2020; Liu & de Rond, 2016). Strategists who accurately predict future states of the world may not exhibit foresight but were simply lucky (Denrell & Liu, 2012). Therefore, accurate predictions of future states of the world—especially of extreme events—may not be viewed as evidence of the strategist's foresight but rather as a form of bad judgment because such predictions systematically neglect the role of chance in the world (Denrell & Fang, 2010). Consequently, chance theories suggest that foresight only plays a minor role in explaining organizational performance. In fact, superior organizational performance itself—especially over long periods of time—is considered to be the result of random processes rather than the strategist's actions (Denrell, 2004; Denrell, Fang, & Zhao, 2013).

From a chance perspective, foresight is considered to be mainly luck as accurate predictions of future states of the world are due to chance. Therefore, foresight should regress toward the mean in one of two systematic manners over time (Denrell, Fang, & Liu, 2019; Denrell & Liu, 2012; Liu, 2020): First, strategists who accurately (inaccurately) predicted future states of the world initially may predict future states of the world less (more) accurately later (weak pattern of luck). Second, strategists who predicted future states of the world most accurately (inaccurately) initially may predict future states of the world less (more) accurately than moderately accurate (inaccurate) strategists later (strong pattern of luck). In other words, strategists who made extremely accurate (inaccurate) predictions regress more strongly toward the mean than strategists who made moderately accurate (inaccurate) predictions. Irrespective of whether foresight follows the weak or strong pattern of luck, chance theories imply that foresight fully regresses to the mean sooner or later and thus, is mostly luck.

Assessing the Relative Plausibility of Foresight Theories: To What Extent Is Foresight Luck?

Cognitive, evolutionary, and chance theories propose plausible explanations for foresight. However, the relative plausibility of cognitive, evolutionary, and chance theories of foresight has received scant empirical attention. While it is empirically challenging to examine the relative explanatory power of these theories in a single study directly, the relative explanatory power may be assessed based on the theories' differing implications for the extent to which foresight is luck. Therefore, one potential approach to better understand the relative plausibility of cognitive, evolutionary, and chance theories of foresight is to quantify the extent to which foresight is luck. In the remainder of this paper, I will apply a regression-toward-the-mean approach to foresight estimates derived from a beta IRT model using data from 265 forecasters who made 40,514 forecasts on 212 questions over two years in a forecasting tournament to quantify the extent to

which foresight is luck.

METHODS

In the following, I will first explain the regression-toward-the-mean approach to quantify the extent to which foresight is luck. Second, I will introduce a beta item response theory (IRT) model to measure foresight. Third, I will describe the data from the forecasting tournament. Finally, I will explain how I analyzed the data from the forecasting tournament using the beta IRT model and how I implemented the regression-toward-the-mean approach to quantify luck in foresight.

A Regression-Toward-the-Mean Approach to Quantify the Extent of Luck in Foresight

I used a regression-toward-the-mean approach to quantify the extent to which foresight is luck (Liu, 2020; Samuels, 1991; Schmittlein, 1989). To illustrate the intuition behind the regression-toward-the-mean approach, suppose we ask 100 strategists to forecast 10 questions about different future states of the world. After the true future states of the world realized, we assign each strategist one score representing their foresight at time 0. Then, we ask the same 100 strategists again to forecast 10 different questions about future states of the world. After the true future states of the world realized, we assign each strategist a second score representing their foresight at time 1.

If foresight is 100% skill, strategists receive the same foresight score at times 0 and 1 (e.g., forecaster A may receive a score of -1 at time 0 and 1, whereas forecaster B may receive a score of 1 at time 0 and 1). Assuming the scores at times 0 and 1 to be identically distributed (i.e., their shapes, means, and variances are identical) and standardized (i.e., mean = 0 and variance = 1), regressing the foresight score at time 1 on the foresight score at time 0 would yield an intercept of 0 and a regression coefficient of 1. Graphically, this corresponds to a 45 degree line intersecting the origin (i.e., $y = x$). In contrast, if foresight is 100% luck, strategists would receive different

scores at times 0 and 1, and the scores would be unrelated to each other. In other words, the foresight score at time 0 would contain no information about the foresight score at time 1. Assuming the scores at times 0 and 1 to be identically distributed and standardized, regressing the foresight score at time 1 on the foresight score at time 0 would yield an intercept of 0 and a regression coefficient of 0. Graphically, this corresponds to the x-axis. If both luck and skill play a role in foresight, strategists would receive different scores at times 0 and 1, but the scores would be related to each other. That is, the first score would contain information about the second score. Assuming the scores at times 0 and 1 to be identically distributed and standardized, regressing the foresight score at time 1 on the foresight score at time 0 would yield an intercept of 0 and a regression coefficient between 0 and 1. Graphically, this corresponds to a line between 45 and 0 degrees intersecting the origin (cf. Figure 1). Note that the regression coefficient corresponds to the correlation between the standardized foresight scores at times 0 and 1 in all these cases.

The regression coefficients (or the correlation between the foresight scores) in these examples may be interpreted as a measure of the percentage of skill involved in foresight. Conversely, one minus the regression coefficient may be interpreted as a measure of the percentage of luck involved in foresight. In other words, the regression toward the mean (Samuels, 1991) can be estimated by regressing the foresight score at time 1 on the foresight score at time 0. This assumes that the relationship between the foresight scores is linear.

However, if the foresight scores exhibit a non-linear (e.g., a quadratic or cubic) relationship, the role of luck in foresight differs depending on the foresight of forecasters at time 0. In these cases, it is worthwhile to explore the regression toward the mean in different subgroups of foresight at time 0 to assess the extent to which the regression toward the mean (i.e., the role of luck) is stronger or weaker in these subgroups (Liu, 2020). This can be done based on the reversion toward

the mean, which generalizes the notion of the regression toward the mean (see Samuels, 1991 for a formal definition): Suppose we split the distribution of foresight scores at time 0 into an upper and lower part based on a cutoff. Assuming that the cutoff lies above the mean foresight of the forecasters at time 0 (i.e., above 0 if the foresight scores are standardized), reversion toward the mean occurs if the mean foresight scores of forecasters in the upper part of the distribution at time 0 revert back toward the mean foresight of the forecasters at time 1 (i.e., toward 0 if the foresight scores are standardized). Analogously, assuming that the cutoff lies below the mean foresight of the forecasters at time 0 (i.e., below 0 if the foresight scores are standardized), reversion toward the mean occurs if the mean foresight scores of forecasters in the lower part of the distribution at time 0 revert back toward the mean foresight of the forecasters at time 1 (i.e., toward 0 if the foresight scores are standardized). This reversion toward the mean can be quantified by dividing the difference between the mean of the foresight scores of forecasters above (below) the cutoff at time 0 and the mean of the foresight scores of these forecasters at time 1 by the mean of the foresight scores of these forecasters at time 0 (cf. Figure 3). This yields the extent to which foresight is luck in percent (see data analysis section for a more formal description).

This regression-toward-the-mean-approach makes four assumptions: First, the foresight of the individual strategist is stable over time. In other words, the foresight of a strategists should not increase or decrease over time as such changes would be interpreted as evidence for luck. This assumption needs to be theoretically justified and is difficult to evaluate by statistical tests or plots. Second, the foresight scores must be identically distributed (Samuels, 1991). Third, the foresight scores must exhibit weak positive dependence—the means of the foresight scores must regress toward the mean of the foresight scores of all forecasters but not beyond (Samuels, 1991). This can be graphically evaluated, for example, by plotting the regression fitted to the data. Fourth, the

measurement of foresight should be free from measurement error. Measurement errors attenuate the relationship between the foresight estimates (Fuller & Hidirolou, 1978), which inflates the estimate to which foresight is luck. Therefore, it is important and worthwhile to optimize the measurement accuracy of foresight to obtain the most accurate quantification possible for the extent to which foresight is luck. Item response theory models represent one way how foresight can be accurately measured. In the following, I will introduce a beta IRT model of probabilistic forecasts to measure foresight as accurately as possible.

A Beta Item Response Theory Model of Probabilistic Forecasts to Measure Foresight

This section introduces a novel beta item response theory (IRT) model of probabilistic forecasts to measure foresight. Probabilistic item response theory (IRT) models assume that the responses of persons to items can be described as a function of their latent abilities, the item characteristics, and random factors (Cai, Choi, Hansen, & Harrell, 2016; Carroll, Primo, & Richter, 2016; Foster, Min, & Zickar, 2017; Hambleton & Swaminathan, 2013; Hambleton, Swaminathan, & Rogers, 1991; Van der Linden & Hambleton, 1997). The beta IRT model describes the probabilistic forecasts of persons on a set of questions over time as a probabilistic function of the latent foresight of persons, the difficulties of the questions, the discrimination of the questions, and time effects. Formally, the beta IRT model may be written as a multilevel model (Hox, Moerbeek, & Schoot, 2017; Raudenbush & Bryk, 2002) with two levels, where the level-1 model can be defined as

$$\begin{aligned}
 Y_{iqt} &\sim \text{Beta}(\mu_{iqt}, \phi_{qt}) \\
 \mu_{iqt} &= \text{expit}(\beta_{\mu 0q} + \beta_{\mu 1q} t_{iqt} + \lambda_q \theta_i) \\
 \phi_{qt} &= \exp(\beta_{\phi 0q} + \beta_{\phi 1q} t_{qt}),
 \end{aligned} \tag{1}$$

where the probabilistic forecast¹ on the true category Y_{iqt} of forecaster $i = 1, \dots, N$ ($N \in \mathbb{Z}^+$ the number of forecasters) on question $q = 1, \dots, Q$ ($Q \in \mathbb{Z}^+$ the number of questions) at time $t_q = 1, 2, \dots, T_q$ ($T_q \in \mathbb{Z}^+$ the number of days question q was forecastable) follows a beta distribution $Beta(\mu_{iqt}, \phi_{qt})$ with mean μ_{iqt} and precision ϕ_{qt} .

The mean μ_{iqt} is modeled as a probabilistic function of the foresight θ_i of forecaster i and the time variable t_{iqt} indicating at what time t forecaster i made their probabilistic forecast Y_{iqt} on question q , where $t_{iqt} = 0$ is the first day question q was forecastable. $\beta_{\mu 0q}$ is the difficulty of question q at time $t_{iqt} = 0$ indicating how easy it was to answer a question correctly at the outset, where larger $\beta_{\mu 0q}$ indicate a higher average probability forecast on the true category of question q ; $\beta_{\mu 1q}$ refers to the time effect controlling for changes in the difficulty of question q over time; and λ_q represents the discrimination of question q indicating how well the question differentiates between different degrees of foresight, which can be reinterpreted as the effect of the latent ability θ_i on the forecast on the true category of question q . In other words, the item discrimination λ_q determines how useful the latent foresight θ_i is for answering item q correctly. The $expit(x) = \frac{\exp(x)}{1+\exp(x)}$ function ensures that the number space of the linear combination $(\beta_{\mu 0q} + \beta_{\mu 1q}t_{iqt} + \lambda_q\theta_i) \in \mathbb{R}$ matches the number space of $\mu_{iqt} \in (0,1)$. The precision ϕ_{qt} (i.e., the inverse of the variance) is modeled as a probabilistic function of the time variable t_{iqt} . $\beta_{\phi 0q}$ is the precision of question q at time $t_{iqt} = 0$; and $\beta_{\phi 1q}$ refers to the time effect controlling for changes in the precision of question q over time. The natural exponential function $\exp(x) = e^x$ ensures that the number space of the linear combination $(\beta_{\phi 0q} + \beta_{\phi 1q}t_{qt}) \in \mathbb{R}$ matches the number space of

¹ Extreme probabilistic forecasts of 1 and 0 are transformed to 0.999 and 0.001 to ensure that the number space of the probabilistic forecasts $f \in [0,1]$ matches the number space of the beta distribution $Beta(\mu, \phi) \in (0,1)$.

$$\phi_{qt} \in \mathbb{R}^+.$$

The second level of the beta IRT model describing the parameters varying across questions can be formally defined as

$$\begin{aligned}\beta_{\mu 0q} &= \beta_{\mu 00} + u_{\mu 0q} \\ \beta_{\mu 1q} &= \beta_{\mu 10} + u_{\mu 1q} \\ \lambda_q &= \lambda_0 + v_q \\ \beta_{\phi 0q} &= \beta_{\phi 00} + u_{\phi 0q} \\ \beta_{\phi 1q} &= \beta_{\phi 10} + u_{\phi 1q}\end{aligned}\tag{2}$$

where the random effects of the corresponding model parameters describing how each question q or person i differs from its corresponding mean parameter are assumed to be multivariate normally distributed with

$$\begin{pmatrix} u_{\mu 0q} \\ u_{\mu 1q} \\ v_q \\ u_{\phi 0q} \\ u_{\phi 1q} \end{pmatrix} \sim \mathcal{N} \left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \tau_{u\mu 0}^2 & & & & \\ \tau_{u\mu 0, u\mu 1} & \tau_{u\mu 1}^2 & & & \\ \tau_{u\mu 0, v} & \tau_{u\mu 1, v} & \tau_v^2 & & \\ \tau_{u\mu 0, u\phi 0} & \tau_{u\mu 1, u\phi 0} & \tau_{v, u\phi 0} & \tau_{u\phi 0}^2 & \\ \tau_{u\mu 0, u\phi 1} & \tau_{u\mu 1, u\phi 1} & \tau_{v, u\phi 1} & \tau_{u\phi 0, u\phi 1} & \tau_{u\phi 1}^2 \end{pmatrix} \right]\tag{3}$$

and the second level of the beta IRT model describing the foresight parameter varying across forecasters can be formally written as

$$\theta_i = \theta_0 + z_i,\tag{4}$$

where the random effect pertaining to the level-2 model of the latent foresight parameter is assumed to be normally distributed with

$$z_i \sim \mathcal{N}(0, \tau_z^2),\tag{5}$$

where τ_z^2 is fixed to 1 to identify the model.

$\beta_{\mu 00}$, $\beta_{\mu 1q}$, λ_q , $\beta_{\phi 0q}$, $\beta_{\phi 1q}$, and θ_0 denote the means of the corresponding parameters

across all items Q or all persons N , respectively. $\tau_{u\mu 0}^2$, $\tau_{u\mu 1}^2$, τ_v^2 , $\tau_{u\phi 0}^2$, and $\tau_{u\phi 1}^2$ describe the variances of the corresponding random effects. $\tau_{u\mu 0, u\mu 1}$, $\tau_{u\mu 0, v}$, $\tau_{u\mu 0, u\phi 0}$, $\tau_{u\mu 0, u\phi 1}$, $\tau_{u\mu 1, v}$, $\tau_{u\mu 1, u\phi 0}$, $\tau_{u\mu 1, u\phi 1}$, $\tau_{v, u\phi 0}$, $\tau_{v, u\phi 1}$, and $\tau_{u\phi 0, u\phi 1}$ refer to the covariations between the corresponding parameters.

The beta IRT model assumes that foresight varies across persons, whereas the difficulty, discrimination, precision, and the respective time effects vary across questions. The beta IRT model also assume the foresight of forecasters to be unidimensional. This means that any systematic differences in the persons' forecasts on the questions result from the same underlying ability like foresight (Hambleton & Swaminathan, 2013). Although the foresight of forecasters may in principle be multidimensional (i.e., forecasters possess different foresight in two or more domains like economics and politics), forecasting tournaments usually consist of questions designed to evaluate the foresight of forecasters in one specific domain. If the latent foresight parameter θ_i is unidimensional and the question responses are independent from each other conditional on the latent foresight parameter θ_i (i.e., the residuals of the questions in the model with a unidimensional foresight parameter θ_i are uncorrelated), the questions are assumed to be locally independent (Hambleton & Swaminathan, 2013). If the assumptions of unidimensionality and local independence are fulfilled, persons can be compared on the latent ability even if they answered different questions as each question measures the same foresight dimension.

While the difficulty and discrimination parameters consider the influence of question characteristics on the forecasts of persons, the time effects in the mean and precision models control for changes in the mean and precision parameters over time. Controlling for time effects is consistent with the idea that the quantity and quality of signals increases over time, which renders accurately forecasting future states of the world easier over time (Schuler et al., 2024).

Although various time effects can be conceived (Bo, Budescu, Lewis, Tetlock, & Mellers, 2017; Merkle, Steyvers, Mellers, & Tetlock, 2016, 2017), research suggests that the time effects in the beta IRT model control for various signal trajectories (Schuler et al., 2024).

IRT models can measure latent traits like foresight more accurately than strictly proper scoring rules as they extract more information about the foresight of persons resulting in somewhat better estimates compared to strictly proper scoring rules following the logic of classical test theory (Hambleton & Swaminathan, 2013). However, this comes at substantially greater computational cost. Therefore, IRT models of probabilistic forecasts are usually only used when measurement accuracy is of paramount importance. As the regression-toward-the-mean approach requires the best measurement accuracy possible to arrive at the most accurate estimate of the extent to which foresight is luck, it was reasonable to estimate foresight with the beta IRT model in this study.

Data

I applied the regression toward the mean approach to foresight estimates derived from the beta IRT model using data from 265 forecasters who made 40,514 forecasts on 212 questions over two years in a forecasting tournament. Forecasting tournaments are prediction polls in which participants make probabilistic forecasts on clearly defined questions about the future state of the real world at a specific point in time (Atanasov et al., 2017; Kapoor & Wilde, 2023; Tetlock & Gardner, 2016; Tetlock, Mellers, Rohrbaugh, & Chen, 2014). For each question, the participants precisely quantify their beliefs about the future by allocating probabilities to a finite number of response categories representing an exhaustive set of answers to the question. For example, forecasters may forecast the question *Will OpenAI release ChatGPT 5.0 by Dec 31st, 2025?* by allocating probabilities to the response categories *Yes* and *No*. Participants can also update their probabilistic forecasts over time until the question is resolved, which enables them to account for

changes in the world. Questions are resolved when the true state of the world is known either before or on the resolution date. When a question is resolved, the forecasts of each participant can be verified as it is clear which response category is the true state of the world (see Atanasov et al., 2017; Tetlock et al., 2014; Tetlock & Gardner, 2016 for more details).²

The forecasting tournament featured questions on geopolitical issues from 2011 to 2014. The tournament featured political questions like *Will Italy's Silvio Berlusconi resign, lose re-election/confidence vote, OR otherwise vacate office before 1 January 2012?* Research on nonmarket strategy (Baron, 1995; Dorobantu, Kaul, & Zelner, 2017) suggests that such political events can change the institutional environment with profound implications for organizational performance. For this reason, geopolitical foresight may be considered important and valuable in strategy. For example, organizations that accurately predicted the election of Donald Trump as the 47th president of the United States well in advance may be better prepared for the impact of his policies on their organization or may have networked with relevant political decision makers sooner to shape his policies according to their interests. The forecasts were made by forecasters who participated in the first and second year of the tournament and were assigned to the control group. Focusing on forecasters in the control group eliminated potential bias in the estimates resulting from the inclusion of forecasters who received treatments during the forecasting tournament. It also ensured the plausibility of the assumption that the foresight estimates are stable over time as the interventions in the treatment groups aimed at improving the forecasters' foresight.

These data consist of probabilistic predictions by numerous forecasters over time on a large number and wide range of independent questions about future states of the world varying in their duration (*min* = 3, *1st quartile* = 40, *median* = 84, *mean* = 122, *3rd quartile* = 172, *max* = 520).

² The website <https://www.gjopen.com/> illustrates what forecasting tournaments are.

Furthermore, the data were collected from the same forecasters in two consecutive years. Therefore, these data were ideal to quantify the extent to which foresight is luck over one year based on a regression-toward-the-mean approach.

Data Analysis

Estimating the foresight of forecasters. To estimate the foresight of forecasters in years 1 and 2, I took a Bayesian approach (Kruschke, Aguinis, & Joo, 2012) using the R package *brms* (Bürkner, 2017, 2021) to fit the beta IRT model (Eq. 1–5) to the data from year 1 and the data from year 2. *brms* translates standard R syntax into Stan code, which is a Bayesian programming language that implements Markov chain Monte Carlo (MCMC) sampling via the adaptive Hamiltonian Monte Carlo (HMC; Duane et al., 1987; Neal, 2011) and the No-U-Turn Sampler (NUTS; Hoffman & Gelman, 2014). HMC and NUTS converge faster than other Bayesian algorithms when fitting high-dimensional models with correlated parameters as in this case (Gelman et al., 2013). More generally, Bayesian data analysis provides a powerful approach to model different likelihoods, explicitly incorporate prior knowledge about model parameters, and draw statistical inferences about model parameters that can be interpreted intuitively (Gelman et al., 2013; McElreath, 2020).

I used the weakly informative default priors by *brms* to specify the beta IRT models (Bürkner, 2017, 2021). The default prior for the fixed mean and variance parameters in the normal and beta IRT models is an improper flat prior over the reals. The intercepts of the variance in the normal and the precision in the beta IRT models represent exceptions because the default for these parameters is a Student-*t* prior with mean $\mu = 0$, scale $\sigma = 2.5$, and degrees of freedom $\nu = 3$. This prior is also applied to all standard deviations³ of the random effects except for the standard

³ The variances τ^2 of the random effects in Eq. 3 and 5 are modeled as standard deviations τ in *brms*.

deviation of the foresight parameter, which I fixed to the constant 1 to identify the beta IRT model. The LKJ prior with parameter $\zeta = 1$ (Lewandowski, Kurowicka, & Joe, 2009) is the default prior for the correlations between the random effects in the beta IRT model.⁴

I fitted the models using four chains, each with 8,000 iterations consisting of 4,000 warm-up and 4,000 posterior samples resulting in a total of 16,000 posterior samples. I set the initial values for the sampler to 0, the maximum tree depth to 10, and the target average proposal acceptance probability during Stan's adaptation period to 0.85. These settings ensured high-quality posterior samples. I assessed model convergence by inspecting the trace plots (Gabry, Simpson, Vehtari, Betancourt, & Gelman, 2019; Gelman et al., 2013) and the potential scale reduction factor $\hat{R}$ (Brooks & Gelman, 1998; Gelman et al., 2013; Gelman & Rubin, 1992; Vehtari, Gelman, Simpson, Carpenter, & Bürkner, 2021) of each parameter in the beta IRT model. The trace plots indicated stationarity, good mixing, and the convergence of the four chains for each parameter. The potential scale reduction factor $\hat{R}$ was 1 with bulk and tail effective sample sizes > 100 per chain for each parameter. This suggests that the algorithms converged successfully to the target posterior distributions of the beta IRT model parameters in both years.

The unidimensionality assumption of the beta IRT model is plausible as all questions in the data were related to geopolitics. Furthermore, the assumption of question independence is plausible as each question did not depend on another question for resolution (i.e., knowing the answer on one question was not required to accurately predict another question). As the unidimensionality assumption and assumption of question independence are plausible, the assumption of local independence is plausible. Therefore, forecasters can be compared on the same foresight dimension even if they forecast different questions. A plot of the question discriminations

⁴ The covariances $\tau_{X,Y}$ of the parameters in Eq. 3 are modeled as correlations $\rho_{X,Y}$ in brms.

λ_q suggested substantial variation among them, which lends support to the assumption of question-specific discriminations. While the credible interval of the time effect in the mean parameter model suggested a positive time effect $\hat{\beta}_{\mu 10} = 0.02$ ($EE = 0.00, 95\% CI = [0.02; 0.03]$), the credible interval of the time effect in the variance parameter model did not support the notion a time effect $\hat{\beta}_{\phi 10} = 0.00$ ($EE = 0.00, 95\% CI = [0.00; 0.00]$). This indicates that while the forecasts became more accurate over time, they did not necessarily become more precise over time. Overall, this suggests that the assumptions of the beta IRT model are plausible.

Quantifying the extent to which foresight is luck. To quantify the extent to which foresight is luck using the regression-toward-the-mean approach, I first extracted the foresight estimates from the beta IRT models of years 1 and 2. Second, I standardized the foresight estimates in each year to ensure that their means and variances are identical (means of 0 and variances of 1). Furthermore, I plotted the distributions of the standardized foresight estimates in each year. These plots showed that the foresight parameter estimates were normally distributed. A two-sided Kolmogorov-Smirnov test suggested that the distributions were identical ($D = 0.07, p = .574$). Overall, this suggests that the foresight estimates of year 1 and 2 are identically distributed. Third, I visualized the relationship between the standardized foresight estimates from year 2 on the standardized foresight estimates from year 1 using locally estimated scatterplot smoothing (LOESS) curves with spans of 0.2, 0.4, 0.6, and 0.8. These plots suggested that there may be a potential non-linear relationship between the foresight estimates from years 1 and 2 that may be captured by a reversion-toward-the-mean approach as the foresight estimates also exhibited weak positive dependence: The mean foresight estimates in year 2 of forecasters with above average as well as forecasters with below average foresight in year 1 reverted toward the mean foresight of forecasters (i.e., 0) but not beyond.

Fourth, I implemented the regression-toward-the-mean approach by regressing the standardized foresight estimates from year 2 on the standardized foresight estimates from year 1:

$$y_i = \beta_0 + \beta_1 * x_i + \epsilon_i$$

where y_i is the standardized foresight estimate of forecaster i in year 2, β_0 denotes the intercept representing the average foresight of the forecasters in year 2 for forecasters with average foresight in year 1, β_1 refers to the effect of the standardized foresight estimate of forecaster i in year 1 x_i on the standardized foresight estimate of forecaster i in year 2 y_i , and ϵ_i is the residual of forecaster i . The regression coefficient β_1 describes the extent to which foresight is skill and the relative regression toward the mean $\rho = 1 - \beta_1$ denotes the extent to which foresight is luck assuming that the relationship between the standardized foresight estimates of years 1 and 2 is linear.

Fifth, to capture the potential non-linear relationship, I fitted a second-order polynomial regression to model the relationship between the standardized foresight estimate from year 2 and the standardized foresight estimate from year 1, including both linear and quadratic terms:

$$y_i = \beta_0 + \beta_1 * x_i + \beta_2 * x_i^2 + \epsilon_i$$

where β_1 denotes the effect of the standardized foresight estimate of forecaster i in year 1 x_i on the standardized foresight estimate of forecaster i in year 2 y_i , and β_2 refers to the effect of the squared standardized foresight estimate of forecaster i in year 1 x_i^2 on the standardized foresight estimate of forecaster i in year 2 y_i . I also fitted a third-order polynomial regression to model the relationship between the standardized foresight estimate from year 2 and the standardized foresight estimate from year 1, including linear, quadratic, and cubic terms:

$$y_i = \beta_0 + \beta_1 * x_i + \beta_2 * x_i^2 + \beta_3 * x_i^3 + \epsilon_i$$

where β_3 refers to the effect of the cubed standardized foresight estimate of forecaster i in year 1

x_i^3 on the standardized foresight estimate of forecaster i in year 2 y_i . I selected the regression model for the final interpretation based on the Akaike information criterion (AIC), Bayesian information criterion (BIC), and likelihood-ratio tests comparing the goodness of fit of these regression models. All regressions were fitted using ordinary least squares (OLS). All assumptions of OLS were fulfilled.

Note that the coefficients in the second-order and third-order polynomial regressions cannot easily be interpreted as a measure of skill or luck because the relationship between the standardized foresight estimates is nonlinear. For example, forecasters with above or below average foresight in year 1 may regress more strongly toward the mean in year 2 than forecasters with average foresight in year 1. In other words, the size of the regression toward the mean depends on the foresight of forecasters in year 1. In this case, the extent to which foresight is luck is quantified by the reversion toward the mean κ . For the upper portion of the foresight distribution as defined by cutoff c with $c \geq \mu$, the reversion toward the mean is defined by (Samuels, 1991)

$$\kappa = \frac{E[X_1|X_1 > c] - E[X_2|X_1 > c]}{E[X_1|X_1 > c] - \mu},$$

where X_1 and X_2 are identically distributed random variables (e.g., the foresight of forecasters in years 1 and 2), $X_1 > c$ selects all forecasters who possess foresight in year 1 above cutoff c , and μ is the population mean of X_1 and X_2 . For the lower portion of the foresight distribution as defined by cutoff c with $c \leq \mu$, the reversion toward the mean is defined by (Samuels, 1991)

$$\kappa = \frac{E[X_1|X_1 < c] - E[X_2|X_1 < c]}{E[X_1|X_1 < c] - \mu},$$

where $X_1 < c$ selects all forecasters who possess foresight in year 1 below cutoff c .

The reversion of the upper portion of the foresight distribution toward the mean can be empirically estimated by

$$\hat{\kappa} = \frac{\sum_{x_i > c} x_i - \sum_{x_i > c} y_i}{\sum_{x_i > c} x_i - \sum_{i=1}^n x_i}$$

and analogously for the lower portion of the foresight distribution by

$$\hat{\kappa} = \frac{\sum_{x_i < c} x_i - \sum_{x_i < c} y_i}{\sum_{x_i < c} x_i - \sum_{i=1}^n x_i}$$

where x_i denotes the foresight of forecaster i at time 0, $x_i > c$ selects all forecasters whose foresight in year 1 is above cutoff c , $x_i < c$ selects all forecasters whose foresight in year 1 is below cutoff c , and y_i refers to the foresight of forecaster i in year 2. Statistical significance may be determined by a bootstrapping approach that estimates 95% percentile confidence intervals of the reversion toward the mean $\hat{\kappa}$ based on 10,000 resamples of the data. In the following, I will report the results of these analyses.

Results

Table 2 shows the results of the first-order, second-order, and third-order polynomial regression models and their comparison. The regression coefficient showed a strong positive relationship between the standardized foresight estimates from year 1 and the standardized foresight estimates from year 2 ($\hat{\beta}_1 = 0.68, SE = 0.05, t = 14.85, p < .001$), which implied that foresight regressed by $\hat{\rho} = 1 - 0.68 = 0.32$ toward the mean foresight of forecasters. This suggests that foresight is 68% skill and 32% luck. This means that the foresight of an individual forecaster regresses on average by 32% toward the mean foresight of all forecasters over one year. Figure 1 shows the regression of the foresight estimates toward the mean foresight of all forecasters from year 1 to year 2 based on the first-order linear regression model. Note that this result may somewhat underestimate the extent of skill and overestimate the extent of luck in foresight because any measurement error in relation to the foresight of forecasters attenuates the regression coefficient (Fuller & Hidirolou, 1978).

 Insert Table 2 about here

 Insert Figure 1 about here

Regarding the model fit, the results show that the first-order linear regression model had the smallest *BIC* (606.40), whereas the third-order linear regression model had the smallest *AIC* (591.85). The likelihood-ratio tests indicate that the second-order linear regression model did not fit the data better than the first-order linear regression model ($\chi^2(1) = 0.13, p = .608$). However, the third-order linear regression model fit the data better than the second-order (and therefore logically also the first-order) linear regression model ($\chi^2(1) = 4.03, p < .01$). Overall, these results suggest that the third-order linear regression model fit the data best, which implied a non-linear relationship between the standardized foresight estimates. This means that the role of luck in foresight seemed to depend on the foresight of forecasters in year 1.

Figure 2 shows the regression of the foresight estimates toward the mean foresight of all forecasters from year 1 to year 2 based on the third-order linear regression model. The third-order linear regression model suggests that the regression toward the mean followed an S-shaped curve, where luck appeared to play a stronger role in high and low foresight but a weaker role in average foresight. To quantify the extent to which foresight is luck depending on the foresight of forecasters, I computed the reversion toward the mean for forecasters with foresight above and below multiple cutoffs. The results displayed in Table 3 suggest that the reversion toward the mean was stronger for forecasters with extreme foresight ($x_i > 3$; 38%, 95% *CI* = [0.11, 0.70]) and exceptionally good foresight ($x_i > 2$; 35%, 95% *CI* = [0.20, 0.73]). Furthermore, the reversion toward the mean was stronger for forecasters with below average foresight ($x_i < -1$; 47%, 95% *CI* = [0.26, 0.73]) and exceptionally poor foresight ($x_i < -2$; 54%, 95% *CI* = [0.31, 0.82]). In

contrast, the reversion toward the mean was weaker for the average forecaster in the upper ($x_i > 0$; 23%, 95% $CI = [0.06, 0.47]$) and lower portion of the foresight distribution ($x_i < 0$; 23%, 95% $CI = [0.07, 0.41]$).⁵ Figure 3 illustrates the reversion of the foresight estimates toward the mean foresight of all forecasters from year 1 to year 2 for different portions of the foresight distribution as defined by cutoff c .

Insert Figure 2 about here

Insert Table 3 about here

Insert Figure 3 about here

Furthermore, Figures 2 and 3 suggest that foresight followed a weak pattern of luck: Forecasters who accurately (inaccurately) predicted future states of the world in year 1 predicted future states of the world less accurately (more accurately) in year 2. However, these forecasters did not predict future states of the world worse (better) than forecasters with above (below) average foresight in year 1. This becomes evident by comparing the mean foresight of forecasters in year 2 who showed extreme foresight ($x_i > 3$) to the mean foresight of exceptionally good forecasters ($2 < x_i < 3$), where the mean difference δ of their foresight is nonsignificant based on 95% percentile confidence intervals of 10,000 bootstraps samples ($\hat{\delta} = 0.53$, 95% $CI = [-0.63, 1.81]$). This is corroborated by comparing the mean foresight of forecasters in year 2 who

⁵ Note that the reversion toward the mean approach only allows for the calculation of effects in upper and lower portions of the foresight distribution as defined by one cutoff. Therefore, it is currently not possible to estimate the reversion toward the mean of average forecasters (i.e., $-1 < x_i < 1$) as this would require two cutoffs for which the reversion toward the mean is not defined (Samuels, 1991). For this reason, estimating the average reversion toward the mean for forecasters in the upper and lower portion of the foresight distribution represents the best estimate of the reversion toward the mean of the average forecaster, albeit the true reversion toward the mean is likely smaller given that forecasters with extreme and exceptionally good foresight exhibit a disproportionately strong reversion toward the mean foresight of forecasters.

showed exceptionally poor foresight ($x_i < -2$) to the mean foresight of below average forecasters ($-2 < x_i < -1$), where forecasters with exceptionally poor foresight exhibited worse foresight in year 2 than forecasters with below average foresight in year 2 ($\hat{\delta} = -0.52, 95\% CI = [-0.98, -0.07]$). Overall, these results suggest that even though luck plays a stronger role in exceptionally good (bad) foresight, such foresight may still signal good (bad) judgment. More generally, these results suggest that even though foresight is to some extent luck, foresight appears to be mainly skill.

DISCUSSION

Motivated by the goal of assessing the relative plausibility of cognitive, evolutionary, and chance theories of foresight, this paper quantified the extent to which foresight is luck using a regression-toward-the-mean approach. The results suggest that foresight is 32% luck, albeit luck plays a greater role among top forecasters (35%) and forecasters with below average foresight (47%) but a lesser role among average forecasters (23%). The analyses also uncovered that foresight follows a weak pattern of luck. In the following, I will discuss the theoretical and practical implications, describe the limitations of this study, and outline opportunities for future research.

Theoretical Implications

The finding that foresight is 32% luck has important implications for cognitive, evolutionary, and chance explanations of foresight: Cognitive theories suggest that strategists possess “boundedly rational foresight” (Gavetti et al., 2012: 9) that is grounded in their mental representations of the world rendering foresight mainly a skill (Csaszar, 2018; Csaszar & Laureiro-Martínez, 2018; Gavetti & Levinthal, 2000). This explanation is consistent with the finding that foresight appears to be mainly a skill but also somewhat luck highlighting that foresight is boundedly rational. It is also consistent with a growing literature highlighting the importance of

mental representations and cognitive processes for foresight (Csaszar & Laureiro-Martínez, 2018; Csaszar & Ostler, 2020; Davis, Eisenhardt, & Bingham, 2009; Davis et al., 2009; Eisenhardt & Sull, 2001; Gavetti & Levinthal, 2000; Kapoor & Wilde, 2023; Lee & Csaszar, 2020; Satopää, Salikhov, Tetlock, & Mellers, 2021; Sull & Eisenhardt, 2015). Although skill may also refer to any other individual factor than cognition like personality characteristics and motivation, research suggests that such factors only play a minor role in the explanation of foresight (Schuler et al., 2024). Thus, the result that foresight is 32% luck may be viewed as evidence for cognitive explanations of foresight.

Evolutionary theories acknowledge that strategists may exhibit myopic foresight grounded in their cognitive preadaptation that enables them to discover opportunities serendipitously (Denrell et al., 2003; Gavetti & Lecuona Torras, 2021; Gavetti & Menon, 2016). This explanation is partly inconsistent with the finding that foresight appears to be mainly a skill as foresight deriving from the serendipitous discovery of opportunities in the environment should be largely luck (Busch, 2020, 2024; Denrell et al., 2003). Although forecasters might have been cognitively preadapted to geopolitical questions, the number, breadth (i.e., diversity of knowledge domains), and time horizon of the geopolitical questions covered in the forecasting tournament make this explanation less plausible compared to cognitive explanations that allow for strategists to anticipate a wide variety of distant future states of the world by learning and adapting mental representations. That being said, cognitive preadaptation and serendipity likely still play a role in foresight to some extent, where cognitive preadaptation contributes to skill.

Chance theories question the existence of managerial foresight by emphasizing the role of random processes in accurately predicting future states of the world (Denrell & Fang, 2010; Denrell et al., 2015; Liu, 2020; Liu & de Rond, 2016). However, this explanation is inconsistent

with the finding that foresight is on average only 32% luck highlighting the role of cognition and cognitive preadaptation in foresight. Furthermore, foresight followed a weak pattern of luck (Denrell et al., 2019; Liu, 2020): The foresight of forecasters who accurately (inaccurately) predicted future states of the world in year 1 predicted future states of the world less accurately (more accurately) in year 2, and this pattern was stronger for forecasters with exceptionally good and bad foresight. Yet, these forecasters still exhibited on average equal (worse) foresight compared to forecasters with moderately accurate (inaccurate) foresight. Thus, this study refines prior research suggesting that exceptional foresight may be a signal of poor judgment (Denrell & Fang, 2010; Denrell & Liu, 2012) by highlighting that extremely good (bad) foresight may also signal good (bad) judgment. Although chance explanations may be less plausible compared to cognitive and evolutionary explanations, they still legitimately emphasize that foresight involves luck. Taken together, even though cognitive theories appear to be more consistent with the finding that foresight is 32% luck than alternative explanations, these results highlight that cognitive, evolutionary, and chance explanations of foresight may be best viewed as complements to each other as each theory emphasizes an important factor in the strategist's foresight.

Practical Implications

The findings in relation to cognitive, evolutionary, and chance explanations of foresight also have important practical implications: Cognitive theories suggest that strategists who displayed exceptionally good foresight may be entrusted with the deliberate design of strategies (Mintzberg, 1978; Mintzberg & Waters, 1985) based on their mental representations of the world to increase organizational performance. For example, organizations may select strategists based on their foresight in forecasting tournaments. Strategists may also seek to learn better representations of the world that can guide their strategic decisions (Csaszar et al., 2024; Csaszar

& Ostler, 2020; Heshmati & Csaszar, 2024). Furthermore, strategists may deepen and broaden their knowledge to serendipitously exploit an increasing range and number of opportunities they see that others cannot see when the environment changes (Busch, 2020, 2024; Denrell et al., 2003). At the same time, strategists should rely on evolutionary principles (Levinthal, 2021; Murmann, 2003; Nelson & Winter, 1982) to guide emergent strategies (Mintzberg, 1978; Mintzberg & Waters, 1985). For instance, strategists may guide the organization by communicating a strategic vision and devise internal selection criteria for projects in the organization (Levinthal, 2017, 2021) that guide the generation and promotion of strategic initiatives (Burgelman, 1983a, 1983b). Furthermore, strategists may design organizational structures to support these activities (Csaszar, 2013). More generally, strategists should remain humble as foresight is always to some extent luck (Denrell et al., 2015; Liu, 2020).

Limitations and Future Research Opportunities

This study is limited in two respects, albeit these limitations represent valuable opportunities for future research. First, this study uses data from a geopolitical forecasting tournament to evaluate the relative merit of cognitive, evolutionary, and chance explanations of foresight. A large and diverse number of questions about the future in this domain provides a good way of assessing the relative plausibility of cognitive and evolutionary theories of foresight. As the number and diversity of the questions increase, cognitive theories become more plausible compared to evolutionary theories as it becomes more difficult to display foresight in a large number of diverse questions simply by cognitive preadaptation. However, future research may subject the relative plausibility of cognitive and evolutionary theories to an even stronger test by quantifying the extent to which foresight is luck using data from multiple domains. Exhibiting foresight across a wide range of knowledge domains (e.g., technology, politics, economics) would

render it even harder to display foresight simply by cognitive preadaptation. Studying the extent of foresight in other domains could also provide insights into how the extent of luck in foresight varies across different domains.

Second, this study used data from two years of a geopolitical forecasting tournament. Although this provides important insights about the regression of foresight toward the mean over the time of up to one year, it remains unclear to what extent foresight regresses toward the mean over longer time horizons (e.g., three to five years). Examining the role of luck in foresight over such long time horizons would require exceptional commitment by forecasters as they would have to participate in two consecutive forecasting tournaments consisting of questions lasting for three to five years. However, examining the role of luck in such distant foresight may provide valuable additional insights for the relative plausibility of cognitive, evolutionary, and chance theories of foresight. Therefore, future research should examine the regression of foresight toward the mean over longer time horizons and evaluate the implications for different explanations of foresight as the relative plausibility of these theories may change with the time horizon of foresight. For example, luck may play a greater role in long-term foresight rendering evolutionary and chance theories more plausible than cognitive theories for extremely distant foresight.

To conclude, cognitive, evolutionary, and chance explanations of foresight each provide unique insights into why strategists may or may not display foresight. While foresight is partly luck, the findings of this paper add to the mounting evidence that foresight is a skill rooted in accurate mental representations of the world. As foresight can be a source of competitive advantage for organizations, understanding why some managers have foresight and how managers may improve their foresight holds great promise for strategy.

REFERENCES

- Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., et al. 2017. Distilling the wisdom of crowds: Prediction markets vs. Prediction polls. *Management Science*, 63(3): 691–706.
- Baron, D. P. 1995. Integrated strategy: Market and nonmarket components. *California Management Review*, 37(2): 47–65.
- Bo, Y. E., Budescu, D. V., Lewis, C., Tetlock, P. E., & Mellers, B. A. 2017. An IRT forecasting model: Linking proper scoring rules to item response theory. *Judgment and Decision Making*, 12(2): 90–103.
- Brooks, S. P., & Gelman, A. 1998. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4): 434–455.
- Burgelman, R. A. 1983a. A process model of internal corporate venturing in the diversified major firm. *Administrative Science Quarterly*, 28(2): 223.
- Burgelman, R. A. 1983b. Corporate entrepreneurship and strategic management: Insights from a process study. *Management Science*, 29(12): 1349–1364.
- Bürkner, P.-C. 2017. brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1). <https://doi.org/10.18637/jss.v080.i01>.
- Bürkner, P.-C. 2021. Bayesian item response modeling in R with brms and Stan. *Journal of Statistical Software*, 100: 1–54.
- Busch, C. 2020. *The serendipity mindset: The art and science of creating good luck*. Penguin.
- Busch, C. 2024. Towards a theory of serendipity: A systematic review and conceptualization. *Journal of Management Studies*, 61(3): 1110–1151.
- Cai, L., Choi, K., Hansen, M., & Harrell, L. 2016. Item response theory. *Annual Review of Statistics and Its Application*, 3(1): 297–321.
- Carroll, R. J., Primo, D. M., & Richter, B. K. 2016. Using item response theory to improve measurement in strategic management research: An application to corporate social responsibility. *Strategic Management Journal*, 37(1): 66–85.
- Csaszar, F. A. 2013. An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organization Science*, 24(4): 1083–1101.
- Csaszar, F. A. 2018. What makes a decision strategic? Strategic representations. *Strategy Science*, 3(4): 606–619.
- Csaszar, F. A., Hinrichs, N., & Heshmati, M. 2024. External representations in strategic decision-making: Understanding strategy's reliance on visuals. *Strategic Management Journal*, 45(11): 2191–2226.
- Csaszar, F. A., & Laureiro-Martínez, D. 2018. Individual and organizational antecedents of strategic foresight: A representational approach. *Strategy Science*, 3(3): 513–532.
- Csaszar, F. A., & Levinthal, D. A. 2016. Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37(10): 2031–2049.
- Csaszar, F. A., & Ostler, J. 2020. A contingency theory of representational complexity in organizations. *Organization Science*, 31(5): 1198–1219.
- Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. 2009. Optimal structure, market dynamism, and the strategy of simple rules. *Administrative Science Quarterly*, 54(3): 413–452.
- Denrell, J. 2004. Random walks and sustained competitive advantage. *Management Science*, 50(7): 922–934.
- Denrell, J., & Fang, C. 2010. Predicting the next big thing: Success as a signal of poor judgment. *Management Science*, 56(10): 1653–1667.

- Denrell, J., Fang, C., & Liu, C. 2015. Perspective—Chance explanations in the management sciences. *Organization Science*, 26(3): 923–940.
- Denrell, J., Fang, C., & Liu, C. 2019. In search of behavioral opportunities from misattributions of luck. *Academy of Management Review*, 44(4): 896–915.
- Denrell, J., Fang, C., & Winter, S. G. 2003. The economics of strategic opportunity. *Strategic Management Journal*, 24(10): 977–990.
- Denrell, J., Fang, C., & Zhao, Z. 2013. Inferring superior capabilities from sustained superior performance: A Bayesian analysis. *Strategic Management Journal*, 34(2): 182–196.
- Denrell, J., & Liu, C. 2012. Top performers are not the most impressive when extreme performance indicates unreliability. *Proceedings of the National Academy of Sciences*, 109(24): 9331–9336.
- Dorobantu, S., Kaul, A., & Zelner, B. 2017. Nonmarket strategy research through the lens of new institutional economics: An integrative review and future directions. *Strategic Management Journal*, 38(1): 114–140.
- Duane, S., Kennedy, A. D., Pendleton, B. J., & Roweth, D. 1987. Hybrid monte Carlo. *Physics Letters B*, 195(2): 216–222.
- Eisenhardt, K. M., & Sull, D. N. 2001. Strategy as simple rules. *Harvard Business Review*, 79(1): 106–16.
- Foster, G. C., Min, H., & Zickar, M. J. 2017. Review of item response theory practices in organizational research: Lessons learned and paths forward. *Organizational Research Methods*, 20(3): 465–486.
- Fuller, W. A., & Hidirolou, M. A. 1978. Regression estimation after correcting for attenuation. *Journal of the American Statistical Association*, 73(361): 99–104.
- Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. 2019. Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182(2): 389–402.
- Gavetti, G. 2012. Toward a behavioral theory of strategy. *Organization Science*, 23(1): 267–285.
- Gavetti, G., Greve, H. R., Levinthal, D. A., & Ocasio, W. 2012. The behavioral theory of the firm: Assessment and prospects. *Academy of Management Annals*, 6: 1–40.
- Gavetti, G., & Lecuona Torras, J. R. 2021. A Neo-Carnegie approach to the agency question: Bridging the evolutionary and cognitive views of strategy. *Strategy Science*, 6(4): 353–359.
- Gavetti, G., & Levinthal, D. A. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1): 113–137.
- Gavetti, G., Levinthal, D. A., & Rivkin, J. W. 2005. Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal*, 26(8): 691–712.
- Gavetti, G., & Menon, A. 2016. Evolution cum agency: Toward a model of strategic foresight. *Strategy Science*, 1(3): 207–233.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., et al. 2013. *Bayesian data analysis*. Chapman and Hall/CRC.
- Gelman, A., & Rubin, D. B. 1992. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4): 457–472.
- Hambleton, R. K., & Swaminathan, H. 2013. *Item response theory: Principles and applications*. Springer Science & Business Media.
- Hambleton, R. K., Swaminathan, H., & Rogers, H. J. 1991. *Fundamentals of item response*

theory. SAGE.

- Heshmati, M., & Csaszar, F. A. 2024. Learning strategic representations: Exploring the effects of taking a strategy course. *Organization Science*, 35(2): 453–473.
- Hoffman, M. D., & Gelman, A. 2014. The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1): 1593–1623.
- Hox, J., Moerbeek, M., & Schoot, R. van de. 2017. *Multilevel analysis: Techniques and applications* (3rd ed.). New York: Routledge. <https://doi.org/10.4324/9781315650982>.
- Kapoor, R., & Wilde, D. 2023. Peering into a crystal ball: Forecasting behavior and industry foresight. *Strategic Management Journal*, 44(3): 704–736.
- Kruschke, J. K., Aguinis, H., & Joo, H. 2012. The time has come: Bayesian methods for data analysis in the organizational sciences. *Organizational Research Methods*, 15(4): 722–752.
- Lee, S., & Csaszar, F. A. 2020. Cognitive and structural antecedents of innovation: A large-sample study. *Strategy Science*, 5(2): 71–97.
- Levinthal, D. A. 2017. Mendel in the C-suite: Design and the evolution of strategies. *Strategy Science*, 2(4): 282–287.
- Levinthal, D. A. 2021. *Evolutionary processes and organizational adaptation: A mendelian perspective on strategic management*. Oxford University Press.
- Lewandowski, D., Kurowicka, D., & Joe, H. 2009. Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9): 1989–2001.
- Liu, C. 2020. *Luck: A key idea for business and society*. Routledge.
- Liu, C., & de Rond, M. 2016. Good night, and good luck: Perspectives on luck in management scholarship. *Academy of Management Annals*, 10(1): 409–451.
- March, J. G., & Simon, H. A. 1958. *Organizations*. New York: Wiley.
- McElreath, R. 2020. *Statistical rethinking: A Bayesian course with examples in R and Stan*. CRC Press.
- McGahan, A. M., & Porter, M. E. 1997. How much does industry matter, really? *Strategic Management Journal*, 18(S1): 15–30.
- Merkle, E. C., Steyvers, M., Mellers, B. A., & Tetlock, P. E. 2016. Item response models of probability judgments: Application to a geopolitical forecasting tournament. *Decision*, 3(1): 1–19.
- Merkle, E. C., Steyvers, M., Mellers, B. A., & Tetlock, P. E. 2017. A neglected dimension of good forecasting judgment: The questions we choose also matter. *International Journal of Forecasting*, 33(4): 817–832.
- Mintzberg, H. 1978. Patterns in strategy formation. *Management Science*, 24(9): 934–948.
- Mintzberg, H., & Waters, J. 1985. Of strategies, deliberate and emergent. *Strategic Management Journal*, 6(3): 257–272.
- Murmann, J. P. 2003. *Knowledge and competitive advantage: The coevolution of firms, technology, and national institutions*. Cambridge University Press.
- Neal, R. M. 2011. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC.
- Nelson, R. R., & Winter, S. G. 1982. *An evolutionary theory of economic change*. Cambridge, Mass.: The Belknap Press of Harvard University Press.
- Packard, M. D., & Clark, B. B. 2020. Mitigating versus managing epistemic and aleatory

- uncertainty. *Academy of Management Review*, 45(4): 872–876.
- Porter, M. E. 1979. How competitive forces shape strategy. *Harvard Business Review*, 57(2): 2–10.
- Porter, M. E. 2008. The five competitive forces that shape strategy. *Harvard Business Review*, 86(1): 78–93.
- Raudenbush, S. W., & Bryk, A. S. 2002. *Hierarchical linear models: Applications and data analysis methods*, vol. 1. Sage.
- Rumelt, R. P. 1991. How much does industry matter? *Strategic Management Journal*, 12(3): 167–185.
- Samuels, M. L. 1991. Statistical reversion toward the mean: More universal than regression toward the mean. *The American Statistician*, 45(4): 344–346.
- Satopää, V. A., Salikhov, M., Tetlock, P. E., & Mellers, B. A. 2021. Bias, information, noise: The BIN model of forecasting. *Management Science*, 67(12): 7599–7618.
- Schmittlein, D. C. 1989. Surprising inferences from unsurprising observations: Do conditional expectations really regress to the mean? *The American Statistician*, 43(3): 176–183.
- Schuler, B. A., Murmann, J. P., Beisemann, M., & Satopää, V. 2024. *Individual foresight: Concept, operationalization, and correlates*. Accepted for publication at the *International Journal of Forecasting*.
<https://www.alexandria.unisg.ch/handle/20.500.14171/121555>.
- Simon, H. A. 1947. *Administrative behavior*. Oxford, England: Macmillan.
- Sull, D. N., & Eisenhardt, K. M. 2015. *Simple rules: How to thrive in a complex world*. Boston: Houghton Mifflin Harcourt.
- Tetlock, P. E., & Gardner, D. 2016. *Superforecasting: The art and science of prediction*. New York, NY: Random House.
- Tetlock, P. E., Mellers, B. A., Rohrbaugh, N., & Chen, E. 2014. Forecasting tournaments: Tools for increasing transparency and improving the quality of debate. *Current Directions in Psychological Science*, 23(4): 290–295.
- Van der Linden, W. J., & Hambleton, R. 1997. *Handbook of item response theory*. Springer.
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. 2021. Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2): 667–718.
- Wang, M. Z. 2023. Changes in industry and corporate effects in the United States, 1978–2019. *Strategic Management Journal*, 44(2): 477–490.

TABLE 1**Summary of Cognitive, Evolutionary, and Chance Theories of Foresight**

	Cognitive Theories	Evolutionary Theories	Chance Theories
Nature of Foresight	Boundedly rational foresight of near, intermediate, and distant future states of the world	Myopic foresight of unique opportunities that can be serendipitously discovered	Foresight may not exist, accurate predictions of future states of the world are due to chance
Explanation for Foresight	Accurate mental representations of the world that enable strategists to predict future states of the world	Cognitive preadaptation (deep domain knowledge) to a future environment that enables strategists to see unique opportunities that other strategists cannot see once the future environment realizes	Random processes in the world determine future states of the world that are accurately or inaccurately predicted by strategists due to chance
Extent of Luck in Foresight	Somewhat luck but mainly skill as strategists can learn and adapt mental representations, albeit their bounded rationality and the inherent randomness of the world preclude perfect foresight	Mainly luck but also somewhat skill as strategists may deepen and broaden their domain knowledge, attend to changes in the environment, and exert effort to seize opportunities in the environment, albeit the strategist mostly cannot control the environment determining what deep domain knowledge becomes valuable	Mostly luck as strategists cannot control random processes and consequently, not predict future states of the world accurately over time

TABLE 2

Results of the First-Order, Second-Order, and Third-Order Polynomial Linear Regression Models and Their Comparison

	Model 1: Linear Effect				Model 2: Quadratic Effect				Model 3: Cubic Effect			
	Model Estimates											
Predictor	$\hat{\beta}$	SE	t	p	$\hat{\beta}$	SE	t	p	$\hat{\beta}$	SE	t	p
Intercept	0.00	0.05	0.00	1.000	-0.01	0.05	-0.24	.814	-0.03	0.05	-0.65	.517
Foresight Year 1	0.68	0.05	14.85	***	0.67	0.05	13.82	***	0.79	0.07	12.13	***
Foresight Year 1 Squared					0.01	0.02	0.51	.612	0.05	0.03	1.90	†
Foresight Year 1 Cubed									-0.03	0.01	-2.75	**
Statistic	Model Fit											
R ²	0.46				0.46				0.47			
F	F(1, 263) = 220.50***				F(2, 262) = 110.10***				F(3, 261) = 77.74***			
AIC	595.66				597.40				591.85			
BIC	606.40				611.72				609.75			
LRT					$\chi^2(1) = 0.13, p = .608$				$\chi^2(1) = 4.03^{**}$			

Note. Each regression model predicted the standardized foresight estimates from year 2 from the standardized foresight estimates from year 1 based on $n = 265$ observations. $\hat{\beta}$ = estimated regression coefficient; *SE* = standard error; *t* = *t*-statistic; *p* = *p*-value; R^2 = coefficient of determination quantifying the variation in the dependent variable that is explained by the independent variable(s); *F* = *F*-test indicating if the regression explains significant variation in the dependent variable; *AIC* = Akaike information criterion; *BIC* = Bayesian information criterion; *LRT* = likelihood-ratio test compares the goodness of fit of two nested models by testing the null hypothesis that the two models fit the data equally well. A significant likelihood-ratio test suggests that the simpler model describes the data significantly worse and therefore, the more complex model should be selected for interpretation. In this case, the likelihood-ratio tests compared Model 2 to Model 1 and Model 3 to Model 2. Robustness checks using robust standard errors and removing one outlier with high leverage in the data did not affect the interpretation of the results.

† $p < .10$

** $p < .01$

*** $p < .001$


TABLE 3
Reversion Toward the Mean Foresight of Forecasters for Different Portions of the
Foresight Distribution

Portion	n	$\hat{\kappa}$	LL	UL	Sig
$x_i > 3$	4	0.38	0.11	0.70	*
$x_i > 2$	8	0.35	0.20	0.73	*
$x_i > 1$	33	0.30	0.10	0.48	*
$x_i > 0$	121	0.23	0.06	0.47	*
$x_i < 0$	144	0.23	0.07	0.41	*
$x_i < -1$	32	0.47	0.26	0.73	*
$x_i < -2$	2	0.54	0.31	0.82	*
$x_i < -3$	1	0.56			

Note. Portion = portion of the foresight distribution of the standardized foresight estimates in year 1 x_i for which the reversion toward the mean was calculated; n = number of forecasters in the portion of the distribution; $\hat{\kappa}$ = reversion toward the mean; LL = lower level of the 95% percentile bootstrap confidence interval based on 10,000 resamples; UL = upper level of the 95% percentile bootstrap confidence interval based on 10,000 resamples; Sig = statistical significance based on the 95% percentile bootstrap confidence interval, where * denotes statistical significance at $\alpha = .05$ if the 95% confidence interval does not include 0. The reversion toward the mean quantifies the extent to which foresight is luck in the corresponding portion of the foresight distribution.

FIGURE 1


**Regression of the Foresight Estimates Toward the Mean Foresight of all Forecasters From
Year 1 to Year 2 Based on the First-Order Linear Regression Model**



Note. The foresight estimates from the beta IRT model in year 1 and year 2 are standardized. The horizontal grey line indicates the mean foresight of the forecasters μ . The grey 45 degree line ($y = x$) shows the expected relationship between the foresight estimates of year 1 and 2 if foresight is 100% skill. The red line is the regression line from the first-order linear regression model including a grey 95% confidence band. As the regression line lies between the horizontal and 45 degree line, the foresight estimates regress toward the mean foresight of all forecasters from year 1 to year 2.

FIGURE 2


Regression of the Foresight Estimates Toward the Mean Foresight of all Forecasters From Year 1 to Year 2 Based on the Third-Order Linear Regression Model.



Note. The foresight estimates from the beta IRT model in year 1 and year 2 are standardized. The horizontal grey line indicates the mean foresight of the forecasters μ . The grey 45 degree line ($y = x$) shows the expected relationship between the foresight estimates of year 1 and 2 if foresight is 100% skill. The red curve is the regression curve from the third-order linear regression model including a grey 95% confidence band. As the regression line lies between the horizontal and 45 degree line, the foresight estimates regress toward the mean foresight of all forecasters from year 1 to year 2, albeit extreme foresight estimates in year 1 appear to regress more strongly toward the mean.

FIGURE 3

Reversion of the Foresight Estimates Toward the Mean Foresight of All Forecasters From Year 1 to Year 2 for Different Portions of the Foresight Distribution



Note. The horizontal grey line indicates the mean foresight of the forecasters. The vertical grey line represents the cutoff c splitting the distribution of the foresight estimates in year 1 into an upper and lower portion. The red square denotes the mean foresight of the forecasters in year 1 and the red triangle indicates the mean foresight of the forecasters in year 2 in the corresponding portion of the foresight distribution in year 1. The difference between the square and triangle illustrates the absolute reversion toward the mean.