An agent model for the High-End Gamers market

Teun Adriaansen[*],
Dieter Armbruster[†*], Karl Kempf[‡] and Hongmin Li[§]

November 12, 2010

Abstract

Understanding the driving forces for the markets of their products is a basic necessity for any business. Quantitative models are either aggregated over large market segments or restricted to utility models of an individual's buying decision. While the aggregate models acknowledge that customer interactions are important they do not model them and hence have no way to adjust their model to changing business environments. This paper develops crucial methodology to bridge this gap between the individual decisions and the overall market behavior using agent based simulations to model the sales of computer chips in the High-End Gamers market. The simulation environment is dynamic and models the succession of 19 products introduced over a 40 month time horizon which includes the recession of 2008 - 2010. Simulated sales are compared to actual sales data and are used to adjust the parameterization of the agents and their environment. Only two agent parameters are sufficient to obtain a very reasonable fit between simulations and data: The amount of money available for the gaming hobby and a parameter related to the gaming skills of the High-End Gamers.

1 Introduction

Success or rejection of a new product in the market is a typical emergent phenomenon. Each consumer decides to buy (or not) based on his/her own perception of the utility of the product. The collection of all of these individual decisions aggregated over the market population over a time horizon defines

[*]Department of Mechanical Engineering, Eindhoven University of Technology, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands, email: a.c.adriaansen@student.tue.nl

[†]School of Mathematical and Statistical Sciences, Arizona State University Tempe, AZ 85287-1804, USA, email:armbruster@asu.edu

[‡]Decision Engineering, Intel Architecture Group, Intel Corporation, 5000 W. Chandler Blvd, MS CH3/10, Chandler, AZ, USA 85226

[§]W.P. Carey School of Business, Arizona State University Tempe, Arizona 85287, USA, email:hongmin.li@asu.edu

success or failure of the product introduction. Every corporation that introduces new products or technologies into the market aims to understand these emergent phenomena to forecast sales, production, revenues and profits.

There exist a large number of aggregate models for the way consumers adopt a new product, the Bass model [2] being the most well known. Conversely, models exist of an individual that makes choices based on the perceived utility of buying a particular product [21]. There is a large gap between the models that parameterize individual behavior and the models that parameterize aggregate behavior. Not much is known on the relationship between the parameters that drive the utility for an individual and the parameters of the aggregate differential equation models. The standard approach a the aggregate level has been to perform statistical fits of the parameters of the aggregate model based on linear or nonlinear regression on sales data to describe e.g. the introduction of color TVs etc. [22]. Such an approach gives no indication about the parameters that describe the utility of the product to an individual.

This paper develops crucial methodology to bridge this gap between the individual decisions and the overall market behavior. We will do this in a dynamic environment that models the succession of several products introduced over a period of time which is subject to the impact of the recession of 2008 - 2010. The specific market we are concerned with is the semiconductor market of high-end performance computer chips that are almost exclusively purchased by people playing competitive games on the Internet. We call this market the High-End Gamers market and the participants in that market the High-End Gamers. This group should not be confused with players using game boxes (WII, Xbox etc). The High-End Gamers market is characterized by the fact that there are only two suppliers in the market (which we refer to as the Blue and Green team, loosely modeled after INTEL and AMD).

We develop Agent Based Simulations (ABS) to model the buying behavior in the High-End Gamers market. We parameterize the ABS to fit sales data for 10 high-end chips released by the Blue Team and 9 high-end chips released by the Green Team over a period of 40 months between January 2006 and April 2009. Each of these chips is characterized by a performance measure and a price path. We obtain reasonable fits between the simulations and the data using only two agent parameters: a monthly income stream and a parameter related to their gaming skills. We discuss further improvements of the model related to secondary traits such as brand loyalty and anticipation of upcoming releases of new chips, and show that their influence is weak.

The model can then be used to study market features that are not readily measurable otherwise. For instance, we can determine the average time between the acquisition of a new chip for High-End Gamers, the influence of the recession on various market segments, and the optimal release timing and pricing for successive generations of chips. Notice that these are conditional properties: given the basic ABS and the resulting population of High-End Gamers that provides the best fit to the data, we can infer these secondary properties for that population. Most importantly ABS' parameterization are intuitive and communicable to management and the sales force.

Since ABS are able to characterize dynamic changes in the market, we believe that ABS offer another tool to understand markets by acting as a test-bed for behavioral theories. For example, it might be possible to use the ABS to forecast sales. Using some of the sales data as training data we forecast the rest of the data. Initial results indicate that forecasts generated by ABS are at least as good as standard forecasting methods used by Intel [30]. Results of a more detailed forecasting study will be reported in a forthcoming paper.

In Section 2 we discuss different modeling approaches for demand of multiple products, in particular the aggregate approach of the Bass model (section 2.1), the individual choice model approach (section 2.2) and the general structure of agent-based modeling (section 2.3). Section 3 discusses the very specialized market for the high-end computer chips and presents the target sales data together with the price and performance properties of the chips. Section 4 describes the agent-based simulation and presents the simulation results. Section 5 compares the sales generated from the simulation model with the actual sales and studies the sensitivity of the match to variations in some characteristic parameters of the ABS. Section 6 concludes this report and lays out future work related to forecasting for other market segments such as servers and laptop computers.

2 Modeling paradigms

2.1 Diffusion: The Bass model

The Bass diffusion model [2] is an aggregate description of the way new products get adopted in the market. It is based on the assumption that there are different classes of buyers called the innovators and the followers. The innovators will buy a product based on its intrinsic value while the followers will buy a product because their neighbors, colleagues, or everybody else is buying the product.

Following this concept Bass develops a difference equation for the number of people N_t that have adopted a new product by the time period t:

$$N_t = N_{t-1} + p(m - N_{t-1}) + q\frac{N_{t-1}}{m}(m - N_{t-1}). \tag{1}$$

Here m describes the total market (note, if $m = N_{t-1}$ then $N_t = N_{t-1}$), q is a parameter associated with the followers group measuring the influence of the people that have already bought the product and p is a parameter associated with the innovator group measuring their propensity to buy which is affected by the intrinsic utility of the product for the buyer and by marketing decisions such as advertising intensity.

Equation (1) is a version of the logistic equation also known as the Verhulst equation in biology that can easily be solved analytically. Typical sales curves $\frac{dN(t)}{dt}$ show a single maximum and the cumulative sales $N(t)$ show a typical saturation curve.

Although the model is approximate and requires the guess of the parameters p, q and m, it has been hugely influential. It has been fitted with some success to the introduction of color TVs, micowave ovens and many other products and has been extended in numerous ways (see [22] for a comparative review of various diffusion methods). While its simplicity is appealing, even after all the extensions have been added, it remains an aggregate model which ignores the characteristics of individual buyers except their status as innovator or follower. Hence, it is very difficult to use the Bass model to answer the *what if* questions: *What if we increase the price of the product? What if we delay the introduction of the new product? What if the general economy turns sour? etc.*

2.2 Choice models

The theories of individual choice assume that an individual, when facing a set of product alternatives, associates a preference level (utility) for each alternative, and chooses the product that generates the highest utility. The choice models then predict the choice probabilities for each alternative based on the form of utility functions assumed. Luce [16] and McFadden [21] show that, if the utility is a linear function of the product-specific attributes, such as quality and price, with a Gumbel-distribution additive noise,

$$u_i = \beta' \mathbf{x}_i + \epsilon_i \ ,$$

then the choice probability of product i have the form

$$\frac{\exp(\beta' \mathbf{x}_i)}{1 + \sum_{j=1}^{M} \exp(\beta' \mathbf{x}_j)}$$

and the probability that the customer chooses not to buy any product is given by

$$\frac{1}{1 + \sum_{j=1}^{M} \exp(\beta' \mathbf{x}_j)} \ .$$

This is often referred to as the Multinomial Logit (MNL) model. It has been used to predict individual choices [19], as well as aggregate market shares [3]. Both the standard MNL, and the related Nested Logit models [18] have been widely employed to fit empirical data in many industries. The typical method for estimating the coefficients in the linear utility function is the Maximum Likelihood Estimates (MLE) method. Applications of the choice model include travel mode selections [4], choices of home heating system [8], coffee purchases [15], subscriptions for local phone services [29], and automobile purchases [12]. A recent paper by [11] models the competition between Intel and AMD for microprocessors using a MNL demand model to study the effect of competition on prices and innovation.

The MNL model, and the more general Nested Logit model, produce demand functions that have simple analytical forms, which makes it easy for fitting empirical data and estimating the parameters, as well as for model-based normative studies (see for example, oligopoly equilibrium analysis by [1], and optimal pricing decisions by [14]). In some variant form of the choice model, an individual may face a budget constraint [5], or the customer population is heterogeneous in the sensitivity toward the product attributes [20]. However, in most applications of the choice-based models, the decisions by the individuals are limited to comparing the linear utility values. In addition, these models are static and do not allow dynamic, rule-based decision makings of the individuals, let alone interactions among the individuals. In practice, situations arise when management wish to learn more about customer behaviors when facing more complex dynamic choice structures; as we discuss next, agent-based models are a more powerful tool toward that end.

2.3 Agent-based simulation

Agent-based simulations have been used extensively in recent years mostly in the context of computational social science [6, 9] and computational economics [28]. Agent-based models provide insight into the general behavior of a system assuming the behavior of its elements is known. The mechanism by which global behavior is generated is called *emergence*. Typically, ABS generate many different emergent phenomena, some of which are obvious and expected and are used as a verification for the model whereas others are unexpected and new. The emergence of the latter is the main reason for performing ABS. Managerially, the unexpected or new emergent phenomena lead to in-depth understanding of the impact of various demand-shaping dynamics.

ABS define autonomous decision-making entities, called agents. An agent typically has a goal to achieve. Each agent individually assesses its state and environment and makes decisions that are either rule-based, or as a result of an optimization calculation or some other decision process. As a result of the decision, the agent may execute a sequence of actions. Repetitive interactions among agents change the state of the agent as well as the environment it is in [17]. The agent-based models simulate the simultaneous operations and interactions of multiple agents in an attempt to recreate and predict the appearance of complex phenomena.

A main feature of ABS is the way agents interact. Interaction may be stigmergic (indirectly via a changing environment), or direct with one or more other agents. Direct interactions may be based on some form of spatial or topological neighborhood:

- Global interaction (every agent interacts with every other agent) based on a global field;
- Local interaction (every agent only interacts with a local neighborhood of other agents);
- Local interaction with some degree of global reach based on the connectivity of a network.

Agent based models have been used more recently to describe diffusion-based models in physical, biological, social and economic settings. However, as noted in Bonabeau [6] there are few real business applications and the emphasis has been on comparing ABS to (differerential) equation based models. For instance ABS are compared to aggregate descriptions based on compartmental models developed in biology (SIR models, susceptible, infected and recovered populations) [26]. Parunak et al [24] compare ABS and systems dynamics models based on ordinary differential equations for supply chains. Alternatively, agents are used to model the dynamical influence of certain decision rules e.g. for auctions in the electricity markets [25, 27] or promotional strategies that support the launch of a product [7]. All of these models are self-referential and at the level of highly abstracted descriptions and are not intended to be quantitative.

We found two notable exceptions to this trend: Goldenberg et al [13] analyze the market introduction of an e-mail software product in four different markets. Focusing on daily sales data they provide a sales tracker framework to parametrize agent-based models. They compare their short term predictions with the actual data and show good agreement. In another paper North et al [23], in collaboration with Proctor and Gamble, apply agent-based modeling to develop a multi-scale consumer market model called the virtual market learning lab as a test bed for consumer market business problems.

Our approach complements and extends these two papers: We are concerned with characterization and forecasting of repeated sales of new generations of a product and the long term behavior of the High-End Gamers market, a high tech market that has some very different characteristics from a consumer product market. In particular, the product and its use that we study is simple enough that we can model its obsolescence which generates sales of the new generations of the product.

3 The High-End Gamers market

The High-End Gamers market is a very specialized market for the fastest computer chips. Only a few of these chips are purchased for a purpose other than playing Internet games. High-End Gamers play games against each other, typically in a one-on-one situation on the Internet. They are usually very competitive and spend a large amount of their time and their disposable income on their hobby.

There are only two manufacturers of these high-end computer chips worldwide. From January 2006 until April 2009 the Blue Team released 10 high-end chips and the Green Team released 9 high-end chips. The estimated sales data of these chips are shown in Figure 1 and Figure 2. Figure 1 shows the total monthly sales of each company. Figure 2 shows the monthly sales data for each chip.

3.1 Performance of high-end chips

The relative performance advantages of these chips are available through published benchmarks. Figure 3 shows performance numbers and manufacturer's

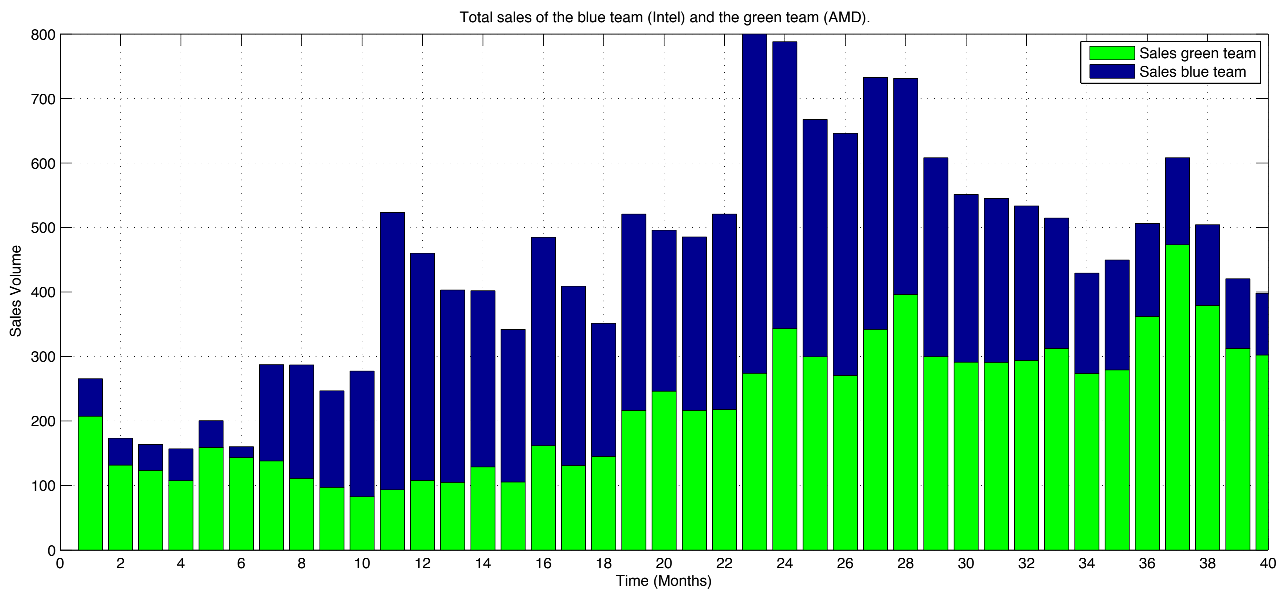

Figure 1: Total monthly sales data of the Blue and Green Team high-end chips (period: January 2006 until April 2009).

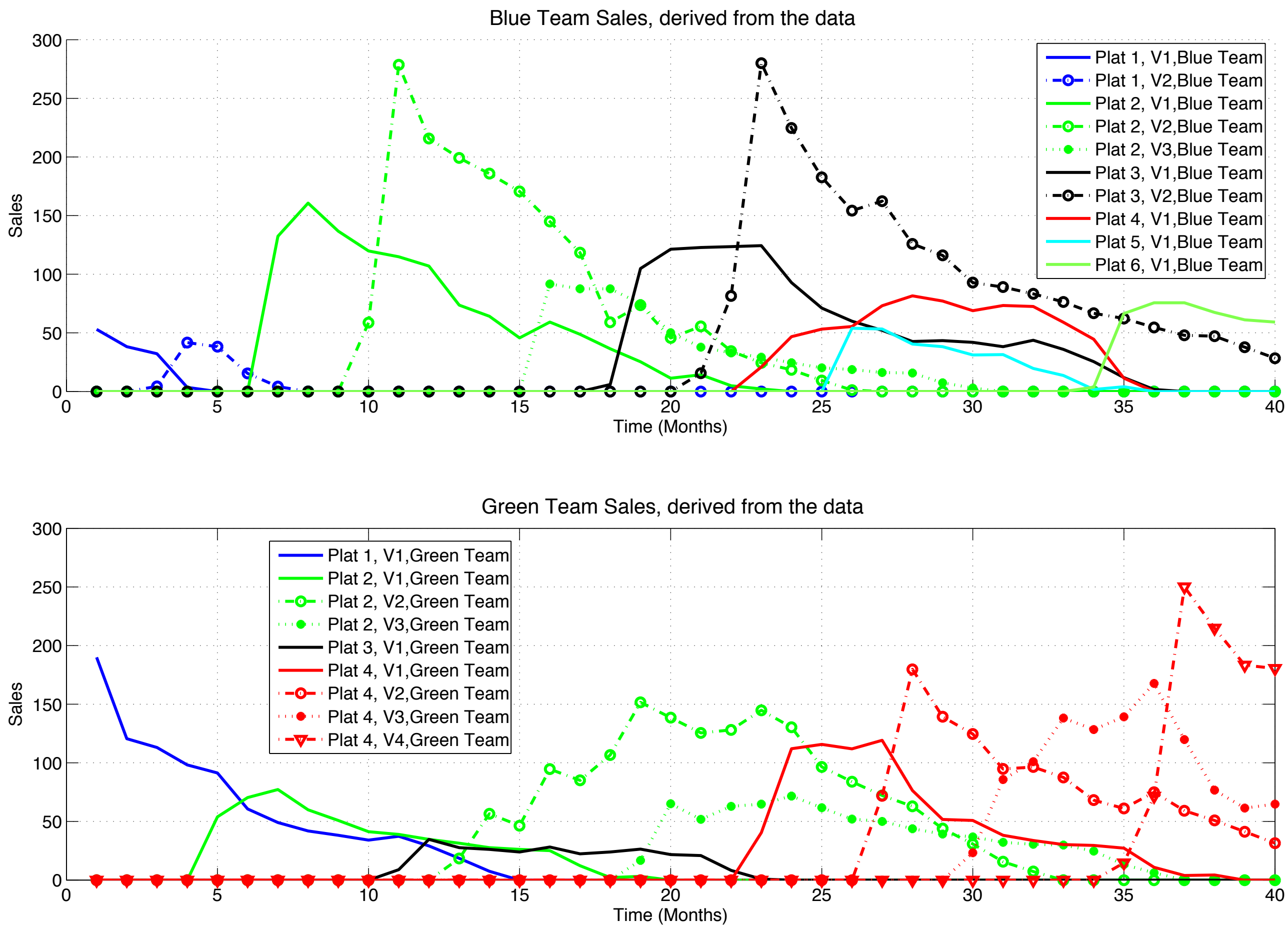

Figure 2: Monthly sales data per chip for the Blue and Green Teams.

release dates for all 19 chips. Note that the green chips were the performance leaders at the beginning of the time horizon while the blue chips had a significant performance advantage towards the end of the horizon.

3.2 Chip vs. platforms

Technological improvement in this market comes in two forms: In some cases, the new chip has a higher CPU frequency than its predecessor, but uses the same chipset, motherboard, graphic card, sound card, etc. We say that the two chips share a platform and they can be exchanged easily. Buying such a chip is a relatively inexpensive improvement of the system. In other cases, platforms are changed and the new chip requires new support hardware. Such improvements are expensive since a gamer has to basically buy a whole new system. Platform changes occur much less frequently than frequency changes. Chips of the same color in Figure 2 represent the same platforms. A change in platform in Figure 3 is represented by a symbol change.

3.3 Prices of high-end chips

After a chip or platform has been released to the market, the chip's price will start to drop at certain moments in time. One of the most important differences between the prices of the green and blue chips is that prices of the green chips did decrease faster. The prices for single chips and for the platforms containing the chip are given in the appendix A.

4 The simulation model

4.1 Basic assumptions

Since real markets are very messy and since there are many different vested interests involved in the characterization of such markets there is a great deal of uncertainty over their basic features. The High-End Gamers market is no exception.

For instance, the actual sales data are not publicly available. While we have access to the real data for the Blue Team they have been altered to protect business interests without changing the fundamental structure of the sales curves. In addition, the actual sales timing for the Blue Team represent sales to retailers and OEMs and not necessarily sales to the end user. However, since we do not have access to the sales data from the OEMs and since we aggregate data into monthly sales, we assume that those sales figures are identical to the sales figures to the High-End Gamer.

The sales data of a chip from the Green Team are estimated using the number of Google searches for this product similar to recent reports on predicting the growth of the H1N1 flu epidemic [10] by tracking influenza related Internet searches. We determined that the Internet searches for a specific INTEL chip and its actual sales have a linear relationship with the sales of many high-end

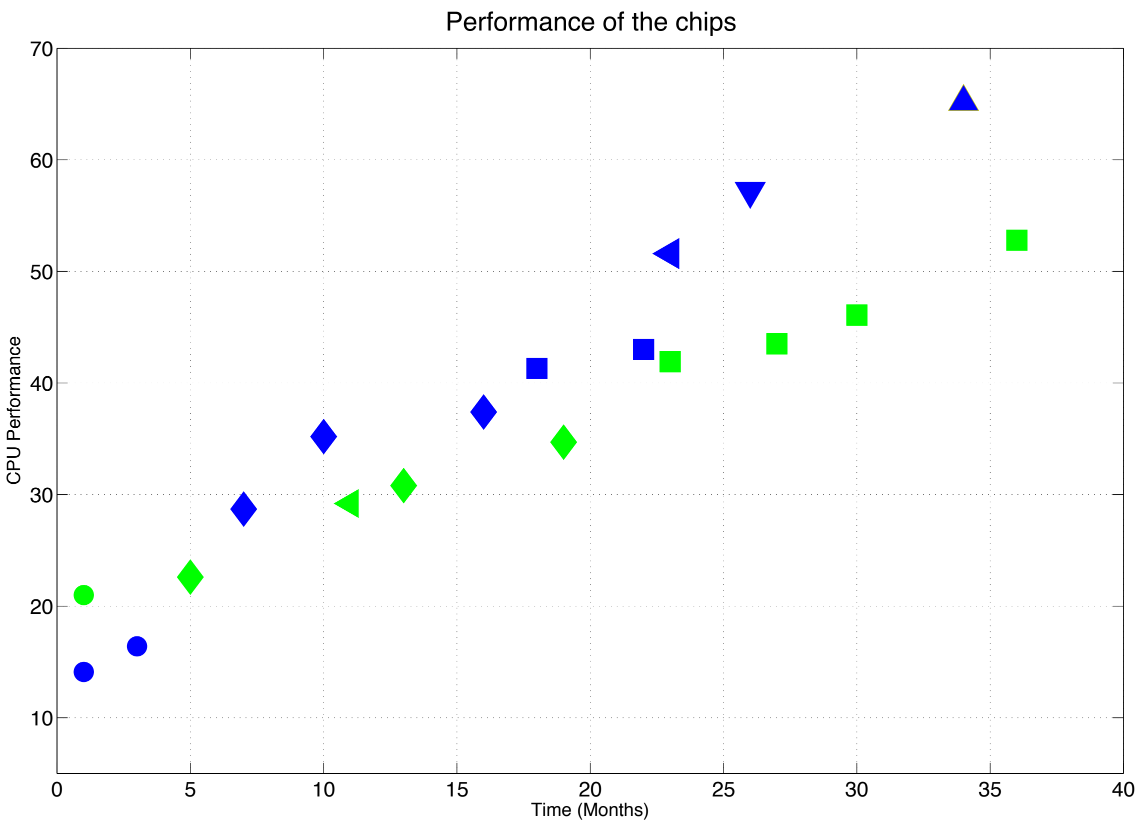

Figure 3: The performance numbers at release date for green and blue chips (period: January 2006 until April 2009). For color and shape legend see Section 3.2

chips and over several orders of magnitude of sales. Based on the Google searches and this linear relationship, we used the regression method to estimate the sales of the Green Team.

Similarly, there is no consensus about the way to measure the performance of a chip. Different benchmarks give slightly different results. We allocate a performance number based on the average over several benchmarks and based on CPU frequency, Dual or Quad core and maximum power consumption.

Finally, we estimate the price history using Internet sites that record such data. These price estimates are a major source of uncertainty in the model. In particular, the variation of the prices for platforms is very high, since most gamers put together their own platform choosing their own components (motherboard, graphic card, etc.).

4.2 A single simulation

Our simulations typically involve a starting population of one to two thousand agents. In each month each agent will play 200 one-on-one games with random pairing against other agents. Games are reduced to a comparison of the performance of the current chip that an agent owns; the agent with the higher-performance chip wins with a probability of 0.85. A single simulation covers a time span of 40 months. Since we use a large number of agents we do not need to run many simulations to get good averages. We found no improvements in the convergence to averages beyond approximately 10 simulations of the time horizon for a fixed parameter set.

4.3 The behavioral model for our agents

The main attributes and actions of agents are:

Losses. Agents will keep a record of the number of losses suffered since their last update of their gaming chip. We have experimented with this feature to let agents remember only the cumulative losses during a recent time interval but found no fundamental difference.

Thresholds. Agents have a loss-threshold. As the losses of a particular agent increase they will reach the threshold for this agent which triggers the agent to search for a new chip. The threshold distribution varies widely over the agent population and is a surrogate for many sociological and psychological traits associated with the population of High-End Gamers. For instance, a low loss threshold, in addition to possibly being a personality trait, may indicate an agent that is a highly skilled gamer and expects to lose only a few games, a high threshold may indicate a novice gamer who expects to lose a significant number of games before blaming hardware for failures.

Income. An agent allocates a certain amount of money per month to this hobby. For short we call this contribution the agent's monthly income. Monthly income accumulates to capital that is currently available to the agent for purchases.

Decisions. Agents have a decision procedure to determine which chip to buy. We assume that the decision of what kind of chip to buy in the High-End Gamers market can be readily described by a small number of characteristic parameters: performance first and price second. Hence, once the agent is on the market it will make a performance evaluation of the available chips and buy the chip with the best performance, provided that the performance rating of the chip is an improvement over its current chip and provided that the agent has the necessary money. At the purchase, the agent's monetary budget will be reduced by the cost of the new chip/platform, its loss tally will be set to zero, the new chip will become its new attribute and a sale of the chosen chip will be registered.

Other attributes. Some of our agents show brand loyalty, expressed by the fact that they will change to a new chip of a different brand only if its performance exceeds a certain improvement threshold. Some of our agents are conservative and will not buy a chip that is brand new on the market. They will wait a few month before they consider the performance and price characteristics of such a chip when making a buying decision. The performance characteristics of an upcoming new chip are typically discussed on the web up to 6 month in advance. Some of our agents will postpone their buying decision to wait for the market release of a promising new chip.

Number of agents. Over the 40 month time horizon the total monthly sales vary dramatically (see Figure 1). They increase from about 200 per month in months 2-4 to almost 800 per month in months 23-24 and decline to about 400 per month in months 39-40. We treat these data as an input into our simulation and change the number of agents participating in games accordingly as shown in Figure 4. We believe that the strong overall growth initially is due to the appeal of High-End Gaming in part fueled by the expansion of high performance Internet connections and by the competition between the Blue and Green Teams which brings ever faster and powerful chips to the market. The reduction in sales from month 24 onwards is probably due to the influence of the recession (see details below).

4.4 Main parameters in the model

Monthly income and the loss threshold are the two main attributes of an agent. Both parameters are characteristic for a particular agent and each agent will be allocated income and threshold parameters at the beginning of the simulation after which they are held constant. The monthly income of 80% of the agents is sampled from a normal distribution (mean=200$, standard deviation 100$). We discard samples with an income of less than 100$ because the income, in relation to the cost of a new platform is too small for the agent to be consistently playing among the High-End Gamers. The monthly income of the remaining

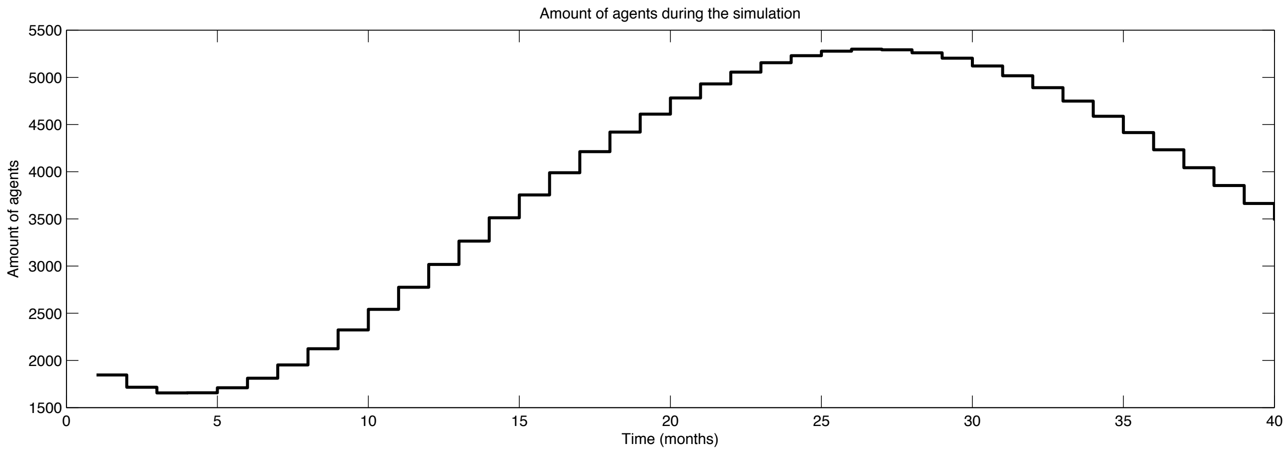

Figure 4: Number of agents during the simulation (initially there are 1800 agents)

20% of the agents is sampled from a Γ-distribution which ensures that there are some agents in the simulation that are very affluent.

The loss thresholds are sampled from a Γ-distribution. There are three distributions, that depend on the monthly income of an agent: one distribution for agents with a low monthly income, one for agents with an average monthly income and one for agents with a high monthly income. This method is used to give agents with higher monthly income on average lower thresholds than agents with lower monthly income. Consequently, agents with a high monthly income will buy on average more chips than agents with a low monthly income. Figure 5a) shows a bimodal realization of the distribution of the loss thresholds for all the agents. Figure 5b) - d) show the associated distribution for the poor, middle income and rich agents.

4.5 Following agents through the simulation

Figure 6 shows a sample time evolution of the attributes for two different agents. Their monthly income are 300$ and 500$, respectively. The figure illustrates some typical features:

- More affluent agents have a lower threshold than less affluent agents. Hence the former buy more often than the latter.
- The number of losses does not increase linearly with time. This reflects the fact that as a chip ages, the other gamers will upgrade and hence the old chip will accumulate losses faster.

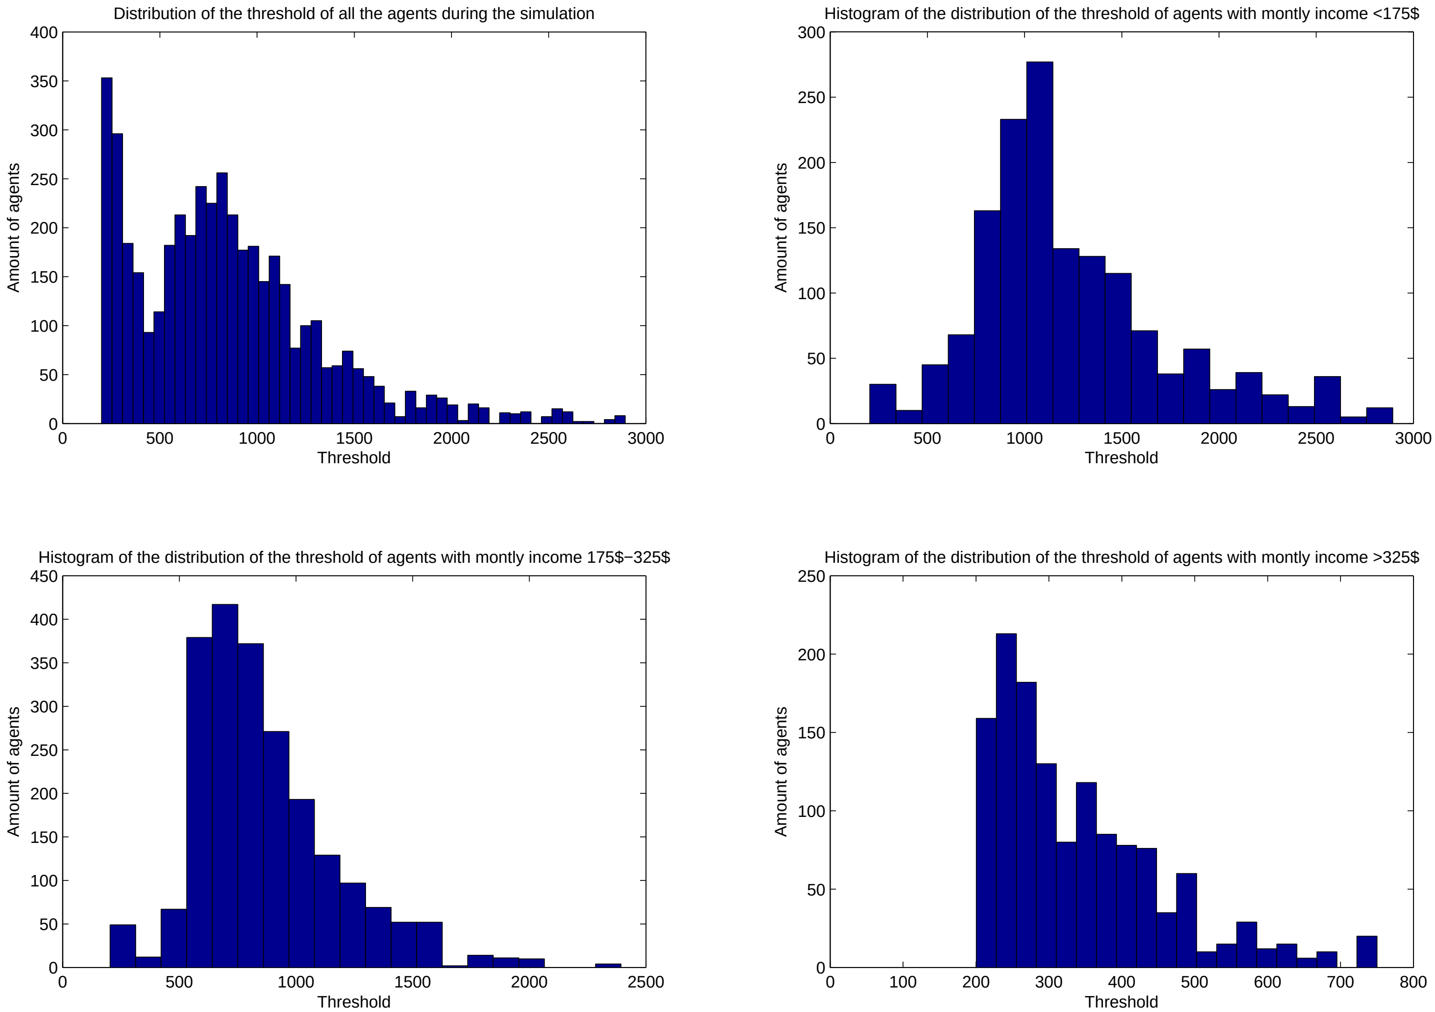

Figure 5: Typical histograms of the agents' loss thresholds. a) all agents, b-d) by monthly income bracket.

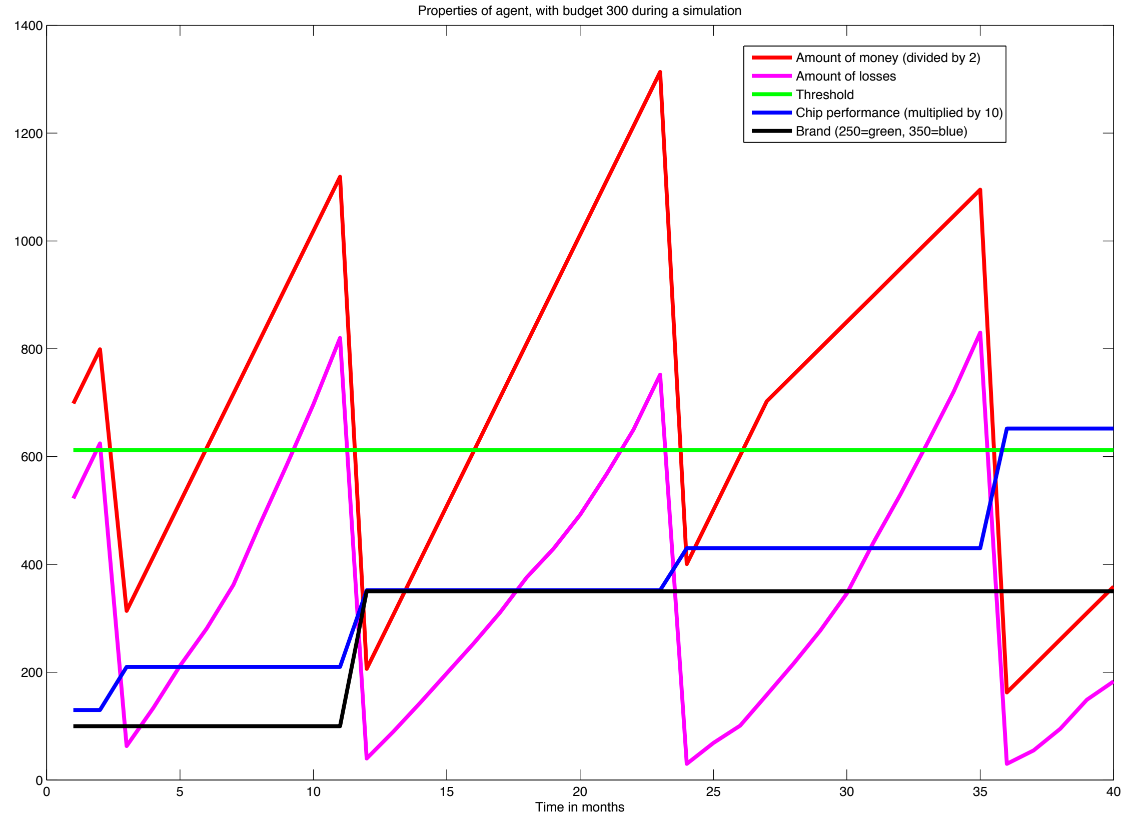

(a) Monthly income: $300

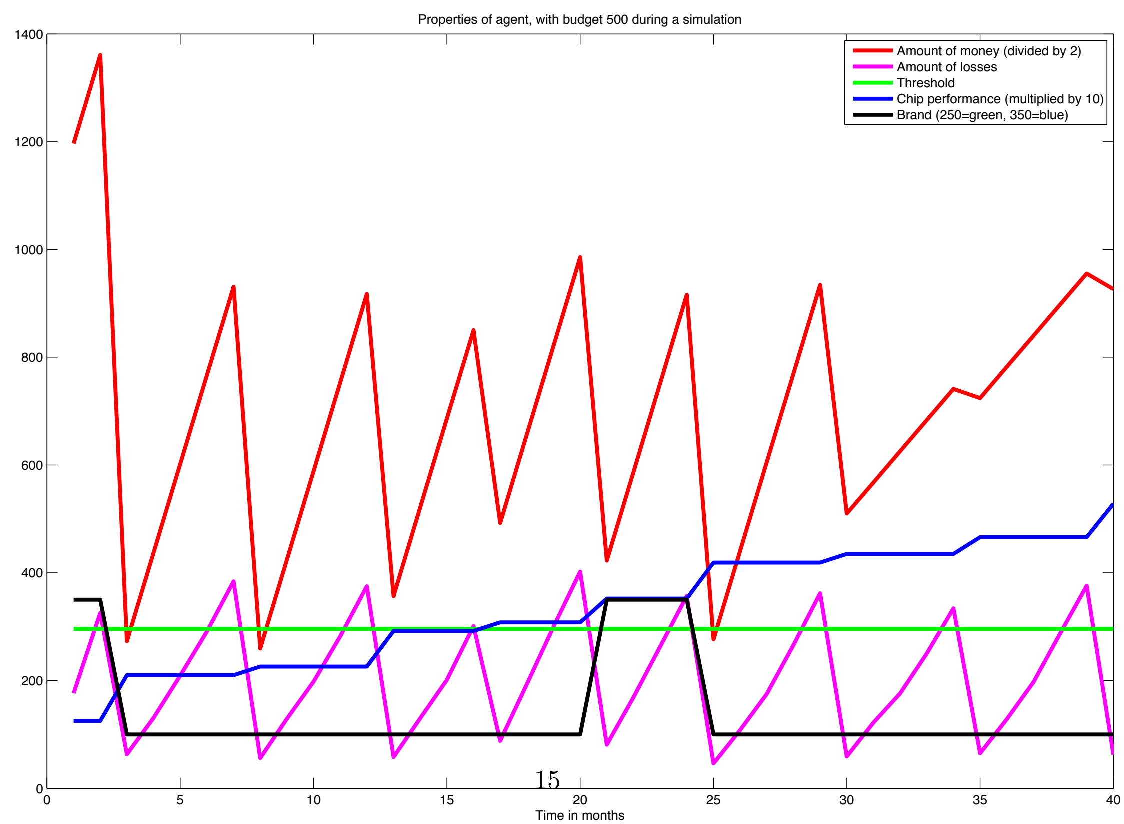

(b) Monthly income: $500

Figure 6: Simulated attributes of 2 different agents as a function of time.

- The recession may reduce the monthly income (Figure 6(a) following month 26).
- Agents (in particular more affluent ones) are able to switch between the brands, buying a new platform in the process. The agent in Figure 6(b) switches in months 2, 20 and 24.

4.6 Refinements

Basing ABS just on income and threshold parameters will give quite reasonable agreement between simulated sales and the data, measured by the mean square error of the monthly sales or of the cumulative sales over all the chips. However, the sales peaks that typically occur in the simulations in the first or second month after release tend to be significantly higher than those in the data, thus skewing the curves for monthly sales figures. We have therefore introduced additional secondary parameters modeling brand loyalty, the willingness of an agent to wait for an upcoming release of a new high performance chip and a conservative attribute leading certain agents to wait until a new chip has been proven in the market place. All those secondary parameters have the effect of reducing the initial peak but they had little influence on the global error measurements over the lifetime of the chip. Specific details of these secondary attributes are given in the Appendix.

The influence of the recession is another theme that can be modeled in more detail: It seems conceivable that due to the economical crisis some gamers drop out of the high-end market altogether and instead purchase regular chips. We model this effect by a decrease of the number of agents participating in the games. Furthermore the remaining potential buyers of high-end chips will also have less money to spend on average due to the economical crisis. Therefore we lowered the monthly budget and the total amount of money of a significant subgroup of the agents. This causes more agents to switch to the lower performing chips of the Green Team towards the end of the time interval under considerations.

5 Simulation output

5.1 Comparison of simulated sales and actual sales

Given the target sales data for the 19 chips of the Green and Blue team, we have run agent based simulations and registered the simulated sales of all the chips. Heuristic parameter optimization produced surprisingly good results. Subsequently, in the high dimensional space spanned by the characteristic parameters for the agents including the secondary parameters, we have tried to generate the best simulations by manually adjusting the parameters to reduce the mean error.

Figure 7(a) shows the total monthly sales for the Green and Blue Team. We have scaled the data such that the total actual sales over all chips in the entire time interval are equal to the total simulated sales. The bars represent the target sales data, the red line represents the average over the simulation results. The error bars indicate the 95% confidence interval for the simulations. Notice that we also adjust the total number of agents participating in the game over time to the total sales curve. Hence, it should not be too surprising that the simulation estimates the total monthly sales quite well. Nevertheless, Figure 7(b) shows the simulation performs very well in terms of replicating the relative proportions of blue and green sales over time.

Figure 8 shows the sales curves for every chip of the Blue Team (results for the Green Team look similar). Figure 9a shows a comparison between the average simulation result, the actual monthly sales and the 95% confidence interval for a typical chip (Platform 3, version 2, blue team). The figure presents the typical shape of most sales curves for the simulation data: The sales ramp up quickly, peak in the second month after the release and then decline gradually. In general, actual sales curves that follow this pattern are matched quite well by the simulation. Sales curves that do not follow this pattern typically matched poorly with the simulated sales. Figures 16 and 17 in the Appendix show the comparison of the sales curves of every chip. Often, an unusual sales pattern could be traced to some incidental sales circumstances that are special for this product.

From a business perspective, obtaining a good estimate of the cumulative sales over time for each chip is arguably more important than the monthly sales. Figure 9b shows the cumulative monthly sales for a chip (platform 2, version 2, Blue Team) together with a curve for the relative error. Since sales data in the beginning are small, the relative error is big but usually it converges rapidly to a value between -20% and + 20%. For most chips the cumulative sales of the simulation within the first few months is higher than the cumulative sales of the real data in the same period, i.e. the simulation overestimates the cumulative sales at the beginning of the lifetime of a chip. Figures 18 and 19 in the Appendix show the cumulative sales curves for all chips.

Figure 10 shows the relative error of the simulated cumulative sales of a chip at the time when the actual sales reach 25%, 50%, 75% and 100% of the total sales of that chip. At the 25% mark the simulation is overestimating the real sales most of the time while at 100% of the total real sales of the chip the distribution of the relative error is much narrower and almost symmetric.

5.2 The High-End Gamers

Having a population of agents whose buying behavior replicates the sales data for these chips reasonably well, we can now query the agent population for their properties and for emergent behavior. For instance (see Figure 11), we can look at the average budget and the average spending per agent, we can determine the average performance of the chip that the gamers own, we can determine the number of platforms that agents are buying (3.15 on average in 40 months) and the average number of chips (platforms and separate chips, 4.42 on average).

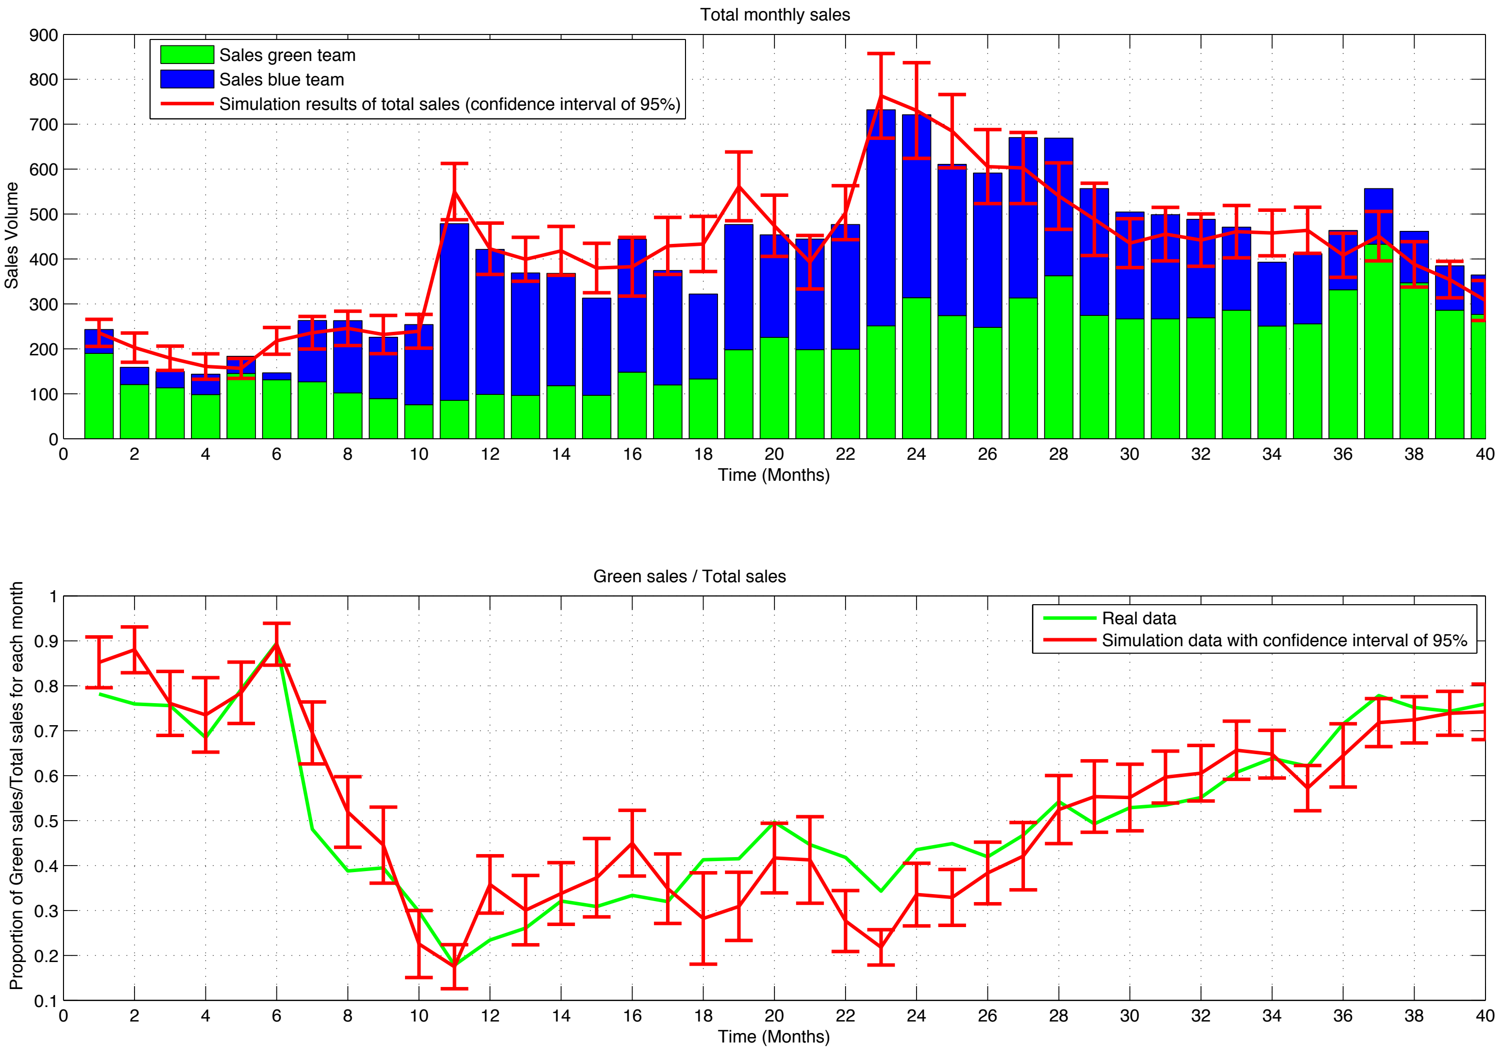

Figure 7: Figure at top: total monthly sales during the simulation and real data. Figure at bottom: green sales as a fraction of the total sales

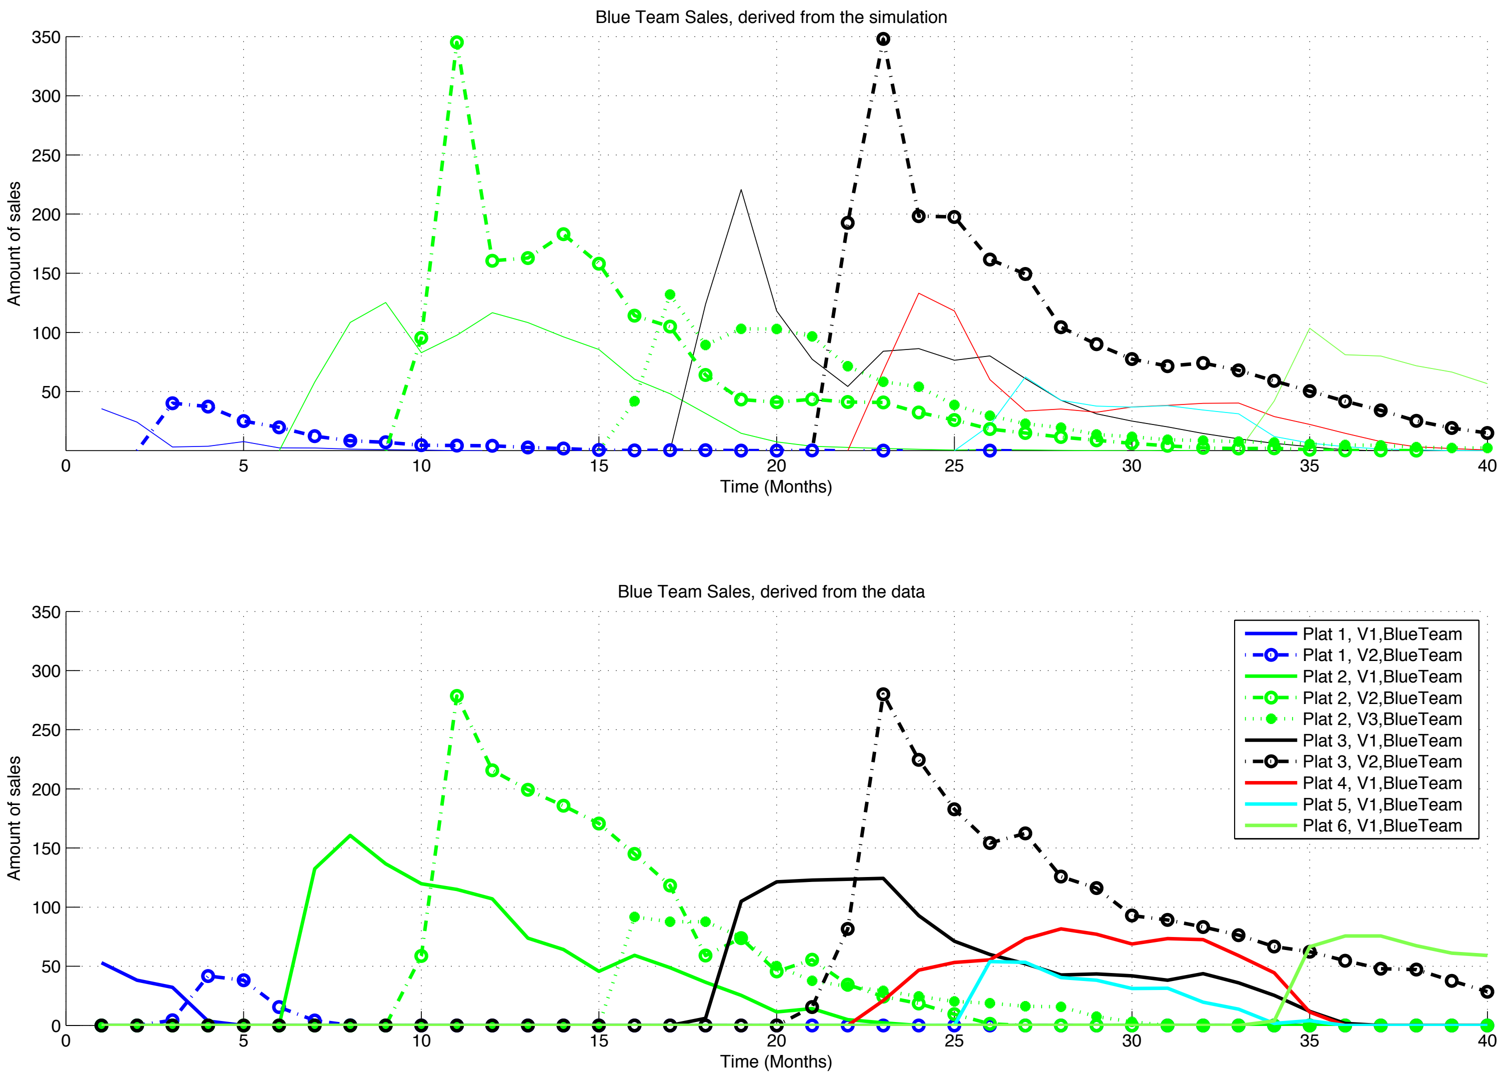

Figure 8: Sales for each chip of the Blue Team (top: simulation results, bottom: data).

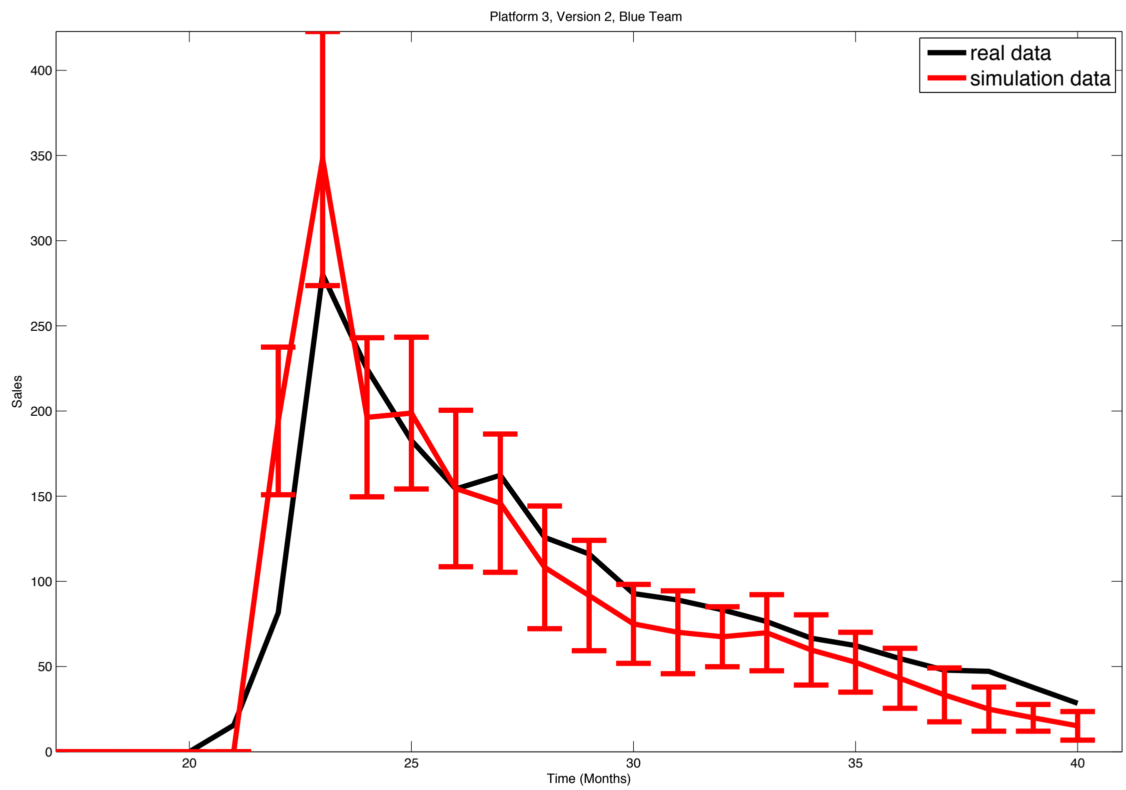

(a) Monthly sales

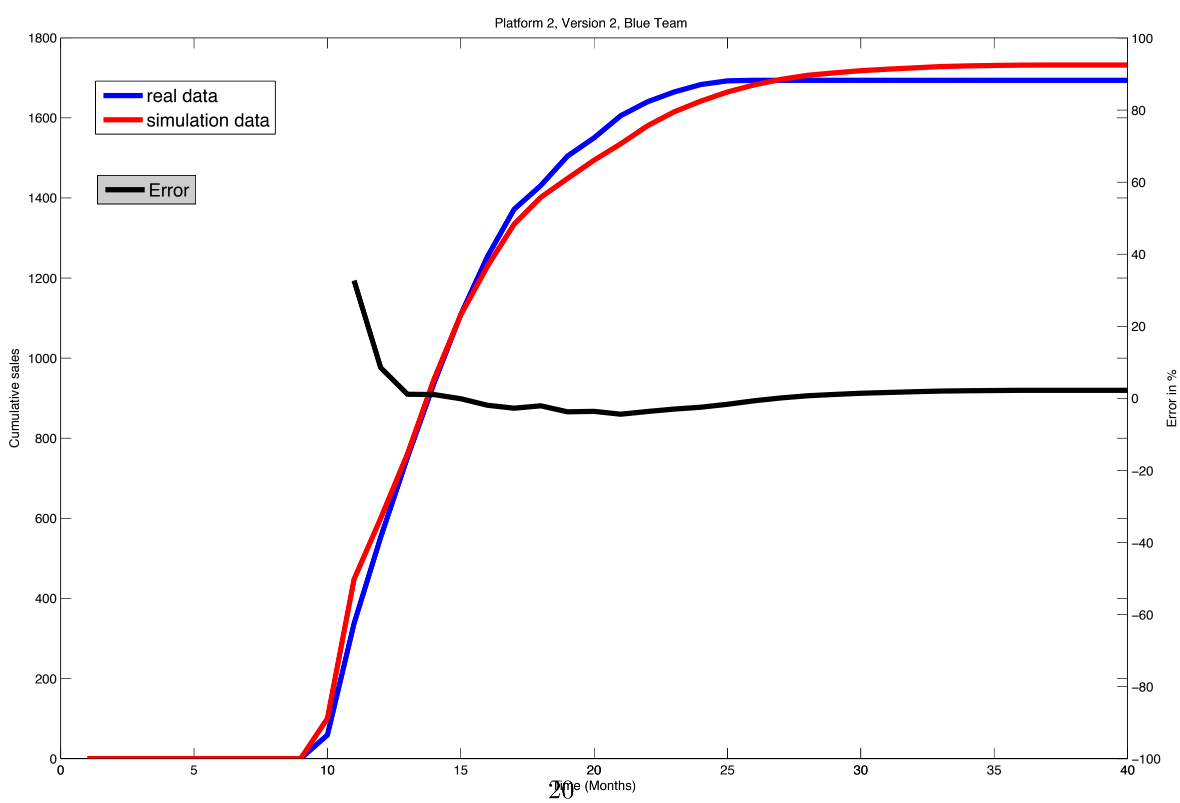

(b) Cumulative sales

Figure 9: a) Average and 95% confidence interval of the simulated sales for two typical chips compared to the actual sales, b) cumulative sales

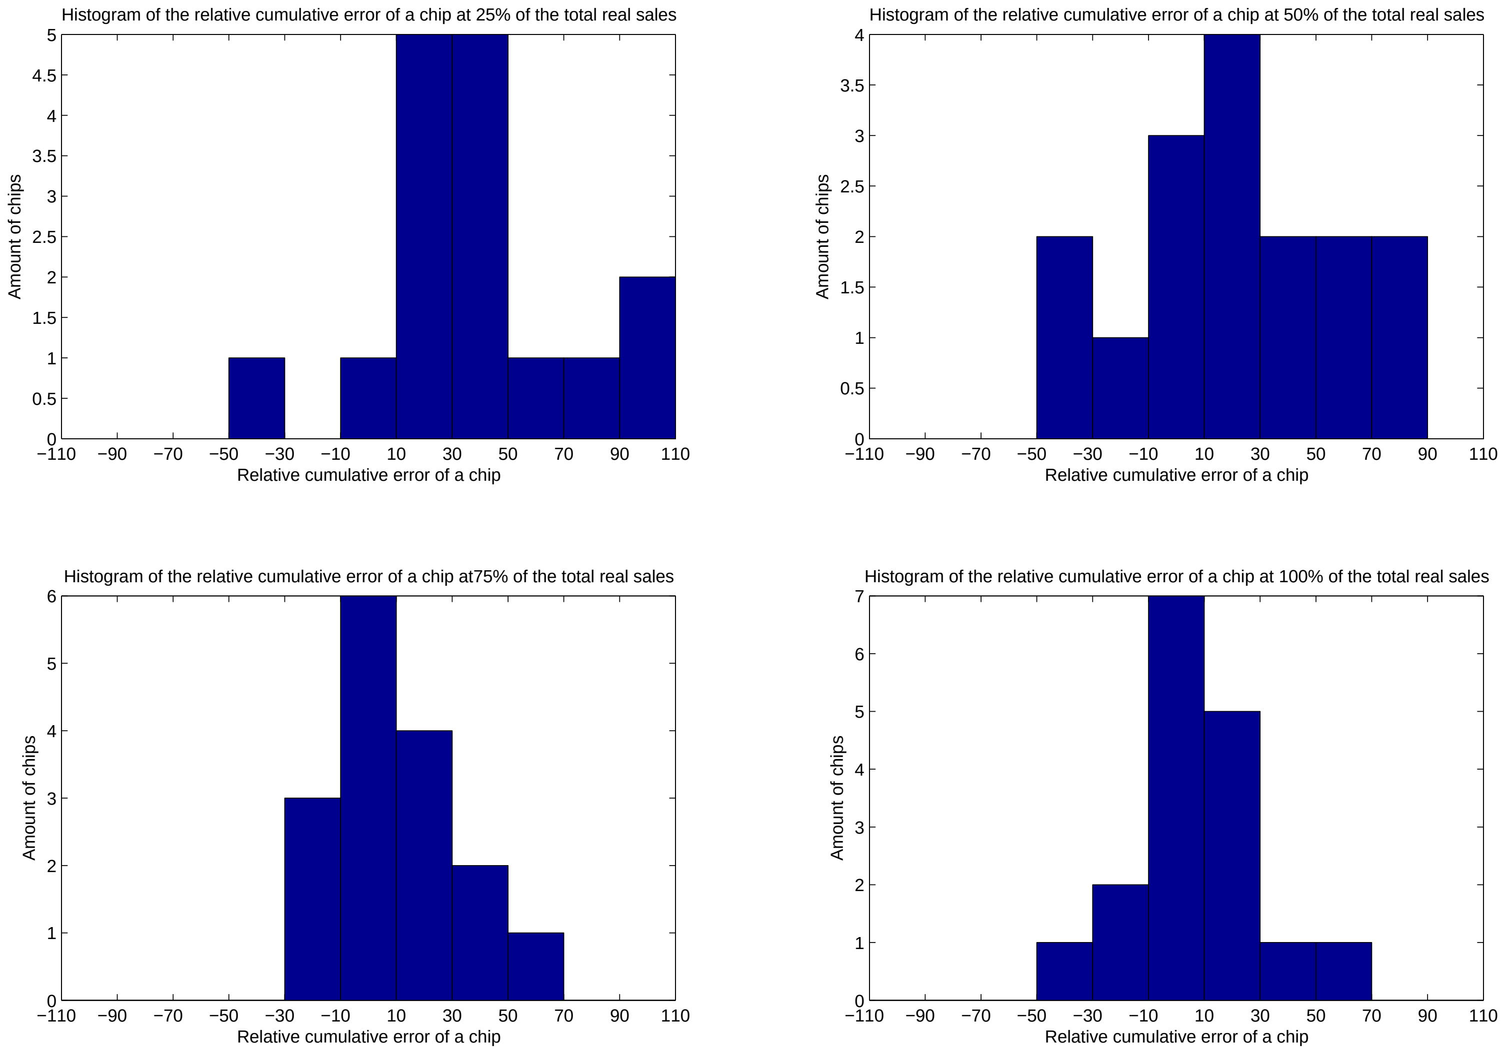

Figure 10: Histogram of the relative cumulative error of a chip at 25%, 50%, 75% and 100% of the total real sales of that chip.

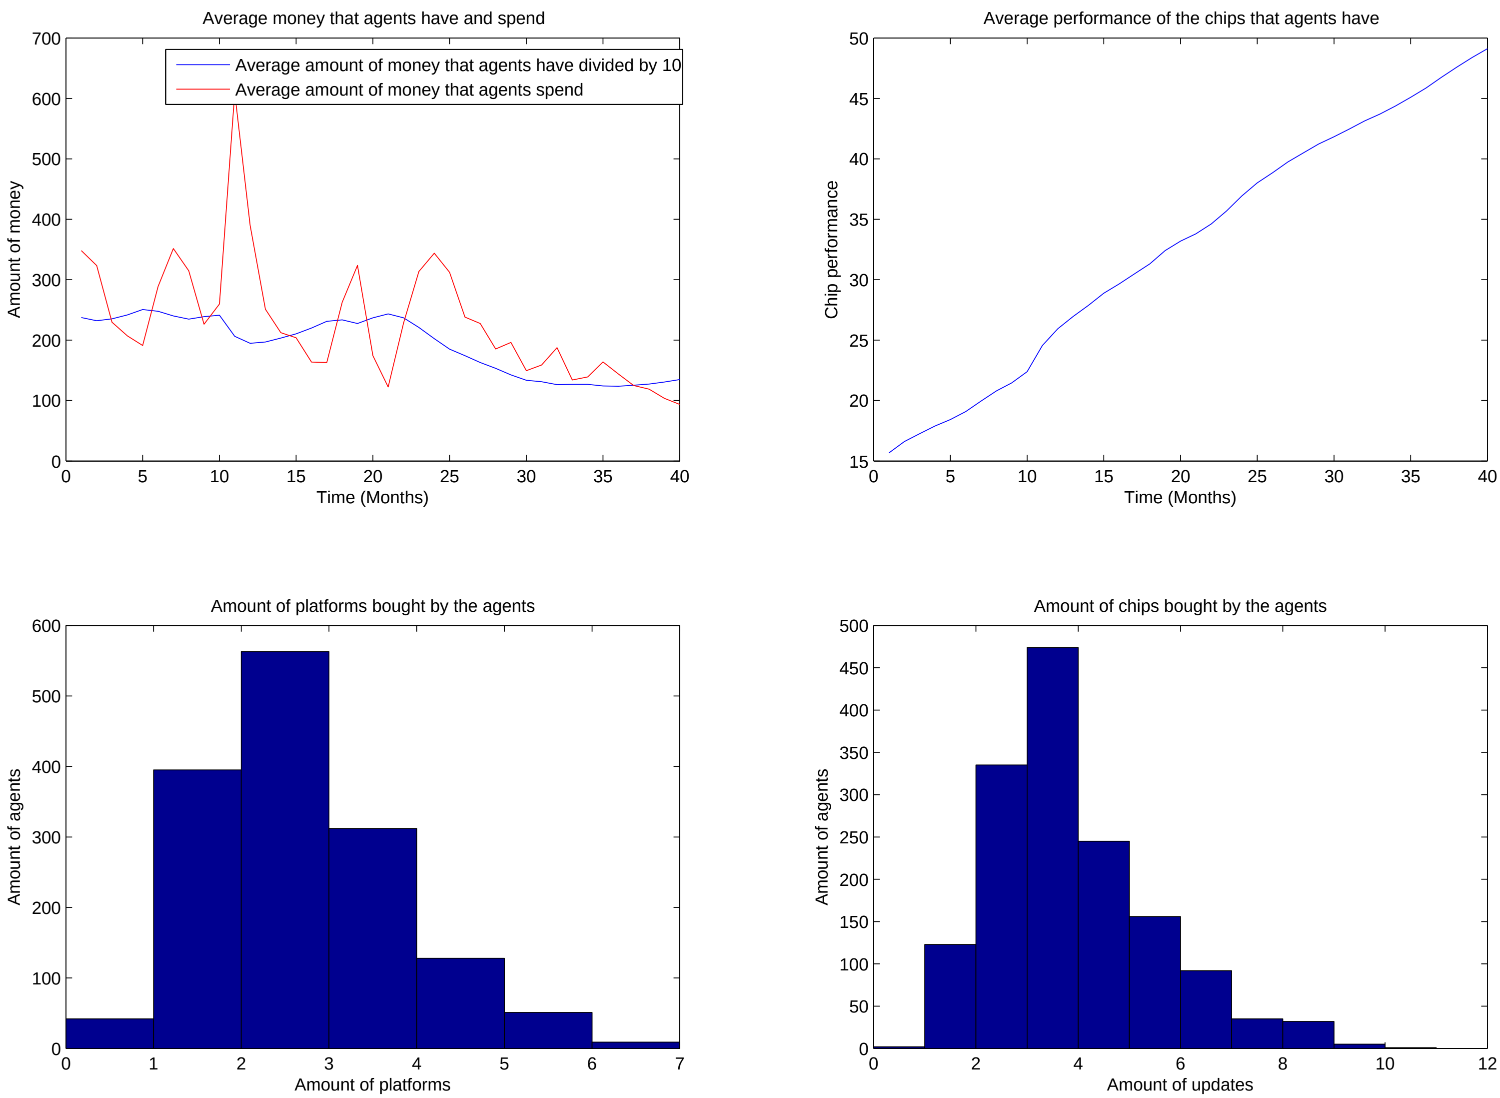

Figure 11: Top left: the average amount of money that agents have and spend during the simulation. Top right: the average performance of chips that agents possess during the simulation. Bottom left: the distribution of the number of platforms bought by the agents. Bottom right: the distribution of the number of chips bought by the agents.

5.3 Robustness of the simulation model

Although the agent simulation is conceptually simple, it still has too many design parameters to allow for a full optimization study. We therefore perform scans in parameter space, keeping all but one parameter fixed and determining the influence of the variation of one parameter at a time. In addition to being a multi-parameter problem, the optimization problem is a multi-objective one. We determine the influence of changing a parameter on eight optimization criteria: (i) the relative error of the total monthly sales (all chips); (ii) the average relative error of the total sales for each chip; (iii) the average relative error of the total sales for each platform; (iv) the average relative error of the sales curves for every chip; (v) the average relative error for the cumulative sales curve for every chip; (vi) the relative error at 25% of total sales; (vii) the relative error at 25% of total sales with a penalty. A penalty of factor 5 will be assigned to the relative errors when the simulation is underestimating the real sales data. (viii) a weighted error consisting of a weight of 50% for the relative error at 25% of total sales with the penalty factor 5 and of 50% for the relative error of the total monthly sales (all chips).

Figure 12 shows the robustness of the simulations relative to these measures against changing the average threshold of the agents up or down by up to 50%. The dotted vertical line in the figures indicates the value used for the scenario presented in section 5.1. We see that relative to most criteria, our best scenario is close to the smallest error measure. The only exception is the relative error at 25% of total sales which will decrease as the average threshold increases. However, if we penalize underestimation that effect goes away. This indicates the opportunity to use the loss threshold to influence the sales distribution during the initial months that the chip is available.

A very similar result can be seen in Figure 13 showing the robustness against changing the average monthly income of the agents up or down by 50%. Since the decision to buy is primarily triggered by the loss threshold, changing the monthly income does not influence the total monthly sales very much. However, it influences the choice of the computer chip and the ability to switch brands because switching brands forces the gamer to purchase an entire system instead of just a chip.

Similar studies can be done to calibrate the parameters characterizing the influence of the recession on the agent's behavior, the loyalty factor to a particular company and other secondary parameters. Overall, we find that relative to most criteria, the simulations are insensitive to those secondary parameters. This confirms our design assumption that performance (characterized by the loss-threshold) and monthly income are the two important parameters characterizing the High-End Gamers market.

Remark: If monthly income is highly important, then prices and the timings of a price decrease become very important too. This will be the subject of further studies.

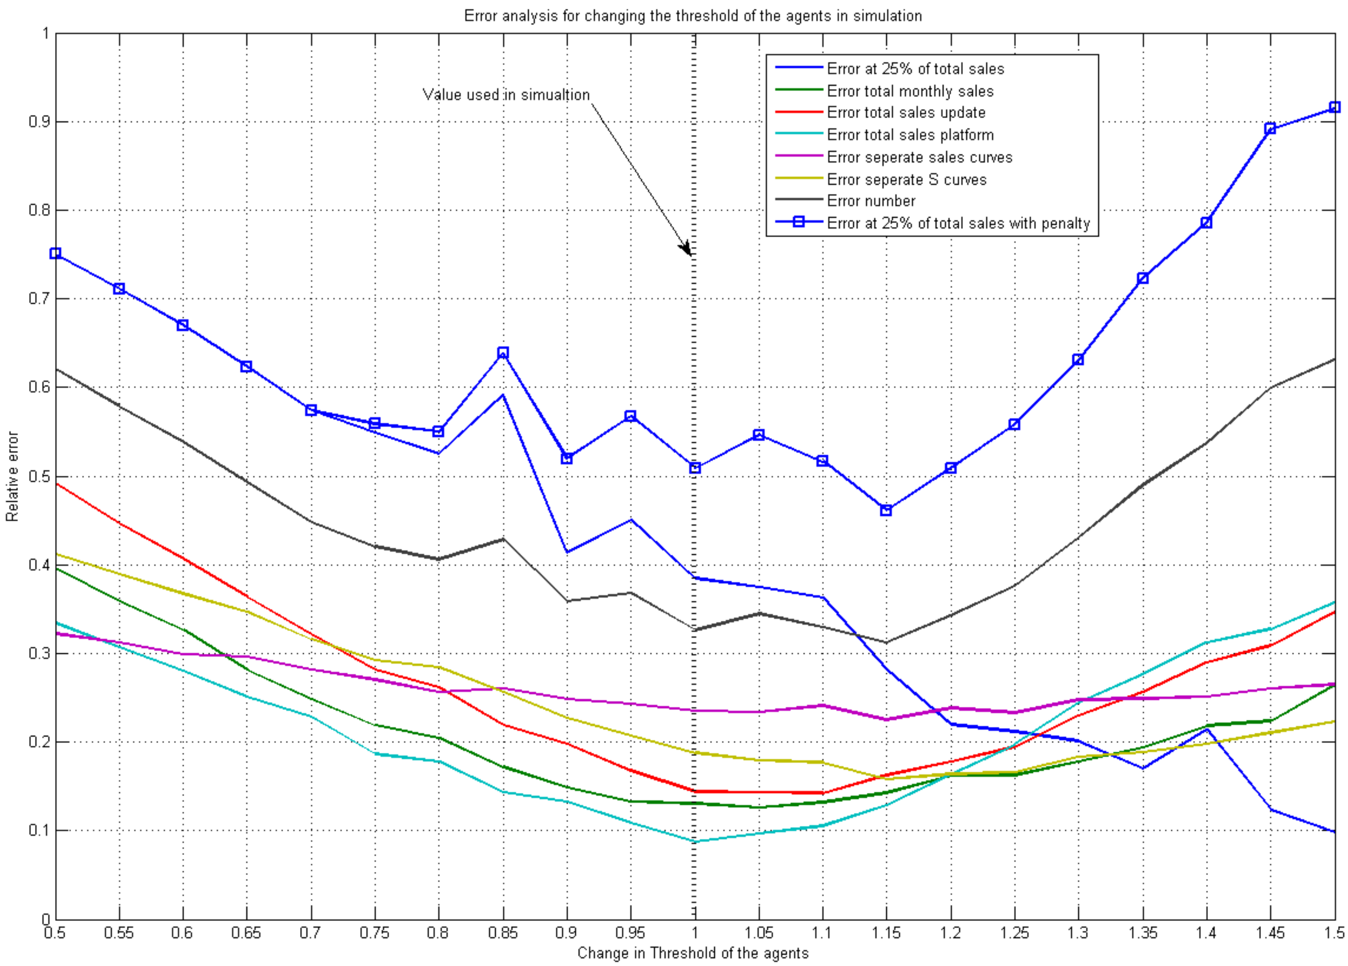

Figure 12: Error analysis for changing the average threshold of the agents.

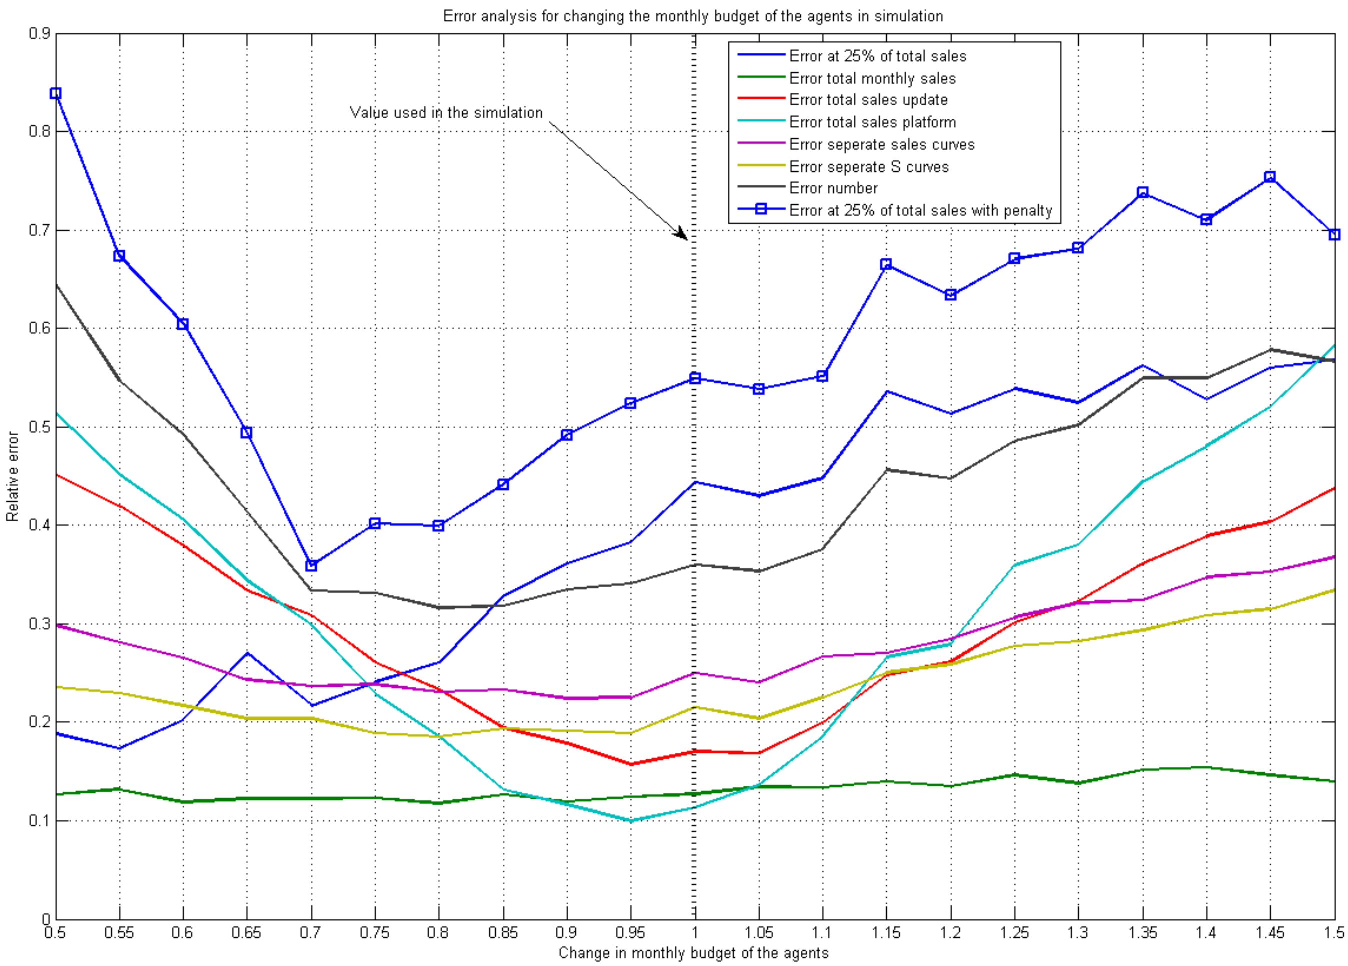

Figure 13: Error analysis for changing the average monthly income of the agents.

6 Conclusion

We have studied the competitive market for the computer chips used by High-End Gamers through an agent based simulation. We have developed an agent model that contains the main features of the market. We showed that the sales curves of 19 chips over about 3 years can be qualitatively and quantitatively approximated using two main parameters representing the agents' income allocated to the gaming hobby and the gamers inclination to improve their computer hardware. Secondary attributes like brand loyalty, anticipative waiting for new chips and other parameters that model the influence of the recession of the gamers' market are necessary to reduce the simulation error of the initial sales spike following the release of a new chip.

Since the data are very noisy and uncertain, the purpose of the ABS was not to create perfect matches between data and simulations. Our goal was to develop a tool that is easy to use for a sales and marketing department to study the dynamical evolution of a market under time dependent changes of its defining characteristics. An ABS model allows us to study *what if* questions dynamically and to explore and refine mental models that exist inside a company. ABS are not intended as a replacement for traditional regression or other studies of the market based on statistical inferences but as a complement. One major limitation of the ABS method is that we have no theoretical bases for the parameter estimates obtained through this method, i.e. we do not know whether they are unbiased or consistent. Neither do we have information regarding the variance and covariance of the parameters. However, the conceptual simplicity of the relationship between the agent attributes and individual behavior of market participants makes acceptance of an ABS model to management and sales force potentially easier than other, more mathematically sophisticated models.

We propose the current work as a prototype approach to modeling other markets. Following that notion we will pursue several future research directions:

- Release times of new chips currently follow roughly a three month period. It is of significant interest to study the market behavior, if the time between releases becomes longer. This raises the option to include the two producers of these chips into the simulation and model their competitive reaction to the market behavior and to each others' actions and hence to merge ABS with game theoretical models.

- The High-End Gamers market is almost insignificant relative to the overall market of computer chips. Hence, its model should be considered as a test run for modeling the much bigger markets of server, desktop, and laptop computers.

- Our immediate interest is to use the understanding of the High-End Gamers market gained from the ABS to study sales forecasting. In particular, we will use some parts of the existing data base as training set to generate the gamers' characteristic attributes and then try to forecast the complement of the data not used for training purposes. The interesting question here is whether the markets and hence the parametrization of the agent attributes are stable over time. If that is the case, then forecasting should be possible. Preliminary results suggest that the market for High-End Gamers is not stable over the whole time horizon but that nevertheless, forecasting is possible within three specific time periods. A detailed study comparing ABS based forecasting to traditional forecasting methods is in the works.

7 Acknowledgements

D.A. was supported by a grant from the Stiftung Volkswagenwerk under the program on Complex Networks and by NSF grant DMS-0604986. D.A and H.L gratefully acknowledge support from the INTEL Research Council. We thank Jay Schwartz (Intel Corporation) for providing us with the Green Team sales via Google searches.

References

[1] S. P. Anderson and A. de Palma. Multiproduct firms: A nested logit approach. *The Journal of Industrial Economics*, 40(3):261–276, 1992.

[2] F.M. Bass. A new product growth model for consumer durables. *Management Science*, 15:215–227, 1969.

[3] S. T. Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242–262, 1994.

[4] C. Bhat. A heteroscedastic extreme value model of intercity mode choice. *Transportation Research Part B: Methodological*, 29(6):471–483, 1995.

[5] G. Bitran, R. Caldentey, and R. Vial. Pricing policies for perishable products with demand substitution. *Working Paper*, 2010. New York University.

[6] E. Bonabeau. Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Science, USA*, 99(supp.3):7280–7287, 2002.

[7] S.A. Delre, W. Jager, T.H.A. Bijmolt, and M.A. Janssen. Targeting and timing promotional activities: An agent-based model for the takeoff of new products. *Journal of Business Research*, 60(8):826–835, 2007.

[8] J.A. Dubin. A nested logit model of space and water heat system choice. *Marketing Science*, 5(2):112–124, 1986.

[9] M.E. Epstein and R. Axtell. *Growing artificial societies: Social Science from the ground up.* MIT Press, Boston,MA, 1996.

[10] Jeremy Ginsberg, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457, 2009.

[11] R. Goettler and B. Gordon. Competition and Innovation in the Microprocessor: Does AMD spur Intel to innovate more? *Working Paper, SSRN: http://ssrn.com/abstract=1084424.*, 2009.

[12] P. K. Goldberg. Product differentiation and oligopoly in international markets: The case of the u.s. automobile industry. *Econometrica*, 63(4):891–951, 1995.

[13] J. Goldenberg, O. Lowengart, and Shapira D. Zooming in: Self-emergence of movements in new product growth. *Marketing Science*, 28(2):274– 292, 2009.

[14] W. Hanson and K. Martin. Optimizing multinomial logit profit functions. *Management Science*, 42(7):992–1003, 1996.

[15] P.K. Kannan and G. P. Wright. Modeling and testing structured markets: A nested logit approach. *Marketing Science*, 10(1):58–82, 1991.

[16] R. D. Luce. *Individual Choice behavior: A theoretical analysis.* Wiley, New York, 1959.

[17] Charles M. Macal and Michael J. North. Tutorial on agent-based modeling and simulation. M.E. Kuhl (ed.) et al., Proceedings of the 2005 Winter Simulation Conference., 2005.

[18] D. McFadden. Econometric models for probabilistic choice among products. *Journal of Business*, 53(3):S13–S29, 1980.

[19] D. McFadden. The choice theory approach to market research. *Marketing Science*, 5(4):275–297, 1986.

[20] D. McFadden and K. Train. Mixed MNL models of discrete response. *Journal of Applied Econometrics*, 15:447–470, 2002.

[21] D. L. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, editor, *Frontiers in Econometrics.* Academic Press, 1974.

[22] N. Mead and T. Islam. Technological forecasting – model selection, model stability, and combining models. *Management Science*, 44(8):1115–1130, 1998.

[23] M.J. North, C.M. Macal, J. St. Aubin, P. Thimmapuram, M. Bragen, J. Hahn, J. Karr, N. Brigham, M.E. Lacy, and D. Hampton. Multiscale agent-based consumer market modeling. *Complexity*, 15(5):37–47, 2010.

[24] H.Van Dyke Parunak, Robert Savit, and Rick L. Riolo. Agent-based modeling vs. equation-based modeling: A case study and users' guide. J.S. Sichman et. al. (Eds): MABS'98, 1998.

[25] I. Praca, C. Ramos, Z. Vale, and M. Cordeiro. Mascem: A multiagent system that simulates competitive electricity markets. *Intelligent Systems*, 18(6):54–60, 2003.

[26] Hazhir Rahmandad and John Sterman. Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. *Management Science*, 54(5):998–1014, 2008.

[27] M. Saguan, N. Keseric, P. Dessante, and J-M. Glachant. Market power in power markets: Game theory vs. agent-based approach. IEEE PES Transmission and Distribution Conference & Exposition, Caracas, Venezuela, August 2006.

[28] L. Tesfatsion. Agent-based computational economics: growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.

[29] K. E Train, D. L. McFadden, and M. Ben-Akiva. The demand for local telephone service: A fully discrete model of residential calling patterns and service choices. *RAND Journal of Economics*, 18(1):109–123, 1987.

[30] S.D. Wu, K. Kempf, S. Shirodkar, A. Mishra, M. Atan, and B. Aytac. Extending Bass for Improved Forecasting at Intel. *Interfaces*, to be published, 2010.

A Prices

Figure 14 and Figure 15 show the price history of all chips and all platforms, respectively, over the 40 month time horizon.

B Secondary attributes and parameters for the agent simulations

Recession. There is considerable disagreement over the start of the recession for the semiconductor industry market. We determine the start of the recession for a particular agent by randomly picking a month between month 21 and 30 of the simulation at which one of the following actions may be taken:

- The agent leaves the game completely.

For the agents that remain in the game

- The monthly income is reduced to between 20% and 80% (uniformly distributed) of an agent's current monthly income.
- The total accumulated budget of an agent is reduced. For 50% of the agents their budget is lowered by 50%. For the remaining agents with budgets of more than 5500$ their budget is adjusted to between 5000$ and 6000$ (uniformly distributed).

Loyalty. Initially each agent is assigned a *loyalty factor*. The *loyalty factor* is uniformly distributed (between 0.0 and 0.4). It is used at the moment when an agent decides to buy and compares the performance of the chips on the market. It lowers the performance of chips from the company not currently used by the agent. The higher the *loyalty factor*, the more loyal the agent is to its (current) company.

Conservative attribute. This attribute models the reluctance of an agent to buy a chip that is brand new on the market. This waiting period between the release of a new chip and the moment an agent might buy this new chip is determined by a gamma distribution ($\alpha = 2$,$\beta = 1.25$). Not all the agents have the conservative attribute: 50% of the agents with a minimal monthly income of 325$ do not have the conservative attribute. Therefore affluent agents on average are able to buy a new chip or platform sooner than non-affluent agents.

Waiting. There are two decisions involved in postponing a purchase, whether to wait for a particular chip and how long to wait for it. The decision to wait for a chip is made by comparing a waiting factor allocated to an agent to a factor allocated to a new chip. Both factors are randomly assigned to an agent or to the chip from a uniform distribution between 0 and 10. If the waiting factor of the agent is higher than the waiting factor of the

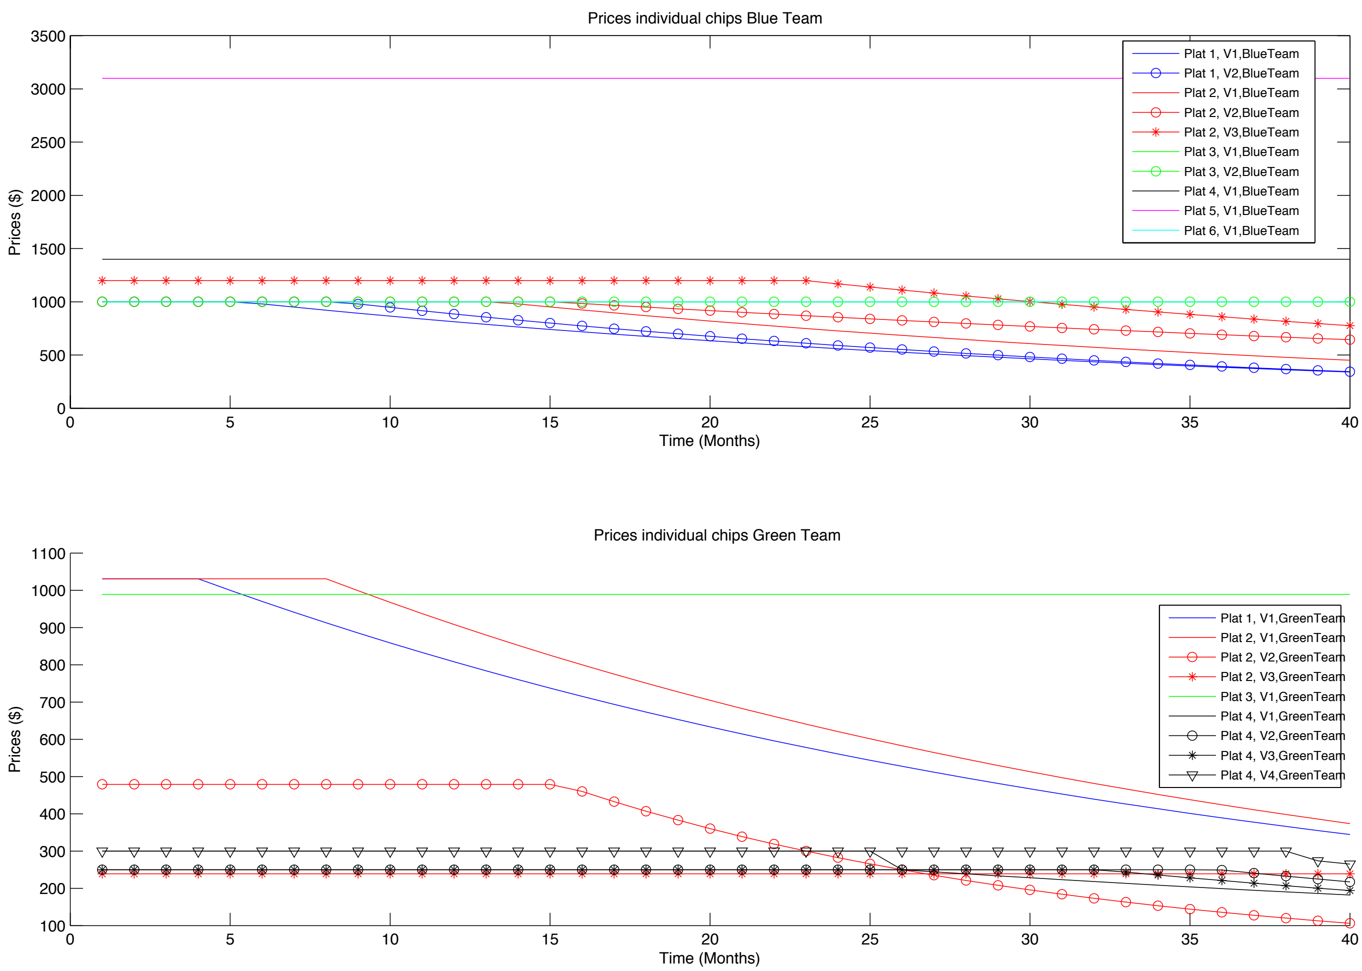

Figure 14: The prices of the Blue and Green Team high-end chips (period: January 2006 till April 2009)

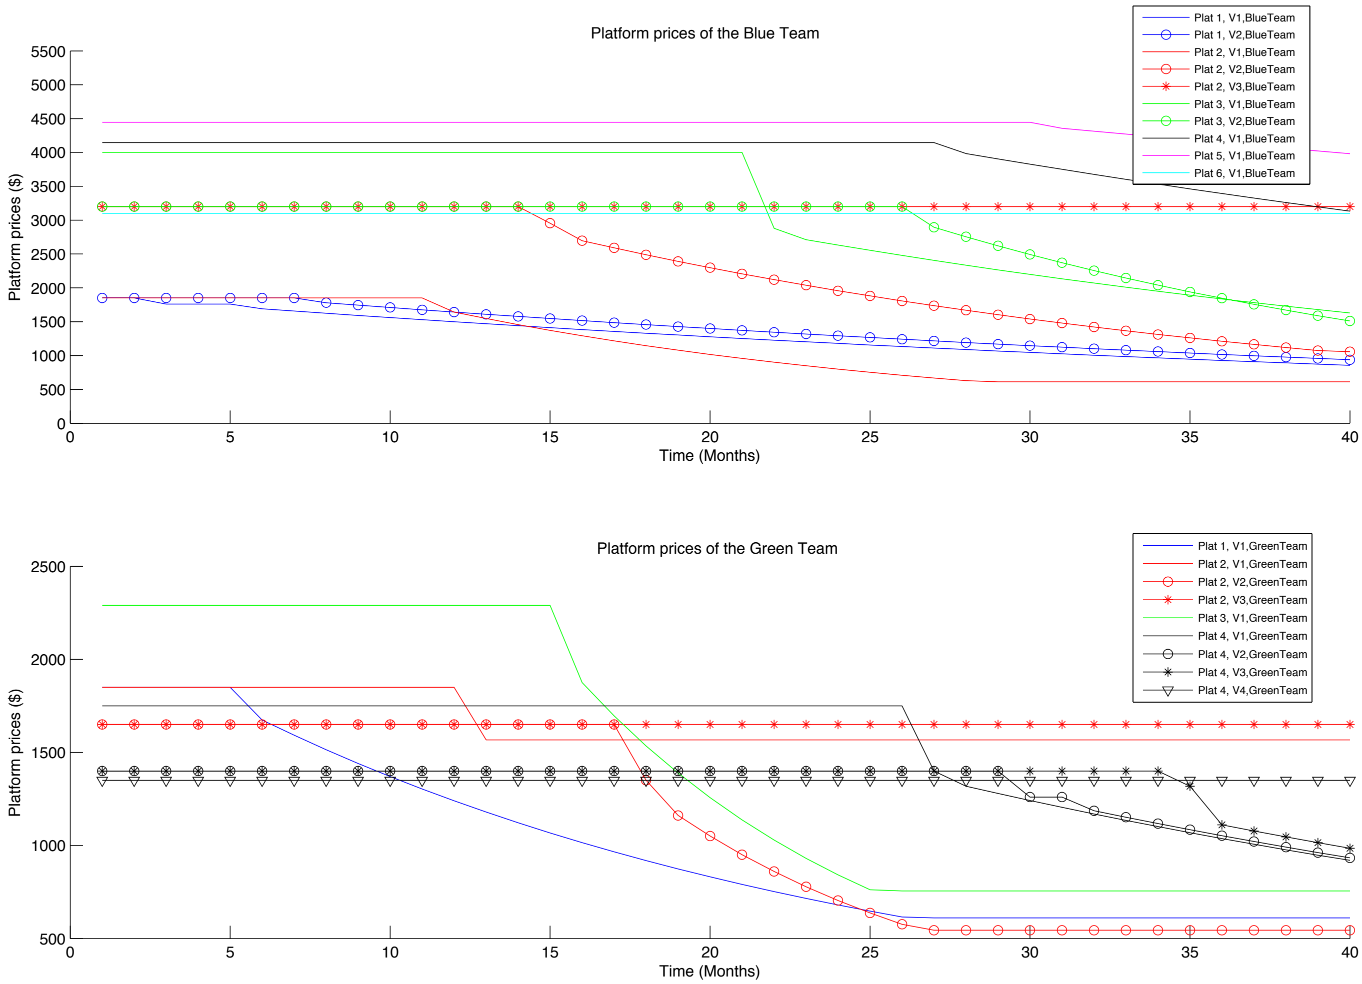

Figure 15: The prices of the Blue and Green Team platforms (period: January 2006 till April 2009

chip, the agent will wait to buy this new chip. Agents with a conservative attribute higher than 5, will never wait. How long an agent will wait is determined by a waiting window that is uniformly distributed between 0 and 4 months.

Initialization of the ABS. At the start of the time horizon, there already existed High-End Gaming and a High-End Gamers market. Hence we must have a starting population that is consistent with the sales in the first few months. Initial conditions are set for the budget, the initial number of losses and the current chip and platform. The initial budget depends on the monthly income of the agent and is normally distributed (μ =monthly income, $\sigma = 2\times$ monthly income). The initial number of losses is determined by the number of losses near the end of the previous simulation. The initial chip and platform are assigned to an agent to generate the best fit to the sales data of the first few chips.

C Detailed sales data

The following figures show the average simulations and 95% confidence intervals for the monthly sales and the cumulative sales of each individual chip. We dropped the simulations for the first two chips of the Blue team since their volume was very small.

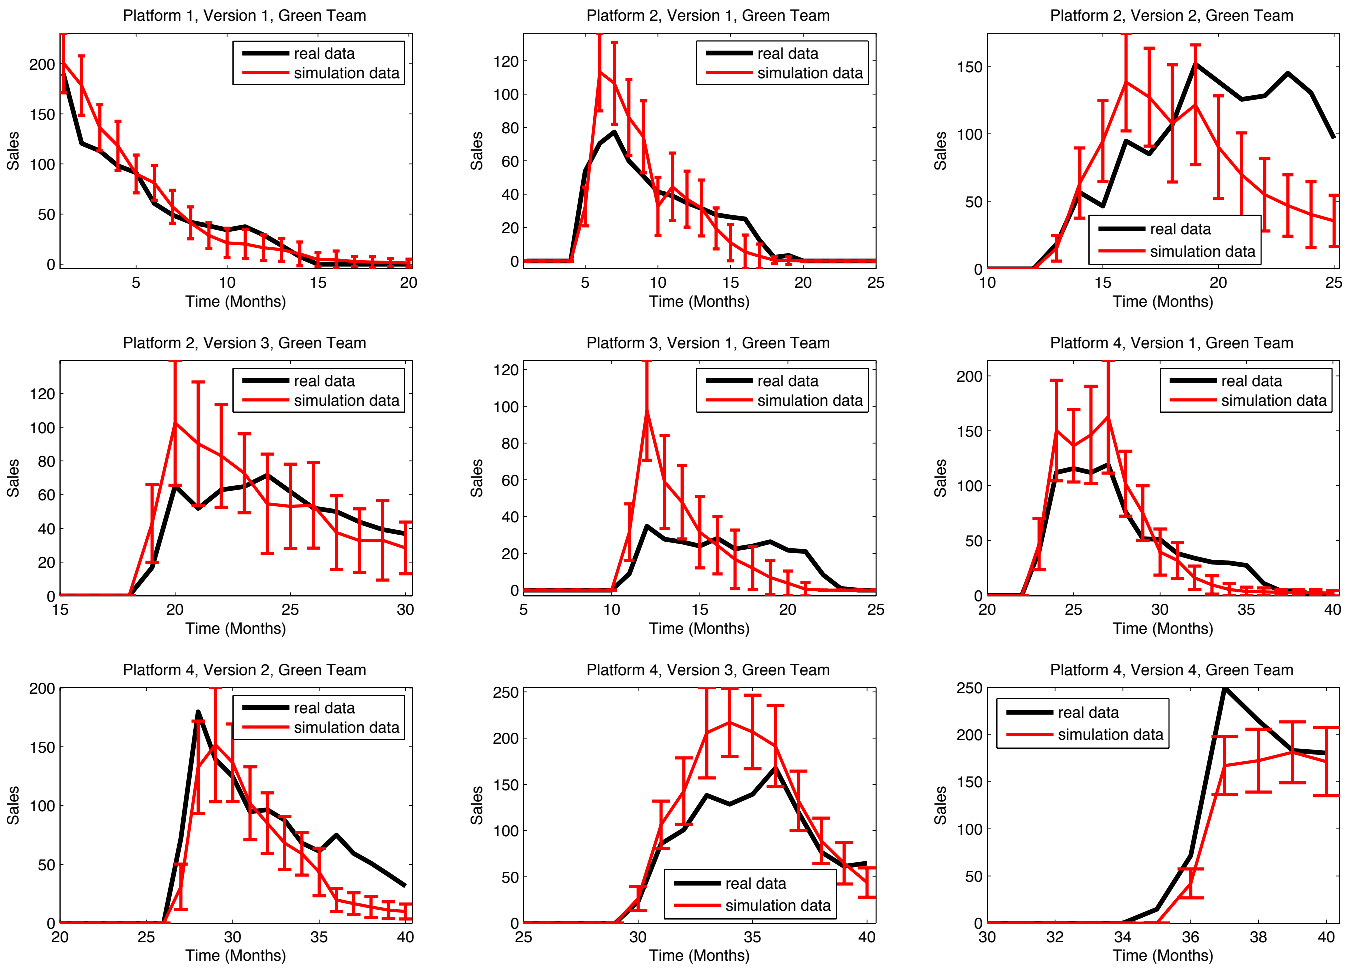

Figure 16: Average sales and 95% confidence interval for each chip of the Green Team during the simulation.

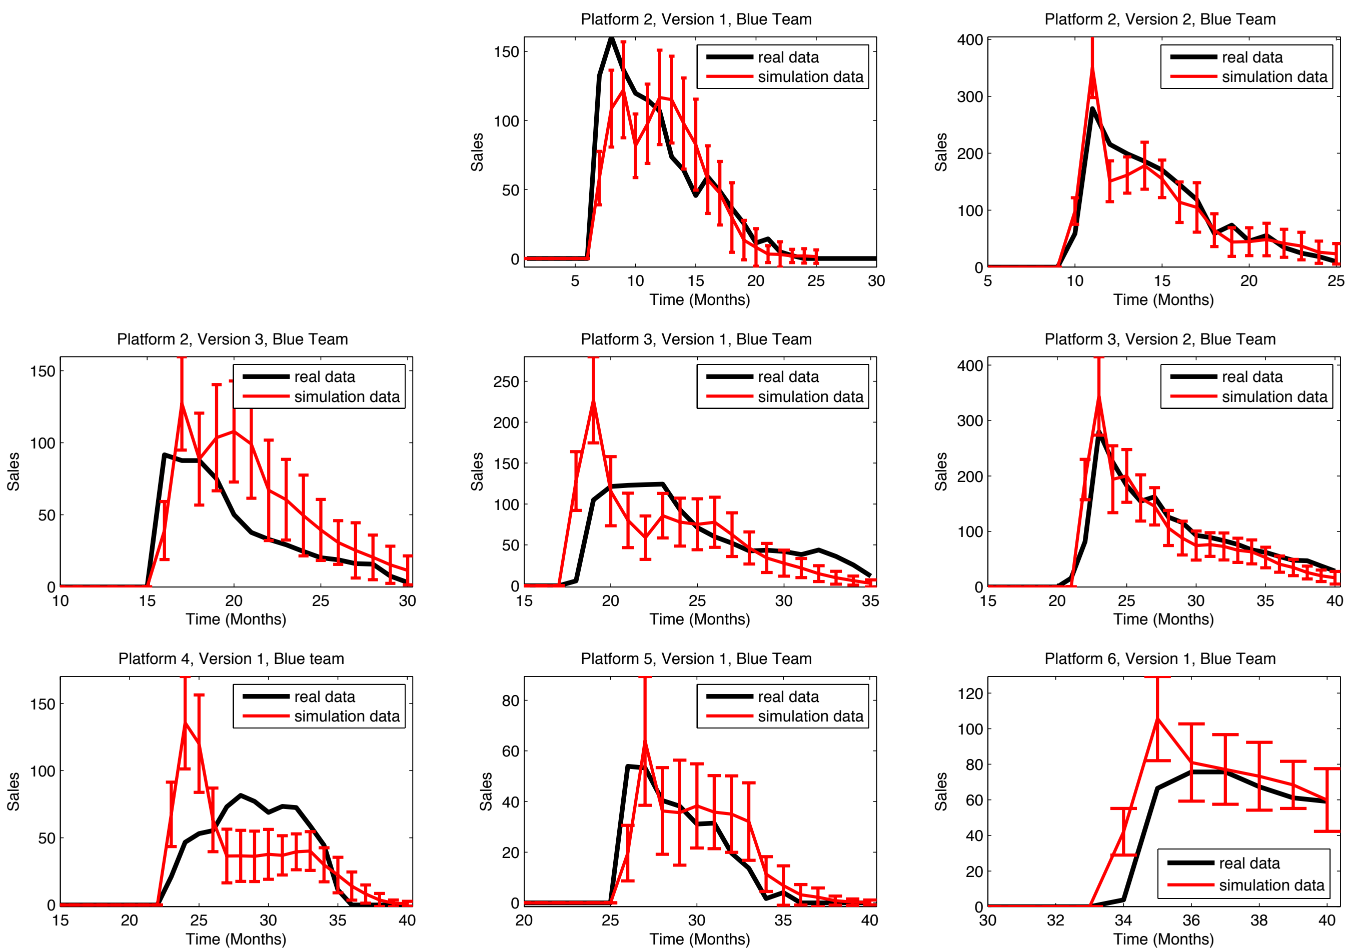

Figure 17: Average sales and 95% confidence interval for each chip of the Blue Team during the simulation.

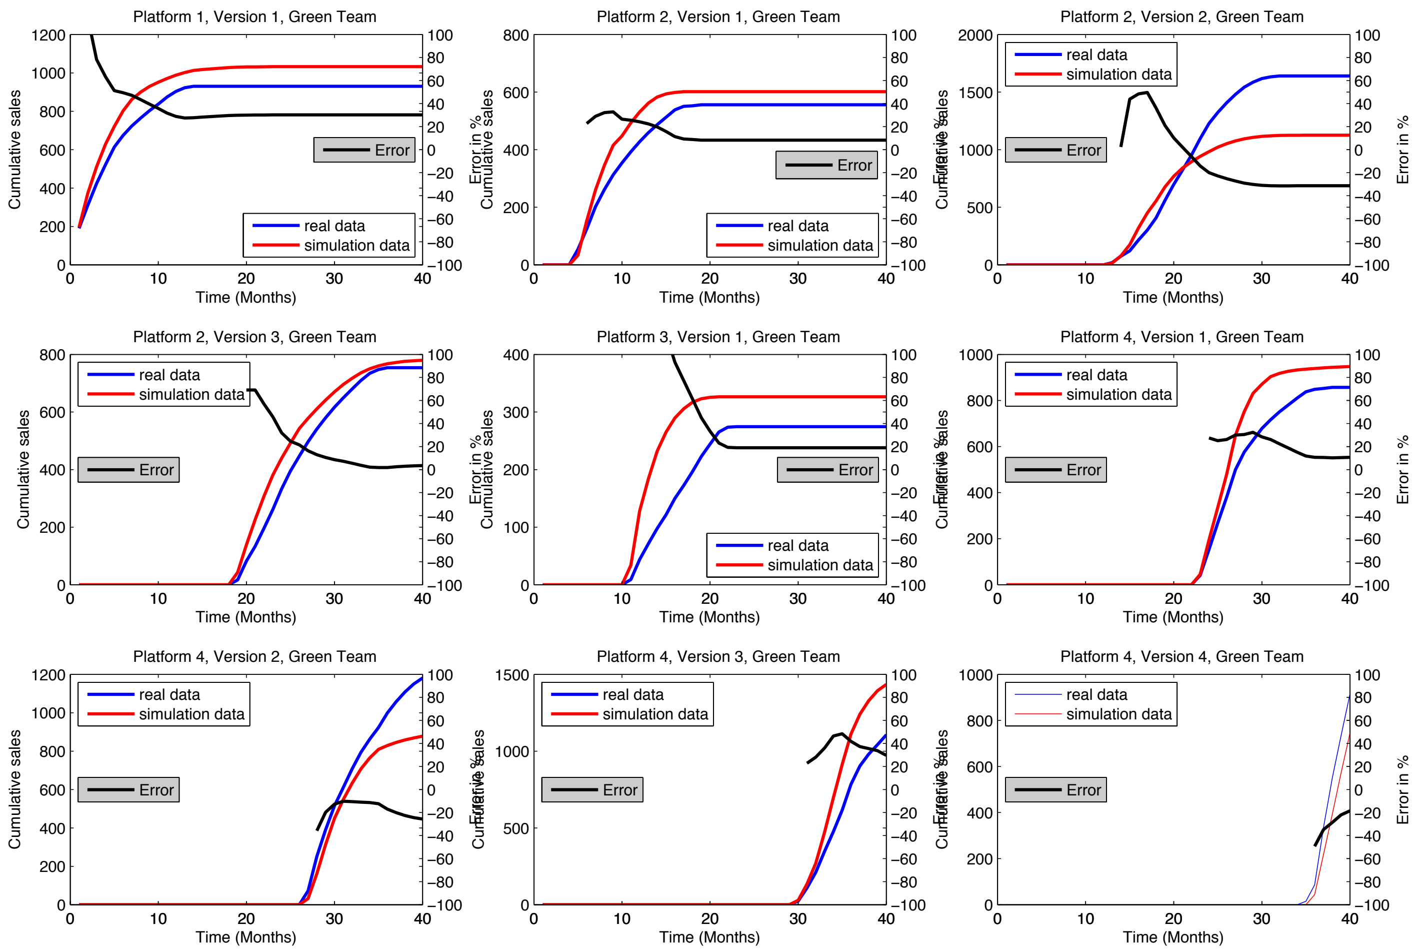

Figure 18: Cumulative monthly sales for each chip of the Green Team, with relative error.

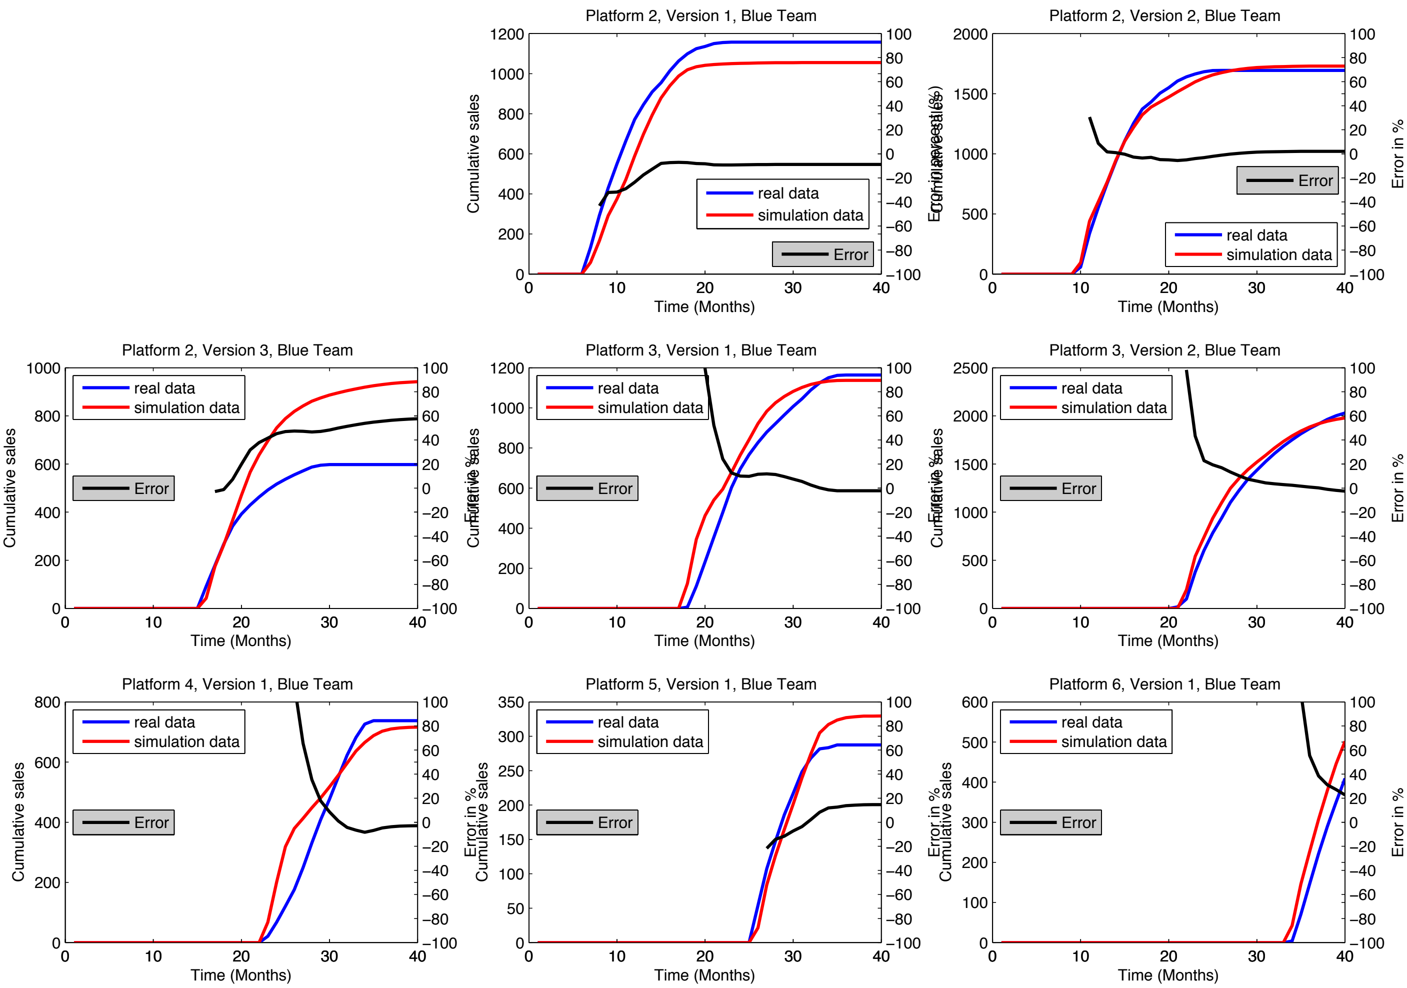

Figure 19: Cumulative monthly sales for each chip of the Blue Team, with relative error.