

Implicit Control Mechanisms and Organizational Learning*

Dennis Campbell
Harvard Business School
dcampbell@hbs.edu

Marc Epstein
Rice University
epstein@rice.edu

Asis Martinez-Jerez
Harvard Business School
asis@hbs.edu

Current Draft: July 2009

ABSTRACT: We link two elements of organizational design, the delegation of decision-rights and the provision of incentives, to the behavior and human capital formation (e.g. learning) of lower-level employees. We exploit a field-research setting in which employees' use of decision rights can be directly measured and traced to performance outcomes. We document implicit incentives in the form of sharp increases in employee termination linked to "excessive" use of decision-rights. We find that these implicit incentives are strongest for employees in business units which monitor the use of decision-rights most closely (e.g. "tight monitoring"). We document two broad behavioral effects for employees that are consistent with implicit incentives in tighter monitoring environments: (1) they are less likely to use discretion in exercising decision-rights; (2) they are less likely to adjust for local information, including historical performance data, in their decision-making. Finally, we document strong learning effects among lower level employees in our setting. However, these learning effects are concentrated among employees in "loosely monitored" business units and almost entirely absent in "tightly monitored" business units. The results are consistent with an experimentation hypothesis in which tight monitoring of decision-rights leads to fewer deviations from prescribed decision rules and less learning.

*We thank Srikant Datar, Ranjani Krishnan, Michal Matejka, Greg Miller, Krishna Palepu, Joe Weber, and participants at the AAA Annual Meeting, Harvard Junior Faculty Conference, the Management Accounting Section Conference, Real Colegio Complutense at Harvard, the Tilburg University Accounting Summer Camp, Universidad Complutense, and the University of Texas for helpful comments. All errors remain our own.

Implicit Control Mechanisms and Organizational Learning

ABSTRACT:

We link two elements of organizational design, the delegation of decision-rights and the provision of incentives, to the behavior and human capital formation (e.g. learning) of lower-level employees. We exploit a field-research setting in which employees' use of decision rights can be directly measured and traced to performance outcomes. We document implicit incentives in the form of sharp increases in employee exits linked to "excessive" use of decision-rights. We find that these implicit incentives are strongest for employees in business units which monitor the use of decision-rights most closely (e.g. "tight monitoring"). We document two broad behavioral effects for employees that are consistent with implicit incentives in tighter monitoring environments: (1) they are less likely to use discretion in exercising decision-rights; (2) they are less likely to adjust for local information, including historical performance data, in their decision-making. Finally, we document strong learning effects among lower level employees in our setting. However, these learning effects are concentrated among employees in "loosely monitored" business units and almost entirely absent in "tightly monitored" business units. The results are consistent with an experimentation hypothesis in which tight monitoring of decision-rights leads to fewer deviations from prescribed decision rules and less learning.

Implicit Control Mechanisms and Organizational Learning

I. Introduction

In this paper, we link two elements of organizational design, the delegation of decision-rights and the provision of incentives, to the behavior and human capital formation (i.e. “learning”) of lower-level employees. A considerable body of accounting research has recognized the delegation of decision-rights and the provision of incentives as two of the most important elements of organizational design (e.g. Baiman et al. 1995; Bushman et. al. 1995; Brickley et. al. 1996; Christie et. al. 2001; Nagar 2002; Abernethy et al. 2004; Campbell et. al. 2009). This literature has largely focused on the joint nature of delegation and incentive compensation choices and has identified several broad determinants of these choices. Considerably less attention has been paid in this literature to how these choices, once determined, influence the behavior and learning of employees – particularly those at lower levels in the organization.

Determining appropriate incentives and decision-rights for lower level employees is an important issue for most organizations. These employees often have unique knowledge about an organization’s individual customers, its local markets, and/or its production processes. For example, most distributed retail and service firms (e.g. department stores, banks) operate across numerous local markets with heterogeneous customer, labor supply, and demand characteristics. In these settings, the relative expertise of headquarters is likely to be diminished vis-à-vis local employees in controlling local operations (Baiman et al. 1995; Campbell et. al. 2009; Casas-Arce et. al. 2009).

The standard prediction from agency theory in these settings is that the firm would delegate decision-rights to local managers to gain the benefits from effective use of local

information without incurring the costs of collecting and transmitting all of this information to top management (Jensen and Meckling 1992; Brickley et. al. 1996; Nagar 2002; Prendergast 2000). At the same time, the firm would provide incentives to use these decision-rights effectively. Incentive alignment can potentially be accomplished through the use of “outcome” measures closely aligned with the firm’s performance or via “input” measures that are more closely aligned with the local employee’s actions (Merchant 1985; Hwang et. al. 2009).

Both of these choices are potentially problematic for motivating the effective use of decision-rights for lower level employees. Incentives based on “outcome” measures provide the strongest motivation for acquiring and using local information effectively. However, particularly for lower level employees, such incentives can impose a great deal of uncontrollable risk for which the employee must be compensated if risk-averse or wealth-constrained. Incentives based on “input” measures can minimize this risk tradeoff, but do little to motivate the use of local information by the employee.

An alternative mechanism for motivating the effective use of decision-rights that appears to be widespread in organizations is direct monitoring of employee use of decision-rights via management reports or other review processes. For example, bank officers may have discretion in underwriting consumer loans, but most banks have in place “exception reports” which flag loans underwritten outside of formal guidelines for further review. Similarly, a local sales representative may have wide latitude in granting price discounts to local clients, but headquarters may have guidelines in place for flagging and reviewing “excessive” price discounting. In this paper, we focus on the behavioral and learning implications of this form of monitoring of decision rights.

When such monitoring mechanisms are in place, employees are likely to limit opportunistic behavior in their use of formal decision rights for fear of detection and possible dismissal from a valued job (Nagin et. al. 2002). Thus, these types of monitoring mechanisms may lead employees to be more cautious about using their discretion to exercise decision rights and we would expect to see them doing so less frequently. The literature on experimentation in organizations would make a similar prediction: employees facing more “evaluative pressure” are less likely to experiment, preferring instead the certainty of managing towards explicit guidelines (Lee et. al. 2004).

The learning implications of monitoring are less clear. As employees use discretion to exercise their decision rights, they may learn over time the conditions under which their discretionary decisions are effective. However, if monitoring leads to less use of discretion by employees, then they are in effect performing fewer “experiments” and have fewer opportunities to learn (Lee et. al. 2004). This “experimentation hypothesis” would predict negative learning implications of increased monitoring. Alternatively, if the threat of detection inherent in these monitoring systems leads employees to be more selective in utilizing decision rights and/or to expend more effort in learning how to use them effectively, then monitoring may lead to enhanced learning. This “selective utilization” hypothesis would predict positive learning implications of increased monitoring. Which of these alternatives is likely to prevail is an open empirical question which we address in this paper.

To investigate these issues, we exploit a field-research setting in which: (1) employees’ use of decision rights can be directly measured and traced to performance outcomes, and; (2) employees work in distinct business units which vary in their monitoring of decision-rights. Specifically, the data for this study come from the MGM-Mirage Holdings group of properties

and casinos. This setting offers a number of opportunities, and research design advantages, for our study. First, it is a setting where these issues are important. Operating employees (casino hosts or “hosts”) are allocated decision-rights to make real-time investment decisions on individual customers in the face of considerable uncertainty. Hosts must decide, often on the spot, how much to award in “comps” (e.g. free meals, discounted or free rooms, etc...) to individual customers based on the host’s assessment of the current and future economic worth of those customers. Second, observable performance information collected by the organization plays an important role in these decisions. Host assessments of the economic worth of customers are based in part on “hard” information in the form of historical customer profitability measurements and in part on “soft” or “local” information in the form of their own subjective assessments from detailed local knowledge of customers.

Third, elements of the actual decisions of operating employees, and outcomes of these decisions, are observable in this setting. The detailed data we obtained from this research site allow us to identify when specific hosts interact with specific customers, the dollar value of “comps” the host decided to award to the customer, and how that customer performs in the future. Thus, we are able to quantify the effectiveness of host decisions by linking their investments in customers (in the form of “comps”) to the future profitability of those customers. We are also able to explore learning in this environment by estimating how the effectiveness of these decisions varies with different measures of employee experience.

Finally, individual properties of MGM-Mirage Holdings vary in the extent to which hosts are monitored and reviewed in exercising their decision-rights over comp awards. Based on field-interviews and quantitative data, we are able to classify properties into two groups: those with relatively frequent or “tight” monitoring of employee use of decision-rights versus those

with relatively infrequent or “loose” monitoring of employee use of decision-rights. Formal decision rights are identical across properties. Thus, we are able to hold formal decision-rights constant while examining the behavioral and learning implications arising from “tight” versus “loose” monitoring of these formal decision-rights.

Using field and quantitative data from MGM-Mirage Holdings, we accomplish three primary objectives in this paper. First, we document how monitoring and enforcement of decision-rights gives rise to implicit incentives in this organization. We document implicit incentives in the form of sharp increases in employee exits linked to the “excessive” exercise of decision-rights. We find that these implicit incentives are strongest for employees in business units which monitor the use of these decision-rights most frequently (e.g. “tight monitoring”).

Second, we explore whether lower level employees appear to respond to the implicit incentives arising from the monitoring and enforcement of decision rights. We document two broad behavioral effects for employees in tighter monitoring environments: (1) they are less likely to use discretion to deviate from decision-guidelines in general; (2) they are less likely to incorporate local information, including historical performance data, when making local decisions. The first effect is straightforward: employees who are more likely to be penalized for mistakes are less likely to “experiment” (Lee et. al. 2004), preferring instead to manage towards explicit guidelines. The second effect persists even though historical performance data is observable to organizational superiors and can be used to justify deviations from organizational decision rules.

Third, and perhaps most important, we identify a link between employee learning and implicit incentives from the monitoring and enforcement of decision-rights. Our results are consistent with the “experimentation hypothesis”. Employees in “tightly monitored” business

units experiment less often by deviating less frequently from explicit decision guidelines and have fewer opportunities to learn how to effectively incorporate local information into their decisions as a result. Consistent with this notion, we find strong learning effects among lower level employees in our setting, but these effects are concentrated among employees in “loosely monitored” business units and almost entirely absent in “tightly monitored” business units.

The primary contribution of our paper is to the accounting literature on performance evaluation. There is, in general, a lack of evidence on the determinants or consequences of the variety of incentive mechanisms operating at lower levels of different organizations. With the exception of a small, but growing, number of studies on promotion and tournament-based incentives for lower level employees (e.g. Baker et. al. 1994, Gibbs 1995, Campbell 2008; Casas-Arce and Martinez-Jerez 2009) most of the studies on incentives for non-executive employees in the accounting literature have focused on short-term bonus plans based on formulaic or subjective weighting of performance metrics (e.g. Banker et. al. 2000, Ittner et. al. 2003; Gibbs et. al. 2004). We add to this literature by documenting the existence of implicit incentives arising from the monitoring of decision-rights as well as the behavioral and learning consequences of these implicit incentives.

While not the primary focus of our paper, our results and the context within which our study was conducted, are also highly relevant for the marketing literature on customer profitability measurement and management. In particular, a growing body of research in this area focuses on the measurement and dynamics of “lifetime” profitability (e.g. Reinartz and Kumar 2000; 2003). Researchers in this area have begun to explore the notion that customer profitability, or at least some of its components (e.g. “cost-to-serve”) are managerial decision variables (e.g. Gupta and Zeithaml 2006; Narayandas 2004). Our results demonstrate that

incentives and decision-rights can have a significant influence on how decentralized decision-makers utilize customer performance metrics (Casas-Arce et. al. 2009). In turn, the way these decision-makers use the information can influence the dynamics of customer performance metrics over time.

The remainder of the paper proceeds as follows. In the next section, we discuss the research setting and data for this study. Empirical tests and results are presented in Section III. Finally, we conclude with a brief discussion in Section IV.

II. Research Setting¹

The data for this study come from six of the major hotel properties of MGM-Mirage Holdings with a common information system for collecting customer performance data. As of the time of this study, the group operates numerous properties within Las Vegas and across other areas across the United States and Asia. Despite having united all these properties under the same corporate group through a series of acquisitions, each property retains its own personality and is managed independently of other MGM-MIRAGE holdings. The typical property in the group provides both lodging and entertainment with hotels ranging in size up to over 5,000 rooms. Properties in the group are full service hotels offering entertainment options in the form of gaming, live shows, restaurants, and bars and nightclubs. Each property has a casino offering a range of activities including table games, slot machines, and video poker.

Properties in the MGM-Mirage group place great emphasis on managing the profitability of their gaming (e.g. casino) customers. Customer treatment is a function of potential profitability. Each property tracks customer gaming behavior and performance, and awards

¹ Description of the empirical setting is partially based on Campbell et al. 2006.

comps to customers, based on the "Players Club" card and associated database. Available to all customers, the players club card is not a payment instrument, debit or credit, but a gaming loyalty card. It is inserted into slot machines in addition to the required currency, and presented at table games in addition to the money being wagered. Data on customer gaming behavior such as amount bet, games played, and time-played are captured electronically in the case of slot-machine play and manually in the case of table game play. The more the cardholder wagers on slots and tables, the more "points" he or she will automatically accumulate on the card. Points can then be redeemed for comps, which could typically include an abatement of the customer's hotel and entertainment costs, coupons for a free meal at a restaurant, or tickets for a show. Points accrued can be redeemed across properties at one hundred restaurants, ten shows, seven spas and salons, and for hotel accommodations.

Customer Profitability

Through its customer information system, MGM-Mirage properties could track 75% of gaming revenues to specific customers. MGM executives commonly refer to customers by their expected amount gambled per trip, a figure which serves as the basic unit of customer profitability, and which acts as the basis for most comping decisions. For example, a hotel suite might be offered to a "\$30,000 per-trip customer." The basic components of profitability are the gaming revenues and the comps that had been offered.

The profitability of a gaming customer is calculated based on the casino's so-called "theoretical win" from the customer—the amount of revenues that the property could *theoretically* make based on the amount the customer bet and the mix of games played—not the "actual win". The primary purpose of measuring the profitability of gaming customer relationships based on theoretical win rather than on what the customer actually won or lost from

the property is to effectively not penalize a customer for luck. One employee explained the logic by noting "...we don't care if you win or lose as long as you give us a shot at your money".

Gaming revenues are recorded and computed differently at slot machines and table games. For slots, electronic gambling machines of various high and low denominations, the play of a customer using a Players Club card could be recorded in great detail. By multiplying the total amount the customer puts into the machine by the theoretical hold percentage of the specified machine, the casino's "theoretical win" from the customer could be computed. This detailed information allows properties to trace around 65% of all slot revenues to individual customers who use the card.

Table games, unlike slot machines, rely on human dealers for their operation. They include games such as blackjack, roulette, baccarat and craps; different table games have different sources of revenue, whether from rakes² or the varying levels of house advantage. The calculation of theoretical win for table games is also based on information recorded onto the computer system by floormen³ through rating slips, which, although not as detailed (or necessarily accurate due to reliance on human judgment to assess amount bet, time played and game speed), nonetheless contain vital information. For a player's gaming to be recorded, he or she needs to ask the floorman to be rated (or the floorman has to identify the customer from prior visits), and the floorman will then manually record the name and card number of the player to begin gathering the necessary information on the customer's play. The theoretical win for table games is calculated as follows: Average bet x Hours played x Hands per hour x House advantage

² Rake refers to the percentage taken by the house in games such as poker where the casino does not put its money into play.

³ Game supervisors.

on the game - assuming the average skill of a knowledgeable player. Properties are able to trace around 85% of all table revenues to individual Players Club customers.

Data captured in the database via the players club card covers more than ten years of operations and more than 8 million customers. In the calculation of customer profitability, costs, such as those of operations and marketing personnel, equipment maintenance and real estate, are not assigned to the customer.

Comps

Comps are an important component of a gaming customer's profitability: because they are used to reward the gaming behavior of customers, comps are considered to be an expense that these customers generate. Comps are tracked in the Players Club database. The way they are recorded into the system differs depending on the nature of the comp. Generally, it is possible to distinguish two types of comps recorded in the database:

Soft costs in the form of complimentary services: Complimentaries in the form of room, food, beverage, entertainment or any other service provided by the hotel are recorded to the customer's Players Club account as expenses incurred on that customer. Typically these services are recorded as expenses at the value at which they are purchased by full-paying customers.

Outside expenditures: Any other complimentary service that implies a payment to an outside provider (for instance, reimbursements for plane tickets) appears as an expense in the Players Club database for its full amount.

Casino Hosts

Individual properties rely heavily on employees known as "casino hosts" in order to initiate and manage customer relationships while simultaneously ensuring the profitability of those relationships. Hosts engage with all segments of customers who qualify for comps. Using a combination of subjective observation, historical customer data, and customer profitability

reports, hosts are able to “red flag” gaming customers and determine their "theoretical win" – the amount that the host expected the casino to win from the customer given the customer's observed or anticipated gaming behavior (e.g. choice of game, average bet, time played, etc...).

Decision-Rights and Incentives for Casino Hosts

Casino hosts ultimately have decision rights on the comps awarded to their customers. The primary decision variable of the casino host is the "comp percentage" – the dollar value of comps awarded as a percentage of the customer's theoretical win on the current trip. Typically, a host will be willing to comp customers up to 40% of their theoretical win on the customer's current trip to the property. Casino hosts are explicitly provided with such decision-rights by property management, and 40% is stated formally as the maximum comp percentage that should be awarded to customers barring exceptional circumstances. The definition of decision-rights for casino hosts in terms of maximum comp percentages of 40% appears to be common practice in the industry. According to one highly experienced senior executive we interviewed "...the general rule of thumb in the industry is 40% for the maximum comp percentage. This has been true for at least the past 20 years".

The level and type of comps awarded to a customer on a given trip are typically agreed upon between the host and the customer at the start of the trip based on the expected gaming behavior of the customer during that trip. Often, the host and the customer establish a sort of “verbal contract” regarding the comps that will be awarded if the customer plays as agreed during the trip. If *ex post* a customer does not play at the level expected by the host, then the host can adjust the comp award downward to meet the 40% threshold. For example, as one host noted “...a customer can sweet talk you to get a [suite], but if his play ends up not warranting it, then I will look at the folio at the end of the trip and make any necessary adjustments to the expenses we will and will not cover.” However, if the host believes that a customer is likely to

be highly profitable in the future through return trips and increased gaming activity, then they can, and often do, use their decision-rights to award comps well in excess of 40% of the customer's theoretical win on the current trip.

When casino hosts exceed the maximum comp percentage of 40% for a given customer-trip, an "exception report" is triggered for review by the CFO of the property. The exception report provides information on the type and dollar value of comps awarded, the customer's current and historical theoretical win, and, for some properties, a written explanation by the host on why the customer was "over-comped" on the current trip. These exception reports appear to provide powerful implicit incentives for hosts to carefully manage comp percentages. Echoing a common sentiment we heard in interviews with casino hosts across properties, one host noted "...I am careful about overextending on comps. I don't do anything I can't justify. I know I have to go into the quarterly meeting and defend my decisions".

Interestingly, all incentives for hosts to directly manage the key decision variable of comp percentage appear to be implicit rather than explicit. We did not find any property that explicitly tied the compensation of casino hosts to the comp percentage. The typical explicit incentive plan for casino hosts is well illustrated by the following plan for one of MGM-Mirage's properties. Casino hosts can receive a percentage of their salary in the form of an annual bonus based on several criteria. 60% of the bonus is determined by total gaming revenues for the property; 15% is based on individual goals on acquiring new, and reactivating "inactive" customers; and 25% is based on subjective evaluation of the performance of the host by senior management of the property. Incentives to award comps to customers are driven by this plan only to the extent that such awards facilitate new customer acquisition, reactivation of inactive customers, and growth in overall gaming revenues. Incentives to limit comp awards to

customers appear to be primarily determined by the exception report review process but may also depend on the extent to which managers consider the level of comps awarded by a host in the subjective evaluation component of the bonus plan.

The Casino Host Decision-Process

As part of this study, we interviewed several casino hosts across properties and also "shadowed" some of them as they performed their work. Our interviews and observations revealed several important features of the decision-process of casino hosts in their interactions with customers as follows:

Static vs. dynamic view of customer relationships: In general, hosts take a dynamic rather than static view of the financial performance of customer relationships. In a representative comment, for example, one host noted "...I usually give it a couple of trips to see what happens. I am looking to see how that customer is likely to perform over the next year or so...". Another explained "...when I overcomp a customer, I am looking at what that will mean for [his or her] profitability for the entire year". Many hosts also spoke of customer loyalty. For example, one host explained "I can usually see within three to six trips whether a customer is loyal". While hosts vary in how they define the appropriate future horizon for assessing customer relationships, it is clear that their expectation of future customer performance shapes their individual comp decisions.

Reliance on hard vs. soft information: Our interviews and observations suggest that hosts vary greatly in the relative extent to which they base their discretionary comp decisions on historical customer information contained in the firm's database versus their own subjective evaluations based on "local knowledge" of customers. When hosts exercise discretion by awarding comps to customers in excess of 40% of the customer's current trip theoretical win,

they are typically doing so based on expectations about the future profitability of those customers. All hosts that we interviewed or observed appeared to form such expectations in part based on directly observable historical data on the customer such as the time since the last trip at the property and various measures of the customer's theoretical win on previous trips (e.g. average over the past 'N' trips, total over lifetime, etc...). Hosts also form expectations about customer performance based on knowledge gained from direct interaction with customers. Often, this "soft" information is in the form of behavioral cues which hosts believe are signals of a potentially valuable customer. For example, one host explained "...I often look for body language. The person who is sitting at the slot machine with his feet up and 500 credits in the machine probably makes sense to focus on...". While all hosts seem to rely, at least in part, on such soft information, the more experienced hosts that we interviewed appeared to express relatively more confidence in their ability to incorporate such information into their decisions. In expressing a common sentiment we heard from interacting with experienced hosts, one host noted that "...good customers just play differently. I just know when it makes sense to overcomp. I have been here a long time". Similarly, when we asked one host several questions about how she makes her comp decisions, she replied "...there is nothing black and white in what we do, it is all grey". Hosts also make decisions based on knowledge gained from repeated interaction with individual customers. Many hosts talked in terms of establishing "relationships" with customers. As one host described "I use discretion based on longevity and what I know about the customer". Another noted that "...it is often difficult to separate out the social relationship vs. the business relationship".

Heuristics for incorporating hard information: In our interactions with hosts, we observed a significant amount of heterogeneity in the heuristics hosts have developed for incorporating

historical customer information from the database into their comp decisions. The following are representative comments that emerged during our interviews with hosts:

I tend to do more of a trip-by-trip evaluation.

I look at the past four trips, throw out the lowest one, and take the average of the remaining three. I'll then update based on current play.

I look at what they have done over the past three trips as long as those trips have been within the past 18 months.

I look at both year-to-date theoretical win and lifetime to date.

I will look at the last three trips but not exclusively. What if their last two trips were two years ago?

I have set up my system to show lifetime theoretical win and theoretical win over the past 12 months.

As demonstrated by these comments, hosts vary considerably in both their views on the appropriate past horizon over which to consider historical customer performance data and in how best to aggregate past information for incorporation into current comp decisions.

Managing expectations: Hosts view managing customer expectations for comps as an important objective. Consistent with their tendency to take a dynamic rather than static view of customer performance, casino hosts tend to think carefully about how their discretionary comp decisions will influence customer expectations for comp levels in the future. As one host explained, "I try to limit overcomping to manage expectations". Another noted, "I tell [customers] when they are overcomped that they shouldn't necessarily expect the same thing next time...". Some of our interviews suggested that this focus on managing customer expectations partly results from the implicit incentives in the exception report review process. Hosts appear to consider not just whether they "overcomp" a customer on a current trip but also

whether comps awarded relative to the customer's theoretical win across multiple trips balance out to prescribed limits (e.g. 40%) over time. As one experienced host explained "I perform well, so I can absorb overcomps. It will all balance out in the end."

Data

The quantitative data for this study are collected from the MGM-Mirage customer information system. The sample used for subsequent analysis consists of 349,887 observations on customer-trips at a specific property with a specific casino host. These customer-host-property-trip level observations occur over the period 1993-2004. The firm's database contains data on over 9 million customer trips over this period. However, the vast majority of these trips are for customers who do not interact with hosts. These customers either choose not to interact with hosts and just use the players club card to accrue points for automatic redemption or they are customers whose level of play does not warrant host interaction. Because, the primary focus of this paper is on the host-level learning and decision-process, we only focus on customer-trips that are assigned to a specific host.

In several cases in the data, one-time customers with very low levels of theoretical win (e.g. due to limited gaming, small bets, etc...) and even zero theoretical win received comps valued in the tens of thousands of dollars. These observations are likely to be on customers who are family members or friends of highly valuable customers. We eliminate these extreme outliers from the data to arrive at our final sample of customer-host-property-trips.⁴

For each trip-level observation we observe the unique identity of the host interacting with the customer, the theoretical win for the customer on the trip, and the dollar value of comps

⁴ Complimentary services awarded to family members and friends of valuable customers can be seen as comps for that customer. However, there is no systematic way in the firm's database to identify which players are comped based on the level of play of a related valuable customer.

awarded to the customer for that trip. Because we observe data from 1993 (the year of inception of the firm's players club program and database) onwards, we can reconstruct each host's history of interactions with individual customers. We exploit this to construct measures of general experience (e.g. cumulative number of all past customer-trips assigned to the host) and specific experience (e.g. cumulative number of past trip-interactions with a specific customer). In many of our tests, we aggregate the individual customer-trip level data at the host-level to form annual "customer portfolio" performance metrics for hosts.

III. Empirical Tests and Results

Classification of Properties in Terms of "Tight" versus "Loose" Monitoring

Table 1 provides evidence that properties differ in the intensity with which they monitor host comp decisions via the exception report review process. We classify properties 1-3 as those with "tight monitoring" and properties 4-6 as those with "loose monitoring" based on the relatively high frequency with which exception reports are monitored in properties 1-3. In properties 1-3, host comp decisions are monitored daily by the CFO of the property and those which exceed the 40% guidelines are flagged for possible follow up. A broader review of each host's overall customer portfolio is conducted at the end of each month in these properties. In properties 4-6, host comp decisions are monitored only once per week and the broader review of a host's entire portfolio is only conducted once per quarter. Adding qualitative support to our classification of properties, interviews with senior management and hosts at each property also revealed that formal written explanations are required from hosts documenting their rationale for "over-comping" customers on a given trip in properties 1-3 but not in properties 4-6 suggesting tighter monitoring in the former.

Implicit Incentives in Properties with Tight vs. Loose Monitoring

Our classification of properties into those with “tight” versus “loose” monitoring of host discretionary comp decisions relies on interviews with corporate and property level management as well as the documented differences in monitoring frequency shown in Table 1. We investigate the validity of our classification in greater detail by examining the existence of implicit incentives in this organization, and the degree to which they differ for properties classified as having “tight” vs. “loose” monitoring practices. We accomplish this by testing the sensitivity of host exit from a property (*EXIT*) to a variety of measures of host performance. Performance-related departure from the organization is a potentially extreme outcome. However, as discussed in the previous section, hosts appear to consider career concerns stemming from the exception report review process when making their comp decisions. To the extent that “excessive” use of decision-rights is linked to departure, this will create implicit incentives for hosts to manage the exercise of those decision-rights.

Hosts face incentives to acquire new, and retain and reactivate existing customers. They also face implicit incentives, via the exception report process, to manage the dollar value of comps awarded as a percentage of theoretical win. To examine the potential strength of these different incentives, and how they vary for properties with tight or loose monitoring, we estimate the following exit-performance relationship:

$$\begin{aligned}
 Exit_{jpt}^* &= \beta_0 + \beta_1 CustomerGrowth_{jpt-1} + \beta_2 TripsPerCustomerGrowth_{jpt-1} + \beta_3 TheoreticalWinPerTripGrowth_{jpt-1} \\
 &\quad + \beta_4 Discretion\%_{jpt-1} + \beta_5 Overcomped_{jpt-1} + \beta_6 ExcessComps_{jpt-1} + \beta_7 ExcessComps \times TightMonitoring_{jpt-1} \\
 &\quad + \beta_8 Experience_{jpt-1} + \sum_{j=2}^6 \gamma^j Property_p^j + \sum_{k=1994}^{2003} \gamma^k Year_t^k + \varepsilon_{jpt}
 \end{aligned}$$

$$Exit_{jpt} = \begin{cases} 1 & \text{if } Exit_{jpt}^* > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Our dependent variable, $Exit_{jpt}$, is set equal to 1 if host j departs from property p during year t . We model the probability of departure as a function of a number of host-level

performance metrics including growth in the number of customers in the host's portfolio (*CustomerGrowth*), growth in the number of trips per customer in the host's portfolio (*TripsPerCustomerGrowth*), and growth in theoretical win per trip for these customers (*TheoreticalWinPerTripGrowth*). Because these metrics are primary objectives of MGM-Mirage's properties, we expect each of them to be negatively associated with the probability of departure. All measures are aggregated at the host-year level.

To gauge the extent to which the exercise of decision-rights is punished or rewarded, we also include in the specification a number of measures of the extent to which hosts use decision-rights in their comp awards to customers. These include: the proportion of individual customer-trips managed by the host during the year in which comps were awarded in excess of 40% of theoretical win (*Discretion%*); an indicator for whether the total comps awarded by a host during the year is in excess of 40% of the total theoretical win across all customers in the host's portfolio for that year (*Overcomped*); and for those hosts whose portfolios are "overcomped" for the year, a measure of the extent to which total comps awarded by a host during the year is in excess of 40% of the total theoretical win across all customers in the hosts portfolio for that year (*ExcessComps*).

The final measure, *ExcessComps*, takes on a value of 0 for hosts who are not "overcomped" in a given year and a value of $[100 * (\text{Total comps awarded by the host for the year} \div \text{Total theoretical win across all customers managed by the host during the year}) - 40]$ for hosts who are "overcomped" for the year. This measure is of particular interest as it captures not only the extent to which a host awarded comps to individual customers outside of the formal 40% guideline but also that the host was not able to manage their customer portfolio so as to absorb any overcomping to individual customers. Thus, this measure partially captures the effectiveness

with which a host exercises her decision-rights. As a result, we expect that, all else equal, high levels of this measure would be associated with a higher probability of departure.

We interact *ExcessComps* with *TightMonitoring*, an indicator for whether the host is employed at any of the three properties which we classify as a “tight monitoring” property in Table 1. If our classification of properties is valid, then we expect that an increase in *ExcessComps* will lead to a higher increase in the probability of departure for properties classified as having “tight monitoring” vs. those classified as having “loose monitoring”. That is, we expect the coefficient, β_7 , on the interaction term to be positive. As additional control variables, we include property indicators, year indicators, and *Experience* measured as the number of years a host has been employed at a property at the beginning of each year.

Table 2 contains results from logit estimation of equation (1). All standard errors are adjusted for clustering of observations within hosts over time prior to inference. Column 1 demonstrates that hosts who are able to grow their customer base, either through reactivating inactive customers or acquiring new ones, are less likely to exit the organization. Similarly, hosts who are able to improve the theoretical win generated per customer trip (*TheoreticalWinPerCustomerGrowth*) face a lower probability of departure. Managers with higher levels of experience face lower departure probabilities. Holding customer growth and other performance metrics constant, hosts who are overcomped across all customers in their portfolios (*Overcomped*) are more likely to leave the organization. This increase in the probability of departure is increased further by the degree to which the host is overcomped (*ExcessComps*). Interestingly, conditional on whether, and the degree to which, the host is overcomped, *Discretion%* is unrelated to the probability of departure. Overall, these results provide evidence of incentives in this organization whereby the exercise of decision rights

(*Discretion%*) is not discouraged *per se*, but where hosts face strong incentives, linked to career outcomes, for managing the effectiveness with which these decision rights are exercised.

The results in column 2 are largely consistent with those in column 1 but also point to the validity of our classification of properties in terms of tight versus loose monitoring. The coefficient on *ExcessComps*×*TightMonitoring* is positive and significant suggesting that “overcomping” is more strongly discouraged in the properties that we classify as having tight monitoring compared to those we classify as having loose monitoring. The estimates reported in column 2 show that a host with the mean levels of all performance measures and operating below the threshold comp percentage of 40% has a 0.3% probability of departing from the organization. An overcomped host in a “loose monitoring” property with the mean levels of all performance measures but operating with a portfolio level comp percentage in the 90th percentile has a probability of exiting the organization that is approximately four times higher at 1.3%. By contrast, a similar overcomped host in a “tight monitoring” property has a probability of exiting the organization of 6.3% – an almost five fold increase when compared to the host in a “loose monitoring” property.

In summary, the results in this section, when combined with the evidence on monitoring frequency shown in Table 1, provide further support for our classification of properties into those with “tight” versus “loose” monitoring. Implicit incentives linked to career concerns (as measured by departure probabilities) are present in properties that we classify as having “loose” monitoring but are strongest in properties we classify as having “tight” monitoring. The qualitative evidence on incentives discussed in section II, coupled with the results in this section, also demonstrates that employees in this organization face strong outcome-based incentives to acquire and retain customers (e.g. customer growth) and increase gaming activity (e.g.

theoretical win per customer trip to a property). In the next two subsections, we explore whether host decision-making is consistent with differences in the implicit incentives arising from tight versus loose monitoring and the implications, if any, for learning in this environment.

Does the Exercise of Decision-Rights Reflect Differences in Implicit Incentives from Tight versus Loose Monitoring ?

In this section, we explore whether any differences in how hosts exercise decision rights reflect implicit incentives arising from tight versus loose monitoring. There are at least two potential effects of tighter monitoring on the exercise of decision-rights: (1) employees may be less likely to use discretion to deviate from guidelines in general; (2) employees may be more or less likely to incorporate local information, including historical customer data, when making decisions about individual customers. The first effect is straightforward: employees who are more likely to be penalized for mistakes are less likely to “experiment” (Lee et. al. 2004), preferring instead to manage towards explicit guidelines.

The second effect is not as straightforward. In many settings, including ours, operational employees have access to local information that can signal when decision-rights should be exercised to exceed guidelines. Such local information can be “hard” (e.g. historical information on customers in the firm's database) or “soft” (e.g. from direct interaction with a customer). The latter is not, while the former is, observable to non-local decision-makers. We might expect that employees in tight-monitoring environments, where they face more pressure to make the “right” decision, would be less likely to incorporate local information into their discretionary decisions. This is particularly true for “soft” local information. In this case, if employees’ discretionary decisions do not pay off in current or future customer performance (e.g. increased sales or retention), then these decisions will be more difficult to justify to superiors in the organization. However, to the extent that “hard” local information can be used as a basis for exercising

discretion, employees may be more likely to incorporate it in their decision making even in tight monitoring environments. In particular, if employees facing outcome-based incentives (e.g. customer growth or retention) feel restrained from making decisions based on “soft” information, then they may be more willing to exercise decision-rights when justified by observable (to local and non-local decision-makers) “hard” information.⁵

Turning first to the question of whether employees in tight monitoring environments are less likely to exercise decision-rights in general, the results presented in Table 3 demonstrate that this is indeed the case. As measured by *Discretion%*, hosts in tight-monitoring properties are less likely to exceed the 40% comp guidelines than those in loose-monitoring properties (mean for tight-monitoring properties=19.6%; mean for loose monitoring properties=29.2%; difference significant at $p<.01$). Hosts in tight-monitoring properties are also almost three times less likely than those in loose monitoring properties to be overcomped at a portfolio level (*Overcomped*) in any given year (mean for tight-monitoring properties=0.139; mean for loose monitoring properties=0.368; difference significant at $p<.01$). Finally, hosts in tight-monitoring properties award substantially less comps relative to the theoretical win of their customer portfolios (*Comp%*) in a given year (mean for tight-monitoring properties=35.6%; mean for loose monitoring properties=59.8%; difference significant at $p<.01$). Overall, the results in Table 3 are consistent with implicit incentives arising from tight versus loose monitoring across properties. By all measures of the extent to which hosts are exercising decision rights, discretionary decisions which deviate from comp guidelines are significantly less prevalent in properties we classify as having tight monitoring.

⁵ For an example outside of our research setting, a sales employee may justify giving a large price discount to a customer purchasing a small quantity of product because historical data on that customer suggests that the current purchase is a one-time deviation from a normal pattern of purchasing large quantities.

Turning next to the question of whether employees in tight monitoring environments are more or less likely to exercise decision-rights based on local customer information, we develop and test an empirical model of host decisions at the individual customer-trip level. As is clear from the qualitative evidence presented in Section II, hosts vary significantly in both how, and the horizons over which, they combine "hard" (e.g. historical information on customers in the firm's database) and "soft" (e.g. direct observation of customer behavior) information in their decisions regarding the amount of comps to award to a customer on a given trip. Capturing the complexities of these decisions in observable data is not trivial. As an approximation, we characterize the casino host decision process as:

$$COMP_{ijpt} = \beta_1 TheoreticalWin_{ijpt} + \beta_2 E_j \left(\sum_{k=1}^{T_i} TheoreticalWin_{ipt+k} \right)$$

where $COMP_{ijpt}$ is the dollar value of comps awarded to customer 'i' by host 'j' at property 'p' during trip 't'. In this characterization of their decision process, hosts determine the dollar value of comps to award a customer based on two pieces of information: the observed theoretical win of the customer on the current trip ($TheoreticalWin_{ijpt}$) and the host's expectation of the

customer's future theoretical win at the property ($E_j(\sum_{k=1}^{T_i} TheoreticalWin_{ipt+k})$).⁶ The customer's

future theoretical win at the property is determined by both the level of theoretical win per trip and the number of future trips to the property by the customer (T_i). If hosts simply award comps to customers based on their observed level of play during the current trip, rather than

⁶ The absence of the 'j' subscript in $\sum_{k=1}^T TheoreticalWin_{ipt+k}$ is intentional. Its absence captures the notion that hosts have incentives to bring customers back to the property even if that customer switches hosts in the future.

using discretion to deviate from the 40% limit on comps awarded per dollar of current trip theoretical win as prescribed by their formal decision rights, then $\beta_2 = 0$ and $\beta_1 \leq 0.40$. Alternatively, if hosts use discretion to deviate from the prescribed 40% limit based on their expectations about the customer's level of future theoretical win at the property, then $\beta_2 > 0$ and $\beta_1 \leq 0.40$.

Consistent with our interviews and observations, we assume that hosts form expectations about future customer performance at the property based in part on historical data on the customer's theoretical win $(\sum_{s=1}^{L_i} TheoreticalWin_{ipt-s})^7$ and in part based on idiosyncratic "soft" information about the customer observed by the host at the property during the customer's current trip (α_{ijpt}). The past number of trips, L_i , considered for each customer depends on the horizon considered relevant by the host. The host may consider the full tenure of the customer relationship with the property in which case L_i would equal the total number of past trips by the customer to the property. Alternatively, hosts may discount information on older trips, in which case L_i would be determined by a shorter time period. The soft information represented by α_{ijpt} can be of several types including direct interactions between the host and customer in which the host inquires about the intent of the customer on current and future trips, inferences the host makes about the customer's appearance and behavior, or any other local information the host

⁷ The absence of the 'j' subscript in $\sum_{s=1}^{L_i} TheoreticalWin_{ipt-s}$ is intentional. Its absence captures the notion that hosts are simply using the customer's past level of performance at the property to make inferences about performance in the future. The past performance of the customer need not be the result of trips in which the customer interacted with the host making the decision on the current trip.

gains outside of the systematic customer data captured in the firm's information system. With this characterization, the host comp decision can be modeled as:

$$COMP_{ijpt} = \beta_1 TheoreticalWin_{ijpt} + \beta_2 \left(\lambda_1 \sum_{s=1}^{L_i} TheoreticalWin_{ipt-s} + \alpha_{ijpt} \right) \quad (2)$$

This characterization of the host decision process has intuitive appeal. Conditional on the customer's current level of play, hosts increase (decrease) the comp awarded when the customer's past level of play is high (low) which may signal that the current trip is a deviation from a pattern of performance established by the customer in past trips to the property. Similarly, hosts increase the comp percentage when they observe local information on the customer that suggests higher future levels of play at the property (α_{ijpt}).

As evidenced by our interviews, the horizons considered relevant for decision making vary considerably across hosts. In our empirical specifications, we choose the relevant past horizon as the 18 months prior to the current trip start date. That is, in equation (2), we set L_i equal to the number of past trips the customer has taken to the property within the 18-month period prior to the current trip. We make this choice for the two primary reasons. First, a customer formally becomes classified as "inactive" if they have not returned for 18 months from their last trip to the property. Whether the customer remains active, and how active they remain, over the 18 months subsequent to the current trip appear to be salient criteria for decision-making in our setting. Second, while many hosts suggested that they consider the entire history of customer data in their decision-making, they almost universally noted that they discount information that is greater than 18 to 24 months old. In the remainder of the paper, for each

customer-host-property-trip observation, we refer to $\sum_{s=1}^{L_t} \text{TheoreticalWin}_{ipt-s}$ as

LagTheoreticalWin. For empirical purposes, we estimate the following version of equation (2):

$$COMP_{ijpt} = \hat{\beta}_1 \text{TheoreticalWin}_{ijpt} + \hat{\beta}_2 \text{LagTheoreticalWin}_{ipt} + \sum_{j=2}^6 \gamma^j \text{Property}_p^j + \sum_{k=1994}^{2003} \lambda^k \text{Year}_i^k + \mu_j + \varepsilon_{ijpt} \quad (2')$$

where μ_j denotes a host fixed effect controlled for through the use of a series of host indicators.

To examine whether employees in tight monitoring environments are more or less likely to exercise decision-rights based on local customer information, we also estimate a version of equation (2') where we allow the empirical weights, $\hat{\beta}_1$ and $\hat{\beta}_2$, on current and historical theoretical win to vary for tight versus loose monitoring properties.

Results from OLS estimation of equation (2') are presented in Table 4. All standard errors are adjusted for clustering of observations within customers over time prior to inference. The results in column 1 demonstrate that, on average, hosts weight both current and past customer information in their comp decisions. The coefficient estimates show that, conditional on past customer performance, hosts on average award \$0.218 in comps per dollar of current trip theoretical win – well within the comp guidelines of all properties. On average, the comp award is adjusted upward by \$0.011 per dollar of historical theoretical win (over the previous 18 months). Based on the coefficient estimate on *LagTheoreticalWin*, a customer's past performance would have to be significantly higher than that on the current trip for a host to substantively shift the average comp percentage to be in excess of the 40% comp guideline. For the median customer-trip in our sample, *LagTheoreticalWin* is approximately two times the current trip theoretical win while it is approximately 16 times current trip theoretical win for

customer-trips in the 90th percentile in our sample. For the median customer-trip, hosts would on average increase the comp percentage (comps relative to current trip theoretical win) by 2.2% (2*1.1%) to 24% whereas for the customer-trip in the 90th percentile, the comp percentage would on average increase by 17.6% (16*1.1%) to 39.4%. Thus, on average, past customer performance would have to deviate significantly from current trip performance for hosts to utilize discretion so as to exceed formal comp guidelines.

Column 2 contains the results from estimating a version of equation (2') which allows the empirical weights on current and historical customer information to vary for properties with tight versus loose monitoring. The coefficient on *TheoreticalWin_{t-k} x TightMonitoring* is negative and significant (coefficient=-0.061; p<0.01). This result is consistent with those in Table 3 and documents that hosts in properties with tight monitoring tend to have lower comp percentages compared with their counterparts in properties with loose monitoring. The coefficient estimates show that, conditional on past customer performance, hosts in properties with loose monitoring on average award \$0.259 in comps per dollar of current trip theoretical win while those in properties with tight monitoring on average award \$0.198 (0.259-0.061) per dollar of current trip theoretical win. The coefficient on *LagTheoreticalWin_{t-k} x TightMonitoring* is negative and significant (coefficient=-0.007; p<0.01). On average, the comp award is adjusted upward by \$0.015 per dollar of historical theoretical win (over the previous 18 months) for loose monitoring properties compared with \$0.008 (0.015-.007) for tight monitoring properties. For the median customer-trip in our sample, hosts in loose monitoring properties would on average increase the comp percentage (comps relative to current trip theoretical win) by 3% (2*1.5%) to 29% whereas for the customer-trip in the 90th percentile, the comp percentage for hosts in these properties would on average increase by 24% (16*1.5%) to approximately 50% – well over the

formal comp guidelines of 40% of current trip theoretical win. The comparable numbers for hosts in tight monitoring properties are a 1.6% ($2 \times 0.8\%$) increase in the comp percentage for the median customer-trip and a 12.8% ($16 \times 0.8\%$) increase for a customer-trip in the 90th percentile of historical theoretical win – neither of these increases would lead hosts to exceed formal comp guidelines on average. Thus, hosts in tight monitoring properties tend to adjust comp awards less in response to historical customer information than their counterparts in loose monitoring properties.

Column 3 contains results from estimating a version of equation (2') which allows the host trip-level comp decision for an individual customer to vary with the performance of the host's entire customer portfolio. Specifically, we add the host-year level variables *Overcomped*, *ExcessComps*, and their interactions with *TightMonitoring* to the specification. *Overcomped* and *ExcessComps* are measured for the year prior to the year of the current customer trip. The qualitative data from our interviews with hosts (discussed in Section II) along with the empirical results linking host-performance to departure in Table 2 demonstrate that hosts face incentives to manage their entire customer portfolios in addition to individual customer relationships. These incentives may lead hosts to vary their comp decisions for individual customer trips based on the extent to which they are overcomped across all customers in their portfolio.

The results in column 3 show that this is the case for properties with tight monitoring but not for those with loose monitoring. The coefficient on *Overcomped* is positive and significant (coefficient=49.95; $p < .01$) while that on *OvercompedxTightMonitoring* is negative and significant (coefficient=-36.7; $p < .01$). This suggests that, in properties with loose monitoring, hosts with overcomped portfolios in the prior year continue to award higher levels of comps conditional on current and past theoretical win. In properties with tight monitoring, there is no

relationship between the customer-trip level comp decision and being overcomped at a portfolio level *per se* (coefficient for tight monitoring properties=49.95-36.7=13.25; F=1.01, p=0.32). However, the extent to which a host's portfolio is overcomped in the prior year (*ExcessComps*) appears to influence the customer-trip level comp decision for hosts in properties with tight monitoring but not in those with loose monitoring (coefficient on *ExcessComps*=1.745, p>.10; coefficient on *ExcessCompsxTightMonitoring*=-2.635, p<.05; sum of two coefficients=-0.89; F=3.70, p=0.054). The coefficient estimates demonstrate that, on average, each 1% increase in the extent to which a host in a tight monitoring property is overcomped at a portfolio level is associated with a \$0.89 decrease in the comps awarded to a customer on a given trip. These results provide evidence that hosts in properties with tight monitoring weight aggregate portfolio level information on their customers when making individual customer-trip comp decisions. However, this effect is relatively small. For the median theoretical win in the customer-trip level sample of approximately \$700, a host in a tight monitoring property with *ExcessComps*=40 (e.g. overcomped at twice the existing guidelines) would reduce the comp percentage on an individual customer trip by only 5.1% ($40 \times 0.89 / 700$).

In summary, the results in this section document three specific patterns which are consistent with implicit incentives from “tight” versus “loose” monitoring influencing the exercise of decision-rights. First, discretionary decisions which deviate from comp guidelines are significantly less prevalent in properties we classify as having tight monitoring compared to those we classify as having loose monitoring. In effect, this leads to less “experimentation” as hosts manage closer to comp guidelines. Second, the comp decisions of hosts in tight monitoring properties are less responsive to past customer performance information than are those of hosts in loose monitoring properties. This is true despite the fact that historical customer performance

is observable to organizational superiors and can provide a basis for justifying discretionary decisions. Finally, while responding less to historical customer performance in their discretionary comp decisions, hosts in tight monitoring properties respond more strongly to aggregate information on their own overall performance compared to their counterparts in loose monitoring properties. In the next subsection, we explore the implications of implicit incentives from tight versus loose monitoring, and the associated patterns in the use of decision-rights, for organizational learning.

Learning and the Tight vs. Loose Monitoring Effect

In this section, we explore the extent of any learning effects due to implicit incentives arising from tight versus loose monitoring. Documenting learning in the decentralized information processing activities of our sample of casino hosts requires that we develop an empirical model to identify how the link between these decisions and performance outcomes varies as hosts gain experience. To the extent that hosts develop skill in incorporating unobservable (to the researcher) local information in their comp percentage decisions, these decisions should be correlated with actual realizations of future customer performance after controlling for observable historical customer performance leading to the following empirical specification:

$$\begin{aligned}
 \textit{TheoreticalWin}_{ijpt} = & \beta_1 \textit{TheoreticalWin}_{ijpt-1} + \beta_2 \textit{Comps}_{ijpt-1} + \beta_3 \textit{Comps}_{ijpt-1} \times \textit{Experience}_{jt} + \beta_4 \textit{Experience}_{jt} + \\
 & + \sum_{j=2}^6 \gamma^j \textit{Property}_p^j + \sum_{k=1994}^{2003} \lambda^k \textit{Year}_t^k + \mu_j + \varepsilon_{ijpt} \quad (3)
 \end{aligned}$$

where ‘i’, ‘j’, ‘p’, and ‘t’ subscript customer, host, property, and time respectively and μ_j denotes a host fixed effect.

Equation (3) is our basis for identifying learning in the customer management decisions of casino hosts in our research setting. If hosts are, on average, skilled at incorporating local information that is informative of future customer performance into their comp decisions, then we expect $\beta_2 > 0$ – hosts deviate from basing comp decisions purely on historical customer data only when future customer performance is high relative to current customer performance. If ability in acquiring and incorporating local information into comp decisions increases as hosts gain experience interacting with customers, then the relationship between these decisions and future performance outcomes should increase with experience implying $\beta_3 > 0$. In some of our estimations, we will also allow the learning effect, β_3 , to vary for properties with tight vs. loose monitoring.

We estimate equation (3) in two ways. First, we aggregate all data up to the annual host portfolio level and analyze whether hosts' investments into their customer portfolios, in the form of comp awards, lead to increased future theoretical win at the portfolio level. This approach will allow us to capture learning effects related to managing a portfolio of customer relationships as opposed to individual customers. For this specification, we measure *Experience* as the number of years a host has been employed at a property at the beginning of each year. *Comps* and *Experience* are mean centered prior to interaction to maintain interpretability of coefficients. We estimate the model using the generalized method-of-moments dynamic panel data model of Arrelano and Bond (1991).⁸ The results are presented in Table 5. Consistent with the notion

⁸ A long literature in econometrics has addressed bias in panel data models with lagged dependent variables and fixed effects in panels with a short time-series dimension (e.g. see Wooldridge 2002). Any attempt at eliminating the host fixed effects, u_i , through either first-differencing or subtracting within-host averages will lead the error term in this equation to be correlated with the lagged dependent variable by construction. This leads to biased coefficient estimates, with the magnitude of the bias decreasing in the time-series dimension of the panel and not the cross-sectional dimension. The generalized method of moments estimator (GMM) developed in Arrelano and Bond (1991) addresses this problem through first differencing the equation to eliminate the u_i and then using all available lags of the *levels* of the regressors dated t-2 and earlier as instruments for changes in the lagged values of the

that hosts are, on average, skilled at incorporating local information that is informative of future customer performance into their comp decisions, the estimate of β_2 shown in column 1 is positive and significant (coefficient=1.38; $p<.01$). The coefficient estimate demonstrates that each \$1 in comps invested in a host's portfolio of customers for the year yields \$1.38 of theoretical win in the next year.

The results in column 2 point to evidence of learning. The coefficient on *CompsxExperience* is positive and significant (coefficient=0.043; $p<.05$) consistent with the notion that hosts gain ability in acquiring and incorporating local information into their comp decisions as they gain experience. The coefficient estimate on *CompsxExperience* documents that the "return" on each \$1 in comps invested in a host's portfolio of customers in terms of future theoretical win increases by \$0.043 per year of host experience. Column (3) contains results from estimation of a version of equation (3) which allows the learning effect to vary for properties with tight vs. loose monitoring. The results show that all learning effects are concentrated in properties with loose monitoring. The coefficient estimate on *CompsxExperience* in the specification in this column (coefficient=0.089; $p<.05$) captures the learning effect for loose monitoring properties. This estimate shows that the "return" on each \$1 in comps invested in a host's portfolio of customers in terms of future theoretical win increases by \$0.089 per year of host experience in loose monitoring properties. The coefficient estimate on *CompsxExperiencexTightMonitoring* (coefficient=-0.083; $p<.05$) captures the differential learning effect for tight monitoring properties. This coefficient estimate suggests that any learning effects are essentially negated for properties with tight monitoring.

dependent variable. Standard specification tests for the Arrelano-Bond estimator, including tests of serial correlation and the Sargan test statistic, fail to reject the validity of these instruments when used in our sample.

Our second approach to estimating equation (3) is to aggregate data at the customer-host-property-year level. Exploiting customer-level data allows us to use alternate measures of experience to capture different types of learning.⁹ Specifically, we estimate a version of equation (3) in which experience is decomposed into general experience (*ExpGeneral*) measured as the cumulative number of all customer-trips handled by the host up to the start of the current year and customer-specific experience (*ExpSpecific*) measured as the cumulative number of trips handled by the host for a specific customer up to start of the current year. Both types of experience may be important. Employees have multiple opportunities to learn about the performance consequences of their discretionary decisions from general experience interacting across customers, but customer heterogeneity may limit the extent to which such learning is transferable across customer relationships. Conversely, employee experience interacting with specific customers should lead directly to better discretionary decisions as employees learn about the performance consequences of their decisions for those customers.

Before turning to estimation of equation (3) using the customer level observations, we first document a simple pattern in the data that is suggestive of learning in the decentralized information processing activities of hosts. Figure 1 illustrates how a measure of the relationship between future customer performance and current comp decisions varies with the general experience level of casino hosts. We measure the "return on comps" (*ROC*) for each customer-host-property-year observation as the total theoretical win for the customer at the property over the subsequent year divided by the dollar value of comps awarded to the customer by a host in the current year. To control for heterogeneity across properties and years, we then adjust this measure by subtracting its property-year level mean from each observation. We form experience

⁹Our objectives could also be accomplished using customer trip level data. However, aggregating data at the customer-year level allows us to avoid any econometric problems arising from overlapping observations in the dependent variable.

portfolios by splitting the sample into 100 quantiles based on *ExpGeneral* and then taking the mean level of the adjusted return-on-comps measure for each portfolio. Figure 1 provides evidence consistent with learning – the ratio of future performance to the current dollar value of comps increases as hosts gain general experience interacting with customers. Hosts in the lowest experience quantiles perform significantly worse than the average for each property-year and their performance does not tend to meet or exceed property-year average performance until their experience levels are in the 10th quantile and beyond.

Table 6 contains results from estimating equation (3) using the customer-host-property-year level data.¹⁰ The results in column 1 provide evidence that skill in acquiring and incorporating local information into comp decisions at least partially arises via learning through general experience interacting with customers. The interaction between *Comps* and *ExpGeneral* is positive and significant at the 1% level. Consistent with the results of the host-portfolio level analyses reported in Table 5, the results in column 2 of Table 6 show that learning effects from general experience are weaker for properties with tight monitoring. The coefficient on *CompsxExpGeneralxTightMonitoring* is negative and significant at the 1% level.

The results in Table 6 are not consistent with learning occurring via experience interacting with specific customers (*ExpSpecific*). Surprisingly, the coefficient on the interaction between *ExpSpecific* and *Comps* is negative and significant at the 1% level. On average, it appears that the quality of discretionary comp decisions declines as hosts gain experience with individual customers. There are at least two potential explanations for this result. First,

¹⁰ For our customer-host-level analyses, we utilize simple OLS with property, host, and year fixed effects rather than dynamic panel data estimation as in our host-portfolio-level tests reported in Table 5. The reason for this is that, holding the number of cross-sectional groups constant, the coefficient bias in fixed effect models with lagged dependent variables tends to 0 as the number of within-group observations increases (Wooldridge 2002). For our customer-host-level analyses, we have 361 unique hosts with an average of 636 customer-year observations per host. Thus, unlike our host-portfolio level dataset which contains at most 10 yearly within-host observations and is susceptible to biased estimates in the absence of dynamic panel data estimation, our customer-host-year level dataset is characterized by a relatively small number of groups (e.g. 361 hosts) each with a large number of within-group observations making this latter dataset ideally suited for estimation via standard fixed effects models.

customers may themselves be learning from repeated interaction about the organization's comp policies and how these policies relate to their own level and type of gaming activity. If this were the case, then customers may become more demanding of comp awards as they gain experience with a property and/or a specific host. In this scenario, we would expect the problem to be attenuated in properties with tight monitoring – where employees are less likely to deviate from decision guidelines and exacerbated in properties with loose monitoring – where employees are more likely to do so. The results in Table 6 provide mixed evidence that this is the case. The coefficient on *CompsxExpSpecificxTightMonitoring* is positive in all specifications, but is only significant in column (3) which excludes lagged theoretical win.¹¹

The second potential explanation for the negative coefficient estimate on the interaction between *ExpSpecific* and *Comps* is that this result reflects the attempts of hosts to dynamically manage the cumulative comp percentage awarded to a customer over time rather than the comp percentage awarded to a customer during an individual time period (e.g. individual trip or year). If hosts overcomp a customer on one trip, they may try to recoup the "investment" by limiting comp percentages on future trips. Similarly, in managing the expectations of repeat customers, hosts may attempt to generally limit comp percentages over time. Figure 2 provides some evidence that this is the case. This figure plots the cumulative comp percentage awarded by a host to a customer against the relationship-specific experience decile of the host. The cumulative comp percentage is defined for each customer-host-property-trip as the dollar value of all comps awarded to a customer by a host during all past trips divided by the total theoretical win for that

¹¹ Note that the estimates provided in column 3 of Table 6 exclude lagged theoretical win. The reason we include this column is to check whether the relatively low baseline "return on comps" (e.g. coefficient on lagged theoretical win) reported in columns 1 and 2 is an artifact of aggregating the data at the customer-year level. If comps awarded during the current year are associated with a persistent increase in theoretical win that arises later in the same year for a given customer (e.g. from repeat trips within the year), then this portion of the "return on comps" would be obscured by the inclusion of lagged theoretical win (particularly when current and lagged theoretical win are correlated as is clearly the case given the coefficients on lagged theoretical win in columns 1 and 2). The results in column 3 suggest that this is indeed the case as the baseline "return on comps" rises to approximately \$1.22 in the absence of lagged theoretical win as an additional explanatory variable.

customer over all past trips with the host. Relationship-specific experience deciles are formed by splitting the sample into deciles based on *ExpSpecific* and then taking the mean level of the cumulative comp percentage for each decile. Figure 2 demonstrates that comp percentages tend to be significantly higher during the customer's first trip with the host, but cumulatively, the comp percentage awarded gradually declines and levels off at approximately 40% -- the maximum comp percentage formally prescribed by the organization. The pattern that emerges in Figure 2 is one in which hosts dynamically manage the total comps awarded to individual customers towards the decision-guidelines prescribed by the firm.

In summary, the results in this section document several patterns in employee learning that are consistent with implicit incentive effects of tight versus loose monitoring. Hosts' comp decisions appear, on average, to be effective – each \$1 awarded in comps leads to more than \$1 in future theoretical win. Furthermore, consistent with learning, the effectiveness of comp decisions increases as hosts gain experience. However, learning effects are concentrated in “loose monitoring” properties and are almost entirely absent in “tight monitoring” properties.

IV. Discussion and Conclusion

We view our paper as among the first attempts to link two important elements of organizational design, delegation of decision-rights and the provision of incentives, to the behavior and human capital formation (e.g. learning) of lower-level employees. To explore these issues, we exploit a field-research setting in which: (1) employees' use of decision rights can be directly measured and traced to performance outcomes, and; (2) employees work in distinct business units which vary in their monitoring of decision-rights.

The first goal of our paper is to document how monitoring and enforcement of decision-rights gives rise to implicit incentives in the organization we study. There is, in general, a lack of evidence on the variety of incentive mechanisms operating at lower levels of different

organizations. With the exception of a small, but growing, number of studies on promotion and tournament-based incentives for lower level employees (e.g. Baker et. al. 1994, Gibbs 1995, Campbell 2008, Casas-Arce and Martinez-Jerez 2009) most of the studies on incentives for non-executive employees in the accounting literature have focused on short-term bonus plans based on formulaic or subjective weighting of performance metrics (e.g. Banker et. al. 2000, Ittner et. al. 2003, Gibbs et. al. 2004). In the organization we study in this paper, we document implicit incentives in the form of sharp increases in employee exits linked to “excessive” deviation from organizationally prescribed guidelines. We find that these implicit incentives are strongest for employees in business units which monitor these deviations most frequently (e.g. “tight monitoring”).

An interesting feature of our research setting is that decision-rights for lower level employees (hosts) are formally the same for all of this organization’s business units: comps awarded to customers as a percentage of the customer’s theoretical win on a given trip should not exceed 40% barring exceptional circumstances. However, the mechanisms through which these decision-rights are monitored and enforced vary in ways that create implicit incentives for employees to exercise these rights in more or less restrained ways. In a sense, this relates to Aghion and Tirole’s (1997) notion of formal versus real authority. Employees in our setting have effectively the same formal decision rights across business units, but their real authority is determined by the way in which superiors monitor and enforce these decision rights.

The second goal of our paper is to explore whether lower level employees appear to respond to the implicit incentives arising from the monitoring and enforcement of decision rights. We document two broad behavioral effects for employees in tighter monitoring environments: (1) they are less likely to use discretion to deviate from guidelines in general; (2)

they are less likely to incorporate local information, including historical performance data, when making local decisions. The first effect is straightforward: employees who are more likely to be penalized for mistakes are less likely to “experiment” (Lee et. al. 2004), preferring instead to manage towards explicit guidelines. The second effect shows that reluctance to experiment persists even though historical performance data is observable to organizational superiors and can be used to justify deviations from organizational decision rules.

The final, and perhaps most distinctive, goal of our paper is to identify the link between employee learning and implicit incentives from the monitoring and enforcement of decision-rights. This is not unrelated to the second goal discussed above. If employees in “tightly monitored” business units “experiment” less by deviating less often from explicit decision guidelines, then they will have fewer opportunities to learn how to effectively incorporate local information into their decisions. Consistent with this notion, we find strong learning effects among lower level employees in our setting with these learning effects concentrated among employees in “loosely monitored” business units and almost entirely absent in “tightly monitored” business units.

In addition to the obvious caveats related to the generalizability of a field-study, we acknowledge that the proxy used in this paper to classify business units in terms of “tight” versus “loose” monitoring is based on a limited amount of data. We have attempted to combine both qualitative (e.g. interview) and quantitative (e.g. descriptive statistics on monitoring frequency and organizational departure regressions) data to validate our classification of business units. However, it remains for us and for future researchers to develop stronger proxies to capture implicit incentives arising from the monitoring and enforcement of decision rights in organizations.

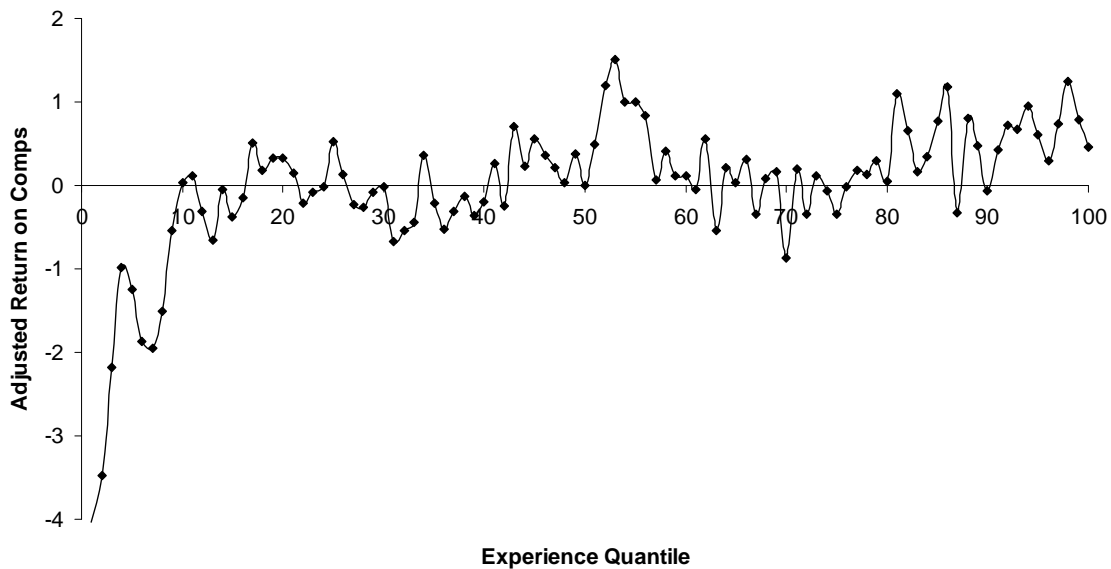
References

1. Abernethy, M.A., Bowens, J. and L. van Lent. 2004. Determinants of Control System Design in Divisionalized Firms. *The Accounting Review* 79 (3): 545-570.
2. Aghion, P. and J. Tirole. 1997. "Formal and Real Authority in Organizations," *Journal of Political Economy*, University of Chicago Press, vol. 105(1), pages 1-29,
3. Anderson, T.W. and C. Hsiao (1982), "Formulation and Estimation of Dynamic Models Using Panel Data," *Journal of Econometrics*, 18: 47-82.
4. Arrelano, M. and S. Bond (1991), "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economic Studies*, 58: 277-297.
5. Baiman, S., Larcker, D., and Rajan, M. 1995. Organizational Design for Business Units. *Journal of Accounting Research*. 33 (2): 205-229.
6. Baker, G., M. Gibbs., and B. Holmstrom. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data," *Quarterly Journal of Economics*, November, 881-920.
7. Banker, R. D., G. Potter, and D. Srinivasan. 2000. An Empirical Investigation of an Incentive Plan that Includes Nonfinancial Performance Measures. *The Accounting Review* 75 (1).
8. Bowman, D., and D. Narayandas. 2004. "Linking Customer Management Effort to Customer Profitability in Industrial Markets." *Journal of Marketing Research* 41 (4).
9. Campbell, D., Martinez-Jerez, A., and Epstein M. "Slots, Tables, and All that Jazz: Managing Customer Profitability at the MGM Grand Hotel." Harvard Business School Case #109-029.
10. Campbell, D. 2008. "Nonfinancial Performance Measures and Promotion-Based Incentives". *Journal of Accounting Research*. 46 (2).
11. Campbell, D., Datar, S., and Sandino, T. "Organizational Design and Control across Multiple Markets: The Case of Franchising in the Convenience Store Industry." *The Accounting Review*. Forthcoming.
12. Casas-Arce, P., and Martinez-Jerez, A. 2009. "Relative Performance Compensation and Dynamic Incentives." *Management Science*. Forthcoming.
13. Casas-Arce, P., Martinez-Jerez, A. and Narayanan V.G. 2009. "Customer Lifetime Value Metrics and Front Line Decision Making at Bankinter." Harvard Business School Working Paper.
14. Christie, Andrew A., Marc P. Joye, and Ross L. Watts (2003): Decentralization of the Firm: Theory and Evidence." *Journal of Corporate Finance* 9, 3-36
15. Gibbs, M. 1995. "Incentive Compensation in a Corporate Heirarchy." *Journal of Accounting and Economics*. 19: 247-277.
16. Gibbs, M., K.A. Merchant, W.A. Van der Stede, and M.E. Vargus 2004. "Determinants and Effects of Subjectivity in Incentives." *The Accounting Review* 79(2): 409-436.
17. Gupta, S. and V. Zeithaml. 2006. "Customer Metrics and Their Impact on Financial Performance." Special 25th Anniversary Issue. *Marketing Science* 25, no. 6: 718-739.
18. Huber, G.P. 1991. Organizational Learning: The Contributing Processes and the Literatures. *Organization Science*. 2(1): 88-115.
19. Ittner, C.D., D.F. Larcker, and M. Meyer. 2003. Subjectivity and the Weighting of Performance Measures: Evidence from a Balanced Scorecard *The Accounting Review* 78(3): 725-758.

20. Lapre, M., and Tsikriktsis, N. 2006. Organizational Learning Curves for Customer Dissatisfaction: Heterogeneity Across Airlines. *Management Science*. 52(3): 352-366.
21. Lee, F., A. Edmondson, S. Thomke, and M. Worline. 2004. "The Mixed Effects of Inconsistency on Experimentation in Organizations." *Organization Science* 15, no. 3: 310-326.
22. Merchant, K. 1985. "Organizational Controls and Discretionary Program Decision-Making: A Field-Study." *Accounting, Organizations, and Society*. 10(1): 67-85
23. Nagar, V. 2002. Delegation and Incentive Compensation. *The Accounting Review*. 77 (2): 379-395.
24. Reinartz, W. and V. Kumar. 2000, "On the Profitability of Long-Life Customers in a Noncontractual Setting: An Empirical Investigation and Implications for Marketing," *Journal of Marketing*, 64 (October), 17-35.
25. ——— and ——— 2003. "The Impact of Customer Relationship Characteristics on Profitable Lifetime Duration," *Journal of Marketing*, 67 (January), 77-99.
26. Wooldridge, J. 2002. *Econometric Analysis of Cross-Section and Panel Data*. Cambridge, MA: The MIT Press.

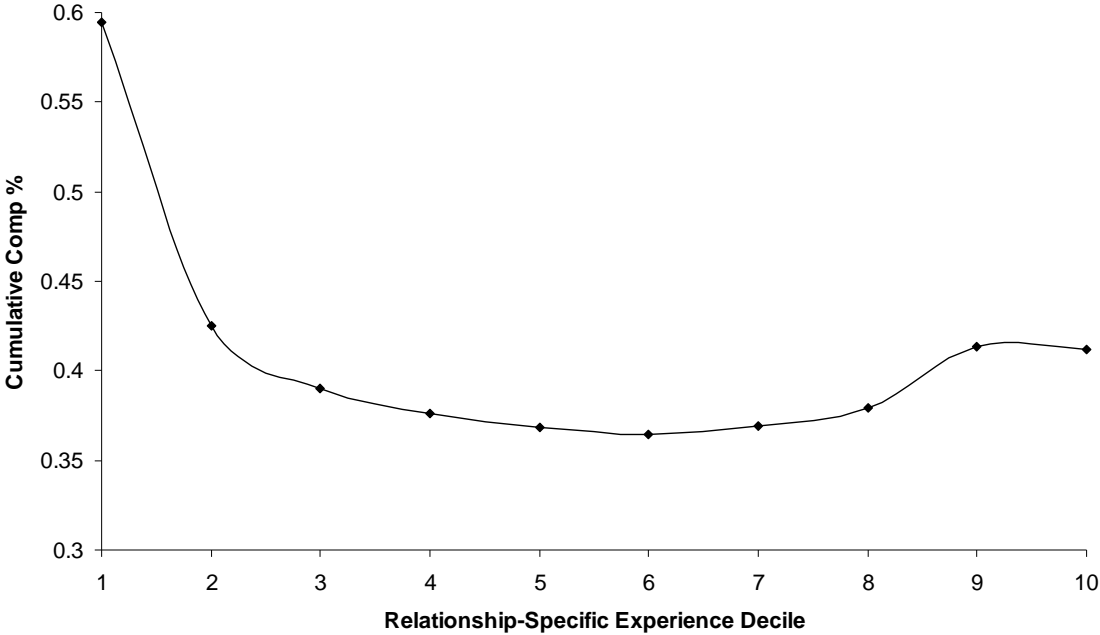
Figure 1

Adjusted "Return-on-Comps" Across Experience Quantiles*



*Experience quantiles based on experience measured as cumulative number of trips assigned to a host; Adjusted return-on-comps measured as $ROC_{itj} - \overline{ROC}_{yj}$, where ROC_{itj} denotes return on comps for customer 'i' on trip 't' at property 'j' and \overline{ROC}_{yj} denotes the mean for ROC across all customer-host trips in year 'y' at property 'j'.

Figure 2
Cumulative Comp Percentage across Relationship Specific Experience Deciles*



*Relationship-specific experience deciles based on experience measured as cumulative number of trips with a specific customer assigned to a host; Cumulative comp % is defined as the total dollar value of comps awarded to a customer by a host over all past trips with the customer divided by the theoretical win of that customer over all past interactions with the host.

Table 1
Property and Host Characteristics

	Properties with:																	
	Tight Monitoring									Loose Monitoring								
	Property 1			Property 2			Property 3			Property 4			Property 5			Property 6		
	Mean	Med.	S.D.	Mean	Med.	S.D.	Mean	Med.	S.D.	Mean	Med.	S.D.	Mean	Med.	S.D.	Mean	Med.	S.D.
Trips per Host	456.7	131	807.5	509.3	140	1043.3	217.9	87	469.9	424.7	133	941.9	37.2	31	82.1	283.6	99	1209.9
Theoretical Win Per Trip	1521.4	1058.9	1491.4	1007.8	627.1	1155.6	423.3	121.5	819.1	868.7	484.6	1108.3	1247.8	853.3	1241.7	1276.1	783.2	1451.3
Data Range	1993-2005			1997-2005			1999-2005			1993-2005			1993-2005			1998-2005		
Number of Unique Hosts	62			20			95			62			39			83		
Number of Host-Years	594			131			445			390			269			422		
Number of Host Exits	16			7			13			22			8			20		
Frequency of Comp Exception Reviews	Daily and Monthly			Daily and Monthly			Daily and Monthly			Weekly and Quarterly			Monthly			Weekly and Quarterly		

Table 2
The Exit-Performance Relation for Properties with Tight vs. Loose Monitoring

	Dependent Variable: <i>EXIT</i>	
	(1)	(2)
<i>Constant</i>	-2.726** (1.15)	-2.612** (1.15)
<i>CustomerGrowth</i>	-1.483*** (0.44)	-1.607*** (0.53)
<i>TripsPerCustomerGrowth</i>	0.39 (0.44)	0.377 (0.42)
<i>TheoreticalWinPerTrip Growth</i>	-0.092** (0.04)	-0.105** (0.05)
<i>Discretion %</i>	-2.186 (1.53)	-2.065 (1.47)
<i>Overcomped</i>	1.424*** (0.44)	1.275*** (0.46)
<i>ExcessComps</i>	0.001** (0.00052)	0.001** (0.00056)
<i>ExcessComps x TightMonitoring</i>		0.004*** (0.00145)
<i>Experience</i>	-1.037*** (0.17)	-1.010*** (0.18)
Property Fixed Effects	+++	+++
Year Fixed Effects	+++	+++
Pseudo R-Squared	0.41	0.42

Implied Probabilities

<i>Overcomped</i> =0, and all other variables at mean	0.003
<i>Overcomped</i> =1; <i>ExcessComps</i> in 90th percentile; Loose Monitoring Property	0.013
<i>Overcomped</i> =1; <i>ExcessComps</i> in 90th percentile; Tight Monitoring Property	0.063

Standard errors in parentheses are adjusted for clustering of observations within hosts over time; * significant at 10%; ** significant at 5%; *** significant at 1% ; +++ denotes jointly significant at the 1% level using χ^2 test; Table reports logit estimates of equation (1) using data on 1,189 host-year observations; *Exit*=1 if host departs from a property in the subsequent year, 0 otherwise; *CustomerGrowth* = annual growth in the number of customers managed by the host over the prior year; *TripsPerCustomerGrowth* = annual growth over the prior year in the average number of trips taken to a property by customers managed by the host; *TheoreticalWinPerTripGrowth* = annual growth over the prior year in the average theoretical win per trip for all customers managed by the host; *Discretion%*=percentage of all customer-trips managed by the host during the year in which comps were awarded in excess of 40% of the trip-level theoretical win; *Overcomped*=1 if total comps awarded by the host to all customers in a given year is greater than 40% of the aggregate theoretical win across all customers managed by the host for that year, 0 otherwise; *ExcessComps*= [100*(total comps awarded by the host to all customers in a given year divided by the aggregate theoretical win across

all customers managed by the host for that year)-40] when *Overcomped*=1 and 0 otherwise. *TightMonitoring*=1 if host is employed at properties 1, 2, or 3 and equals 0 otherwise; *Experience* = number of years of host experience at property as of the start of the year.

Table 3
Use of Decision-Rights for Properties with Tight vs. Loose Monitoring

	Properties with:			t-test for Difference
	All Properties	Tight Monitoring	Loose Monitoring	
<i>Discretion %</i>	23.7 (21.8)	19.6 (19.3)	29.2 (23.8)	4.43***
<i>Overcomped</i>	0.235 (0.424)	0.139 (0.346)	0.368 (0.482)	7.35***
<i>Comp %</i>	45.8 (1.2)	35.6 (1.0)	59.8 (1.4)	3.21***

Table reports mean for each host-year level variable across 2,251 host-year observations; Standard deviations in parentheses; *** significant at the 1% level; *Discretion%*=percentage of all customer-trips managed by the host during the year in which comps were awarded in excess of 40% of the trip-level theoretical win; *Overcomped*=1 if total comps awarded by the host to all customers in a given year is greater than 40% of the aggregate theoretical win across all customers managed by the host for that year, 0 otherwise; *Comp%*= total comps awarded by the host to all customers in a given year is divided by the aggregate theoretical win across all customers managed by the host for that year.

Table 4

Determinants of the Trip-Level Comp Decision for Properties with Tight vs. Loose Monitoring

	Dependent Variable:		
	Comps		
	1	2	3
<i>TheoreticalWin</i>	0.218*** (0.002)	0.259*** (0.004)	0.259*** (0.004)
<i>TheoreticalWin x TightMonitoring</i>		-0.061*** (0.005)	-0.061*** (0.005)
<i>LagTheoreticalWin</i>	0.011*** (0.001)	0.015*** (0.002)	0.015*** (0.002)
<i>LagTheoreticalWin_{t-k} x TightMonitoring</i>		-0.007*** (0.002)	-0.007*** (0.002)
<i>Overcomped</i>			49.951** (20.008)
<i>Overcomped x TightMonitoring</i>			-36.667* (21.100)
<i>ExcessComps</i>			1.745 (1.129)
<i>ExcessComps x TightMonitoring</i>			-2.635** (1.209)
Host Fixed Effects	+++	+++	+++
Site Fixed Effects	+++	+++	+++
Year Fixed Effects	+++	+++	+++
Number of Host-Customer-Trips	220,223	220,223	220,223
R-Squared	0.41	0.42	0.42

Standard errors in parentheses are adjusted for clustering of observations within customers over time; * significant at 10%; ** significant at 5%; *** significant at 1% ; +++ denotes jointly significant at the 1% level using χ^2 test; Table reports OLS estimates of equation (2') using host-customer-trip level data; *TheoreticalWin*=theoretical win generated by the customer on the current trip; *LagTheoreticalWin* = Cumulative theoretical win generated by the customer over the 18 months prior to the current trip start-date; *Overcomped* and *ExcessComps* are measured at the host-year level and are defined in the notes to Table 2. *TightMonitoring*=1 if host is employed at properties 1, 2, or 3 and equals 0 otherwise.

Table 5
Learning and the Return on Comps for Properties with Tight vs. Loose Monitoring

	Theoretical Win		
	1	2	3
<i>Theoretical Win</i> _{t-1}	0.062 (0.045)	0.059 (0.045)	0.06 (0.045)
<i>Comps</i> _{t-1}	1.383*** (0.152)	1.241*** (0.164)	1.370*** (0.167)
<i>Comps</i> _{t-1} <i>x</i> <i>Experience</i> _{t-1}		0.043** (0.017)	0.089*** (0.023)
<i>Comps</i> _{t-1} <i>x</i> <i>Experience</i> _{t-1} <i>x</i> <i>Tight Monitoring</i>			-0.083*** (0.025)
<i>Experience</i> _{t-1}	110.353*** (18.603)	112.739*** (18.536)	91.853*** (19.060)
Year Indicators	+++	+++	+++
Number of Host-Years	1,720	1,720	1,720
Number of Unique Hosts	324	324	324

Standard errors in parentheses are adjusted for clustering of observations within hosts over time; * significant at 10%; ** significant at 5%; *** significant at 1% ; +++ denotes jointly significant at the 1% level using χ^2 test; Table reports Arrelano-Bond dynamic panel data estimates of equation (3) using host-year level data; *TheoreticalWin*=aggregate theoretical win across all customers managed by a host during the year; *Comps*=aggregate comps awarded by host to all customers managed by that host during the year; *TightMonitoring*=1 if host is employed at properties 1, 2, or 3 and equals 0 otherwise; *Experience* = number of years of host experience at property as of the start of the year.

Table 6
General and Specific Learning for Properties with Tight vs. Loose Monitoring

	<i>TheoreticalWin_{t+1}</i>		
	(1)	(2)	(3)
<i>TheoreticalWin_t</i>	0.460*** (0.0060)	0.458*** (0.0060)	
<i>Comps_t</i>	0.254*** (0.0200)	0.240*** (0.0210)	1.223*** (0.0200)
<i>Comps_t x ExpGeneral</i>	0.0004*** (0.0001)	0.0006*** (0.0002)	0.002*** (0.0002)
<i>Comps_t x ExpSpecific</i>	-0.003*** (0.0010)	-0.003*** (0.0010)	-0.007*** (0.0010)
<i>Comps_t x ExpGeneral x TightMonitoring</i>		-0.001*** (0.0003)	-0.003*** (0.0003)
<i>Comps_t x ExpSpecific x TightMonitoring</i>		0.0001 (0.0010)	0.003** (0.0010)
<i>ExpGeneral</i>	4.757*** (0.2810)	4.857*** (0.2840)	7.132*** (0.3390)
<i>ExpSpecific</i>	27.910*** (1.5980)	27.698*** (1.6000)	51.718*** (2.2840)
Host Fixed Effects	+++	+++	+++
Site Fixed Effects	+++	+++	+++
Year Fixed Effects	+++	+++	+++
Observations	229,861	229,861	229,861
R-squared	0.34	0.34	0.25

Standard errors in parentheses are adjusted for clustering of observations within customers over time; * significant at 10%; ** significant at 5%; *** significant at 1% ; +++ denotes jointly significant at the 1% level using χ^2 test; Table reports OLS estimates of equation (4) using customer-host-year level data; *TheoreticalWin*=customer's total theoretical win during the year; *Comps*=total comps awarded by the host to the customer during the year; *ExpGeneral*=Cumulative number of trips managed by the host up to the start of the year; *ExpSpecific*=Cumulative number of trips for a specific customer managed by the host up to the start of the year; *TightMonitoring*=1 if host is employed at properties 1, 2, or 3 and equals 0 otherwise.