



Intel[®] Technology Journal

The Spectrum of Risk Management in a Technology Company

Using Forecasting Markets to Manage Demand Risk

Using Forecasting Markets to Manage Demand Risk

Jay W. Hopman, Information Technology Innovation & Research, Intel Corporation

Index words: prediction markets, forecasting, planning, performance incentives

ABSTRACT

Intel completed a study of several generations of products to learn how product forecasts and plans are managed, how demand risks manifest themselves, and how business processes contend with, and sometimes contribute to, demand risk. The study identified one critical area prone to breakdown: the aggregation of market insight from customers. Information collected from customers and then rolled up through sales, marketing, and business planning teams is often biased, and it can lead to inaccurate forecasts, as evidenced by historical results.

A research effort launched in 2005 sought to introduce new methodologies that might help crack the bias in demand signals. We worked with our academic partners to develop a new application, a form of prediction market, integrated with Intel's regular short-term forecasting processes. The process enables product and market experts to dynamically negotiate product forecasts in an environment offering anonymity and performance-based incentives. To the extent these conditions curb bias and motivate improved performance, the system should alleviate demand miscalls that have resulted in inventory surpluses or shortages in the past. Results of early experiments suggest that market-developed forecasts are meeting or beating traditional forecasts in terms of increased accuracy and decreased volatility, while responding well to demand shifts. In addition, the new process is training Intel's experts to improve their use and interpretation of information.

INTRODUCTION

Demand risk is implicit to manufacturing businesses, but for high-tech firms it poses a particularly strong threat. As product lifecycles shrink and new generations of technology enter the market more quickly, achieving strong top- and bottom-line results hinges on estimating overall demand and product mix as accurately as possible. Products with manufacturing lead times of months or even quarters are all the more critical to forecast correctly because last-minute inventory adjustments are limited or sometimes just not feasible. Our study of multiple

generations of product transitions discovered that producing high-quality demand forecasts is difficult to achieve consistently and that mistakes can be quite costly [1].

Managing demand risk is critical to Intel's success, but it is only one of many business challenges the company faces. Across the organization, teams grapple with questions such as how many units of products x, y, and z customers will demand at certain prices, how much factory capacity should be funded, which products should be brought to market, which features and technologies should be included in new products, and when new products will be ready for production and distribution. Interviews with employees trying to answer these questions reveal a common issue: belief that they do not have the best available information and insight to guide business decisions.

Tackling demand risk and other challenges requires moving information around decentralized organizations in new ways. If employees across Intel's many functional groups have information and insights that can help inform our planning and forecasting decisions, we need a way to aggregate that information and turn it into intelligence. Prediction markets are a potential solution to this problem and have been written about extensively for the past five to ten years. Our research discovered that, despite the buzz around prediction markets, the integration of prediction markets and similar Information Aggregation Mechanisms (IAMs) into organizational forecasting processes is still in its infancy. Popular stories on prediction markets still frame the potential as being greater than the demonstrated value, and reports of usage at companies such as Hewlett Packard, Microsoft, Google, Eli Lilly, and others suggest that application is often viewed as experimental and that markets are largely separate from other organizational forecasting processes [2, 3].

While our research of prediction markets is growing to explore more business problems over time, the area we first tackled at Intel is demand forecasting. We have developed and piloted an IAM that is integrated into our

regular forecasting processes and, through this development, have considered many questions and ideas about designing markets real companies can use to address real problems. It will take extensive research and experimentation to answer these design questions, but we are encouraged that even trial solutions based on the experience of other researchers, feedback from our business partners, and our own intuition are producing good results.

CHALLENGES TO ANTICIPATING MARKET DEMAND

Since 2001, we have been studying the release of current and historical products. We have tracked the evolution of forecasts, factory production, and inventory for many major product releases and studied how the signals flow through teams across Intel's organization. Our methods have included both quantitative analysis of our data sets and interviews with personnel in groups that work with these data sets to understand policies, strategies, and perspectives on the product transitions.

We learned that calling demand correctly for new products—and the products the new products replace—is a formidable task. Four fundamental sources of noise cause difficulty in determining true market demand: *current data*, such as orders and inventory; *market assessment*, such as intelligence and consensus on how appealing products and promotions (and competing products) might be to the market; *market objectives*, the goals Intel has for its products, such as unit sales, average price, market segment share, and technology leadership; and, *strategic plans*, such as the decisions about which products and stock keeping units (SKUs) to sell, how to price them, and how to take advantage of technological and manufacturing capabilities. Nearly all the pitfalls we have discovered in forecasting demand can be linked—with the benefit of hindsight—to one or to a combination of these factors. The question, of course, is how to account for these factors in advance to systematically and repeatedly do the best possible job of forecasting and planning.

The fundamental problem in managing forecasts is twofold. First, the hard data being created, judged, and passed from group to group lack credibility based on past performance, so each group feels the need to adjust the information based on any number of experiences and heuristics. Groups preparing to publish data are aware of how other groups will likely judge the data and are therefore prone to gaming the system, i.e., adjusting numbers in anticipation of future judgment.

Second, data sets themselves do not really convey any specific meaning. Meaning can be inferred from how the data compare to expectations or previously published

data, but numbers in enterprise applications or spreadsheets cannot explain the strategies Intel and its customers are employing or the uncertainties they are facing. Decentralized organizations must find a means of transmitting business context; in other words, instead of transmitting mere data sets, they must transmit information and intelligence from employees who have it to employees who need it to make decisions and plans. We learned that Intel has many informal networks that attempt to move that knowledge across the organization, but these networks have many failure modes: turnover of employees in key positions, limited bandwidth of each individual and team, and difficulty systematically discovering the important information to be learned (stated differently, whom to include in the network).

Our research has led to three methods that are being used at Intel today. One focuses on market assessment and uses data from across the organization to score factors affecting ramp rates (Product Transition Index). The second ties market objectives, strategic plans, and market assessment, identifying risks and developing contingency strategies to improve coordination and cooperation (Transition Playbook) [4]. And, the third (IAM) paradoxically uses the most structured of the three methods to promote transmission of the most unstructured information, i.e., any and all information participants feel is relevant to developing a forecast.

The source of demand uncertainty for Intel begins with biased signals from the interaction with our customers. Customers typically signal strong demand for popular upcoming products. In fact, if Intel fulfilled all bookings (advance orders) for all customers, the result would most often be substantial oversupply. Customers want to assure supply and be certain that a competitor does not procure an unfair share, so the condition of “phantom demand” develops. Orders are inflated to keep the playing field level across customers and so that each can lock in as much supply as possible in the event of a shortage. Conversely, orders for new products are sometimes deflated, signaling that customers do not want to go to the new product too quickly. Perhaps they prefer the prior product for any number of reasons, or they believe low demand might lead to price reduction, or in some cases, new technologies and supporting components are relatively scarce and will increase in supply (decrease in price) over time. Whatever the exact cause, a study of orders and forecasts developed by Intel's geographical sales organizations shows that the volatility of these signals is large, and it is not unusual for the forecasts to be over 20% high or low.

Once geographical (“geo”) forecasts are published, a central business planning group is responsible for publishing official demand forecasts that guide the supply

network. The geo forecasts are one input considered by this team, but many other factors including models of worldwide sales growth, Intel's share of the market segment, product mix by any number of attributes, sales versus price point, historical product ramp rates, and various inventory data (for instance, work in process, finished goods, customer stocks) are used to produce official forecasts. While historical results show that the central business planning team does reduce the volatility of the geo data and often achieves improved accuracy, their track record shows that overcalling or undercalling sales, especially during product ramps, is not as rare as we might hope. These missed calls can lead to significant surpluses or shortages that take money right off the bottom line. Intel's factories, keen not to get caught in these situations, do not always build to the official forecasts. They also use models to help maintain proper inventories, smooth production, and achieve high operating efficiency, but our research has found no evidence to date that this final judgment improves demand fulfillment systematically or repeatedly.

The challenge of demand forecasting is real and costly. Demand risk is among the greatest threats facing Intel and other manufacturing firms day to day. To demonstrate how formidable demand risk can be, the following are actual situations we have discovered:

- Various groups across Intel estimated sales of a new product over an initial period after launch to be anywhere from one million to four million units.
- Two similar products (common architecture) were released within a quarter of one another in different (essentially non-competing) market segments. One resulted in a shortage, the other in a surplus.
- Geo forecasts for one new product were as low as 13.5 million units for a fixed period, while official forecasts were guiding the factories to build 26.5 million and sales targets were 27 to 28 million.
- Two products were projected to sell over 10 million and below 7.5 million units during a future period. In three months the forecasts flipped to under 5 million and over 11 million, respectively.
- Sustained growth in the mobile PC segment beginning in 2004 caught the whole industry by surprise. Four quarters of year-over-year growth, roughly double what was expected, made for a tough supply picture.

Certainly not all misses are this extreme. Intel's forecasting teams routinely perform quite well given the challenge of their task, often achieving forecasts with less than +/- 5% error. However, sustained high performance does not make up for each isolated miss that costs the

company millions to, conceivably, hundreds of millions of dollars. Everyone involved in forecasting at Intel continually strives to achieve better performance across the board, and we are always exploring new approaches that might bring improvement.

MARKET MECHANISMS AS FORECASTING TOOLS

In essence, all markets are prediction markets. The value of assets traded in a market depends on information that is not fully revealed and will not be known for some time, if ever. Market valuations are explicitly or implicitly predictions of that unknown information, perhaps the future value of a commodity, the expected cash flows generated by a firm, or the outcome of a potential corporate merger.

While commodities futures are often used as financial instruments to hedge long or short positions, the markets also reward traders with better information while punishing those with worse information. Giving traders incentives to reveal good information is the core function of prediction markets, even where no underlying assets, in the traditional sense, are available to be traded. Prediction markets trade future events or outcomes, and the settling process amounts to using a documented and published formula to determine winners and losers and to pay out incentives. Many experiments and real-world tests show that market mechanisms can be implemented simply to create predictions and that these systems work rather well.

Perhaps the best known of all prediction markets are the Iowa Electronic Markets, which enable traders to forecast the outcomes of future elections. The power of these markets to generate forecasts accurate and stable enough to inform decision makers has been demonstrated for nearly two decades [5]. Another set of experiments at Hewlett Packard demonstrated the ability of prediction markets to call future sales more effectively than traditional forecasting processes [6].

In our research at Intel we are extending the idea of prediction markets to create "forecasting markets," which are essentially prediction markets or similar IAMs integrated into the company's standard, ongoing forecasting processes. Participants reveal not just an expected outcome but a series of expected outcomes for the same variable over time. So, the forecasting market captures individual and collective assessments about trends such as increasing or decreasing demand just as weather forecasts anticipate warming and cooling trends.

We believe that three factors enable markets to outperform other types of forecasting systems and more effectively move information from source to decision maker. First, the features of anonymity and incentives

work together to draw out good information. The experiments of Kay-Yut Chen and Charles Plott at Hewlett Packard suggest that people provide the best information when rewarded to do so and when protected from potential ramifications of expressing their honest opinions. Incentives encourage participants to search for the best information they can find and reward trading behavior that is unbiased. Anonymity helps prevent biases created by the presence of formal or informal power, the social norms of group interaction, and expectations of management. We found many individuals at Intel who told us that their opinions sometimes differ from stated targets or unstated expectations. Looking back at forecasts that were off substantially, we have been told that teams sometimes did not believe the forecast they published but were pressured, perhaps overtly, to adjust forecasts upward or downward. To the extent anonymity and incentives curb bias and motivate the hunt for good information, they should improve the signal created by market mechanisms.

Second, the simple mechanism of aggregating data through a survey or market has two remarkable properties. It smoothes results over time, which is great for guiding supply, and it tends to produce a group forecast more accurate than the forecast of at least a majority of individual participants. A study by Scott Page demonstrated that even among a fairly homogeneous group this effect holds true. In the context of forecasting selection order in professional sports drafts, he found that averaging the individual forecasts of experts soundly outperformed the forecasts of any individual [7].

Finally, in many forecasting examples it has been found that increasing the diversity of a pool of participants increases the accuracy of the collective forecast. As long as each additional participant brings some information, adding more, diverse opinions improves the collective judgment. This condition holds true in many cases because good information tends to be positively correlated and sums, while errors are often negatively correlated and cancel [8].

DESIGN CONSIDERATIONS AND ELECTIONS

The first steps toward implementing a new IAM are finding business problems to address and teams interested in gathering better intelligence to solve those problems. In the context of demand forecasting, we started by partnering with two teams responsible for developing forecasts for product families. We determined that quarterly unit sales—with rules to define exactly which sales are included or excluded—would be useful to forecast with an IAM. With agreement on the result to be forecasted, the design process begins.

We have found through the development of current and upcoming IAM implementations that design considerations can be organized into five categories: interface, information, incentives, integration, and inclusion. The Appendix “Five categories of considerations for designing Information Aggregation Mechanisms” lists examples of design questions that should be evaluated within each category. Since we have found that design choices in one category often depend on choices in other categories, the five categories are developed more or less in parallel. Many companies implementing markets may start by designing the interface or simply assume that the only available IAM mirrors the stock market with regular, continuous trading periods and double auction trading. In fact, many interfaces exist, and choosing the best one should be guided by many other considerations. To demonstrate the application of the five categories of design considerations, the remainder of this section covers the design process for our original pilot market.

We began our design process by considering *inclusion*. As a first experiment we wanted to enable the central business planning team that creates the official forecasts to generate a collective forecast using an IAM. We would compare their collective forecasts created through the IAM to their current and historical collective forecasts created through Intel’s standard processes. We invited this team, and other participants who had a global perspective on the business and function more as analysts rather than sales or marketing staff. It was a relatively homogeneous pool of experts (but not without differences of opinion) and about as unbiased a group as we could put together, and we felt it would be a good baseline for future experiments that would tend toward greater diversity and bias. The total pool invited numbered from 20 to 25.

We carefully weighed how the IAM process would *integrate* with the workflow and processes of our participant pool. Knowing that the official forecasts are published monthly and that the potential participants are quite busy with that process for nearly two weeks out of the month, we decided against a continuous market; instead, we elected to time a snapshot IAM at the midpoint between official forecast publication dates. This scheme would maximize participation while effectively doubling the beat rate of new forecasts. The official and IAM forecasts would leapfrog each other, each outcome feeding the other process roughly two weeks later. (Since having the official forecasts and IAM forecasts influencing each other was unavoidable, we decided at least to make their interaction systematic.)

Structuring the *information* in the market was simple, as we decided to mirror the structure of the regular forecast. Each market would create separate forecasts for unit sales

of a product family in the current quarter, the next quarter, and the quarter after that. A packet to be sent to all potential participants before each market was developed. It included the definition of the results to be forecast, how incentives would be awarded, instructions for using the forecasting application, the current official forecast, and a small set of (already available) information such as historical sales and current orders. The market interface itself would provide some information during and after each snapshot. Once actual results were determined, prizes would be announced to individual winners, and a list of prizes awarded, showing amounts but not recipient names, would be published to the whole participant group. At no point would lists of participants be published, giving all participants the option of anonymity.

The *interface* design was based on the experience of our academic partners and the design choices in the above sections. We knew which information we wanted to collect and that we wanted to use a monthly snapshot to collect it, so the team opted for a synchronous Web-based application that seemed a good fit. It is essentially a survey mechanism that enables each participant to create a probability distribution of unit sales while watching others enter their distributions. Participants can learn from the aggregate forecast of the group while continuing to invest their own individual budgets into the offered investments, each corresponding to a range of potential unit sales. This method had demonstrated solid results in laboratory experiments outside Intel and can develop a complete forecast in as little as 15 to 20 minutes.

The behavior of participants in the IAM is based on the way *incentives* are awarded. Once the actual result is known, investments made in the range containing the result are placed in a drawing for cash prizes. Each participant's chances of winning prizes are proportional to his or her share of all investments in the winning range. We wanted to avoid extremes of all incentives going to a single winner or dividing incentives among all winning tickets, because we did not want to encourage participants to concentrate their investments too narrowly or spread them too broadly. Incentives were a hot topic in the design phase—how large should they be? In the end we settled on an amount significant enough to attract and retain interest (we hoped) but not large enough that employees might shirk other job responsibilities.

Our overall design structures each investment as a decision based on both the individual's expectations for the outcome and the aggregate group prediction. Participants weigh owning lower percentages of more likely outcomes against higher percentages of less likely outcomes. In the end, we believe the system works well because having each participant weigh the conditional probabilities of various outcomes creates a robust

collective forecast. And, the final outcome of the market—the amount of investment in each potential range of unit sales—forms a probability distribution based on the intelligence of the entire group.

We analyze not only the collective forecast but also the transaction records of individual traders. Assessing trading behaviors and inferring strategies from the behaviors helps us understand how the systems work, under what conditions they might work well, and why certain types of participants or investment strategies may contribute more or less toward a good (or bad) outcome. Over time we also expect to use the observed data to determine whether the formal outcome is the best possible forecast the market had to offer. Perhaps other information based on the demonstrated knowledge and track records of participants, individually or grouped by function, geography, or experience, will lead us to be able to handicap traders by the knowledge they impart to the system over time.

RESULTS

We are using three primary measures to assess the performance of our markets: accuracy, stability, and timely response to genuine demand shifts. Having run pilot markets for approximately 18 months, we are starting to get a sense for how the markets are performing. Although the market forecasts and official company forecasts are not independent, it is nonetheless interesting to compare the signals and then assess how effectively they are working together. In terms of accuracy, the markets are producing forecasts at least the equal of the official figures and as much as 20% better (20% less error), an impressive result given that the official forecasts have set a rather high standard during this time period with errors of only a few percent. In the longest sample to date, six of eight market forecasts fell within 2.7% of actual sales. The accuracy of the official and market forecasts has been remarkably good, well within the stated goal of +/- 5% error for all but a few individual monthly forecasts. Until more results are generated over time we will not be able to determine the extent to which this strong performance stems from the introduction of the market forecasting process. It is also possible that sales were unusually easy to forecast. Regardless, specific results from the pilots have shown the value of the market forecasts and are leading us to believe the markets are having a positive impact.

On one occasion we saw the first market for an upcoming quarter's sales vote "no confidence" on the prior official forecast. Ranges of potential sales in the IAM are structured so that the prior official forecast is roughly centered in the set of ranges. A "no confidence" vote occurs when all investments from participants come in

either above or below that official forecast, meaning that the group believes there is a 100% chance of falling on one side of the official forecast. The only time this has occurred the market forecast was correct. The official forecast published prior to the market forecast was off by over 10%, and the market led it in the right direction.

Much like public stock markets, we have seen our IAMs react quickly and decisively to strong news and then take time to assess and properly discount it. One IAM dropped by 4.7% and then bounced back to almost exactly where it was before the drop. This was not accidental. A rash of cancelled orders and bad news that appeared to signal softening demand turned out to be an aberration, and the market needed time and additional information to make that call. These sorts of sudden shifts are unusual. In fact, the IAM forecasts are quite stable, with as much as 20% less fluctuation from month to month than the official forecasts during the same period. The business planning team responsible for the official forecast observed that the market signals were more stable and implemented a new process step to try to filter noise from each new official forecast.

Surprisingly, the market forecasts are not necessarily improving as the forecasting horizon shrinks. Although we will need a longer history of data to draw a firm conclusion, we have some evidence that the forecast is as likely to get worse in the final month before the actual result is known as it is to get better. The reason, as we understand it today, is that as the amount of signal goes up rapidly toward the end of the period, the amount of noise goes up rapidly as well. As the amount of information explodes and the time to assess it shrinks, it would not be a surprise to see humans unable to tell the forest from the trees. Fortunately, forecasts out in the 3-8 month horizon, which provide the factories ample opportunity to plan product starts, are performing quite well.

Another key set of results is feedback from owners of the official forecasts, as well as market participants. Discussion with the owners has centered around learning to produce better official forecasts from the market results. The value or credibility of the results has never been questioned; in fact, the one month we were late publishing the market results brought reminder e-mails from the owners. Not long into the pilots the owners began discussing new markets for other key forecasts. Clearly, they are seeing the value of this new data source. Participants have been quite positive as well. Quotes such as "I enjoy participating in the trials" are common. Another trader cited the IAM process as a welcome break from the often mundane job of forecasting: "I think it's great we're doing this simply because it makes work more fun and incentivizes us to do our homework and make the right call, which should lead to better results." We are also

amused that although we never publish the list of participants and winners, everyone knows who participated and who won.

Based on the results and word-of-mouth advertising, interest in expanding the research into new parts of the business is growing. We expect the number of forecasting markets to quadruple in the next three months. More implementations producing more data will accelerate our pace of learning.

CHALLENGES

The main challenge in implementing IAMs in a corporation, as with many innovations, is securing buy-in that the time invested is worth the potential benefits. It helps that certain teams are forward thinking and some of these same teams have been burned by poor forecasting performance in recent years. We generally look for those customers first. In fact, we have had little trouble finding volunteers—teams—that want to try something new. At present we have as many teams wanting to run experiments as we can accommodate.

As we propose market mechanisms to aid with forecasting, potential participants and managers have most often expressed three concerns: incentives, anonymity, and groupthink. Regarding incentives, why does it make sense to pay for performance when employees are already paid to do their jobs? This is an interesting question because most businesses think nothing of offering commissions for sales. Do businesses not already pay the sales force, and should they not be selling anyway? We learned that the first time an Intel factory achieved all of its performance targets across a suite of metrics was when a program offered direct incentives, i.e., cash to each individual employee, for that precise outcome. We do not feel it is out of line to offer forecasters incentives for performance or general market participants incentives for good information. The potential value of the improved forecast is orders of magnitude greater than the cost of the incentives.

The feature of anonymity is somewhat incongruous with Intel's culture of direct, constructive confrontation. If employees disagree they engage and resolve their differences. Allowing employees to participate in systems without identification (to others, not to the research staff running the system) is foreign and may be difficult for some employees to swallow. However, Intel is also a company that values results, and there is room in the culture for improvement.

Can IAMs enable or even cause groupthink? A classic approach toward defeating groupthink is assigning private roles to individuals. For instance, everyone on a team gets a card, and everyone knows that some cards say "devil's

advocate.” With some individuals assigned the role but no one knowing who those individuals are, everyone is able to dissent with less fear of reprisal. A market system where all participants are anonymous and incentivized for performance takes this approach to the limit, freeing individuals to express themselves. Interestingly, although some IAMs that enable participants to observe the group forecast develop could potentially lead to artificial consensus, in all market-like mechanisms the primary opportunity to win and win big comes from being right when everyone else is wrong. This feature certainly helps prevent too great a consensus.

A few more specific challenges have also been faced. Running synchronous IAMs across a global corporation is a problem, given that it is always 2 a.m. somewhere. Teams are reluctant to schedule anything out of normal hours, and it is challenging to find a good time for any large group of people to do something together. This issue is forcing us to consider asynchronous approaches as well.

Another issue has been dealing with a world made up of local geographies. If global sales are the sum of several geographies’ sales, how does one tap local knowledge to forecast the global outcome? We have found our experts within a geography reluctant to try to forecast global results because they feel they do not have enough information to perform the task. That leaves three choices: limiting the markets to global forecasts and participants with a global view, running multiple markets specific to local geographies, or swaying the local experts to participate in a global forecasting market. In the latter case, participation is a critical consideration. If sales are 50% geo A, 30% geo B, and 20% geo C, do we need participation roughly proportional to sales from each geo? Or, is a result weighted by recent sales preferable to the formal market result, which is weighted by participation?

Two remaining challenges we have identified are scalability and long horizons. Forecasting total sales for a product family is valuable, but it does not address the mix of products or SKUs within those products. The market solution probably cannot scale to forecasting all SKUs, and it may not even be suited for that task. Perhaps the right balance is forecasting total product family sales and key products—new or of strategic importance—that will have the greatest impact on financial performance. Regarding horizons, markets are better suited to the short term. Incentives lose power if the payoff is too remote, and feedback is important for driving participation and performance. Forecasting a result within a few quarters seems to work, but over a year begins to feel like a stretch. We are experimenting with alternative market structures that might help forecast the distant future while paying incentives more quickly.

SUMMARY AND CONCLUSIONS

Demand risk is a serious threat to bottom-line performance at Intel and other manufacturing firms. Our studies identified numerous cases where poor information flow led to poor forecasts, which in turn led to decreased business performance.

Markets, and more generally IAMs, promise to help companies address demand risk and other business challenges by improving organizational information flow. Based on results to date, our IAM implementations appear to have had a desirable impact on forecast accuracy and stability. The key drivers that we believe have led to strong performance are 1) anonymity and incentives, which encourage honest, unbiased information, 2) the averaging of multiple opinions, which produces smooth, accurate signals, and 3) feedback, which enables participants to evaluate past performance and learn how to weigh information and produce better forecasts.

Although greater diversity in our participant pool may improve the collective forecast, many ways to increase diversity also increase the potential for bias in our real-world scenarios. Crowds have demonstrated the ability to solve problems such as estimating the weight of a steer [8] or choosing the winner of an upcoming election. But, the prediction may not turn out so well if the new diverse opinions come from those who will profit from selling a heavier steer or from members of the election campaign team for one of the candidates. We hope to explore this issue in upcoming phases of our research.

Our framework for designing IAMs is enabling us to systematically develop new solutions for a number of business problems, and experience, be it in the form of successes or failures, will make us more effective designers. Of particular interest are forecasts that tend to break the simplest IAM designs, predictions with long horizons or predictions whose outcomes may never be known. For instance, was product A better to bring to market than product B? We are defining solutions to these problems today and will soon be testing them in our organization.

Many business processes in use today are neither perfectly effective nor efficient; yet, they are the lifeblood of the organizations that use them. IAMs are a new approach toward business management, promising, and at the same time frightening to potential adopters. As with many such innovations, starting small and running in parallel to existing processes are keys to success. As our trials are demonstrating excellent results at remarkably low cost, expanding their use at Intel is a natural and expected outcome.

APPENDIX

Five Categories of Considerations for Designing Information Aggregation Mechanisms

Information: What is the result to be forecast, how is it defined, when is it known, and to what precision is it known? What range could the result cover, and what granularity of forecast is material to the business? What level of granularity might participants be able to forecast? What information is provided to participants in advance of the market, during the market, and after the market? Where is the line between providing a baseline to improve inputs and providing an anchor that might undermine accurate information? What types of analyses will the information produced by the market enable, and which business decisions will be informed by that analysis?

Integration: Which business processes are related to the market? What is the timing of key decisions or events that will inform the market or be informed by the market? Which other processes attempt to forecast the same result, and should the market function independently or coordinate with the other processes? Based on business cycles or other processes and workflow, when are participants available or busy?

Inclusion: Who should participate in the market? How many participants are needed to achieve good results? Should participants have local and specific views or more aggregate views? Should groups that demonstrate bias in other forecasts participate, and would they bring the same biases to the market? Can people across wide ranges of time zones participate together, and will participation skew results? Might anyone outside the firm participate?

Interface: How will individuals interact with the market? Will the market be continuous or provide snapshots? Is participation synchronous or asynchronous? What level of anonymity is provided? How do traders convert their knowledge and preferences into data and, ultimately, collective forecasts?

Incentives: What will motivate participants to enter the market, and what will motivate strong performance? How do incentives compare to salaries, awards, or other incentives within the corporation? To what extent are strong performance and bragging rights incentives? Will management support the incentives? Are systems available to pay the incentives out without undue cost? Will those processes scale to large numbers of participants?

ACKNOWLEDGMENTS

Thanks to our academic partners whose creativity and insight has yielded new approaches toward solving tough problems: Feryal Erhun, Paulo Goncalves, Jim Hines, Blake Johnson, Thomas Malone, Charles Plott, and Jim Rice.

Thanks to partners in Intel who have driven, supported, and challenged this research: Nicholas Day, Jim Kelloso, Karl Kempf, Adam King, Rich Krigger, Jeff Loose, David McCloskey, Mary Murphy-Hoye, Chris Parry, Jason Seay, and Paul Thomas.

REFERENCES

- [1] Erhun, F., Gonçalves, P. and Hopman, J., "Managing New Product Transitions," *Sloan Management Review*, vol. 48 num. 3, Spring 2007, pp. 73-80.
- [2] Kiviat, B., "The End of Management?," *Time*, July 6 2004 at <http://www.time.com/time/magazine/article/0,9171,660965,00.html>*
- [3] King, Rachael, "Workers, Place Your Bets," *Business Week*, August 3, 2006 at http://www.businessweek.com/technology/content/2006/tc20060803_012437.htm*
- [4] Hopman, J., "Managing Uncertainty in Planning and Forecasting," *Intel Technology Journal*, vol. 9, Issue 3, 2005, pp. 175-183.
- [5] Berg, J. E. and Rietz, T. A., "Prediction Markets as Decision Support Systems," *Information Systems Frontiers*, 5(1), 2003, pp. 79-93.
- [6] Chen, K.-Y. and Plott, C. R., "Prediction Markets and Information Aggregation Mechanisms: Experiments and Application," *California Institute of Technology*, 1998.
- [7] Page, Scott E., *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press, Princeton, NJ, 2007, pp. 210-212.
- [8] Surowiecki, James, *The Wisdom of Crowds*, Doubleday Press, 2004, New York, New York.

AUTHOR'S BIOGRAPY

Jay Hopman is a researcher and strategic analyst in the IT Research Business Agility cluster. His research interests include product management, decision making, forecasting, planning, and organizational coordination. Jay received a B.S. degree in Computer & Electrical Engineering from Purdue University and received an

MBA degree from the University of California, Davis. His e-mail is jay.hopman at intel.com.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel Leap ahead., Intel Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel's trademarks may be used publicly with permission only from Intel. Fair use of Intel's trademarks in advertising and promotion of Intel products requires proper acknowledgement.

*Other names and brands may be claimed as the property of others.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS[®] Ready mark under license from Palm, Inc.

Copyright © 2007 Intel Corporation. All rights reserved.

This publication was downloaded from
<http://www.intel.com>.

Additional legal notices at:
<http://www.intel.com/sites/corporate/tradmarx.htm>.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

developer.intel.com/technology/itj/index.htm