Charles University
Center for Economic Research and Graduate Education
Academy of Sciences of the Czech Republic
Economics Institute

CERGE-EI

Discussion Paper Series

THE IMPACT OF FINANCIAL INCENTIVES ON TASK PERFORMANCE: THE ROLE OF COGNITIVE ABILITIES & INTRINSIC MOTIVATIONLABOUR

Ondrej RYDVAL

Discussion Paper No. 2003 – 112

December 2003

P.O. Box 882, Politických vězňů 7, 111 21 Praha 1, Czech Republic
http://www.cerge-ei.cz

THE IMPACT OF FINANCIAL INCENTIVES ON TASK PERFORMANCE:

THE ROLE OF COGNITIVE ABILITIES & INTRINSIC MOTIVATION

Ondrej Rydval*

CERGE-EI

December 2003

Abstract

Economists widely assume that financial incentives represent the dominant and effective stimulator of human productive activities. Yet this assumption may be wrong in the case of mental production processes. This is because intrinsic motivation as well as financial stimuli foster cognitive effort, and because the impact of effort on performance is also moderated by individual cognitive abilities. My ultimate aim is three-fold: first, to formalize a model of cognitive production that would incorporate these insights; second, to test the derived predictions econometrically on laboratory and field data; and third, to apply the gained insights in a broader economic context. This paper provides a literature review necessary as a starting point for carrying out the proposed research.

* CERGE-EI, Charles University, Academy of Sciences, POB 822, 111 21 Prague, Czech Republic, Email: Ondrej.Rydval@cerge-ei.cz

While preparing this paper, I have benefited from the guidance of Prof. Andreas Ortmann, and from the World Bank Second Year Research Fellowship for the year 2003.

1. Introduction

Economists widely assume that financial incentives represent the dominant stimulator of human productive activities. In mental production tasks, however, the underlying mechanism through which incentives affect human behavior is relatively unclear. We have limited knowledge of how financial stimuli interact with intrinsic motivators in inducing cognitive effort, and we know even less about how effort combines with cognitive abilities in determining performance.

As an illustration, consider a classical agency model of work compensation (e.g. Milgrom and Roberts, 1992). The manager's goal is to induce a worker's effort by designing an optimal effort compensation scheme. Doing so may be complicated by moral hazard problems associated with unobserved effort and random shocks to output, but provided that such problems are solved, incentives are normally assumed to map smoothly into effort and performance. As explained below in more detail, at least two qualifications are in order. First, incentives might induce the worker's effort but at the same time "crowd out" her intrinsic satisfaction from performing the task, so the overall effort level might not change much. And second, even if incentives do stimulate her effort, the worker's performance might not respond much if she does not possess some of the skills necessary for completing the assigned task.[1]

This paper presents a starting point for further exploring these issues both theoretically and empirically by providing a thorough literature review of the issues at stake. The forthcoming theoretical work will focus on formalizing the cognitive production process and in turn on deriving predictions for incentives effects therein. This will involve extending the

[1] Admittedly, the issues mentioned so far have to some extent been addressed in economics, psychology and other fields. For example, economic research on incentives in organizations deals

existing labor theory of cognition (see below) for two features: the impact of intrinsic motivation on effort, and the role of cognitive abilities in production. My empirical work will concentrate on generating experimental data to estimate the cognitive production function, and to test how individual cognitive abilities moderate the impact of incentives on performance.

Since incentives lie at the heart of economics, this research direction promises a wide spectrum of theoretical and practical applications. The most immediate contribution will be to the current methodological debate about the use of incentives in experiments and – through the "omitted variable" cognitive capital – the internal as well as external validity of experimental results. As illustrated above, my research also closely ties into the literature on the use of incentives in organizations. Furthermore, better understanding of incentive effects and ability constraints in cognitive production should be informative for several areas of public policy such as funding of higher education.[2]

This paper is organized as follows. In the next section I attempt to embed the research on incentive effects in cognitive tasks in a broader context, pointing out parallels between economics and psychology research. I also briefly look at related literatures that have the cognitive production process in their core but have never modeled it explicitly. Section 3 presents a powerful empirical illustration of why considering intrinsic motivation issues and especially the cognitive ability constraint is important. Section 4 provides a more thorough theoretical and empirical background for my research, organizing the review of existing literature around the so called capital-labor-production framework (henceforth KLP

---

with skill differences among workers and skill acquisition (see Gibbons, 1998, for a survey). Below I discuss in detail why the treatment of these issues has been inadequate.

[2] Labor economists have long struggled with separating private returns to education and work experience (in terms of earnings) from returns to cognitive ability, which has partly precluded the determination of an optimal level of educational funding. See later Section 6 for details.

framework, or KLP). Section 5 outlines the research road ahead. The concluding section elaborates on the above mentioned applications of the proposed research.

2. Embedding the proposed research in a broader context

The effect of financial incentives on human behavior has received widespread attention in experimental economics and neuro-biology[3], as well as in the literature on provision of incentives in organizations.[4] Most informative about the impact of incentives on cognitive performance seems the debate among experimentalists.[5] Economists widely use performance-based rewards in their belief that incentives are necessary to stimulate subjects' cognitive effort which in turn ensures that decision errors are largely avoided and performance is measured reliably. In contrast, psychologists usually pay participants credit points or flat fees, arguing that their intrinsic motivation is sufficient, and that imposing financial incentives can distract or even crowd out a priori intrinsic goals.

This debate has motivated but also sharply divided both theoretical and empirical work on incentive effects. On the one hand, psychologists have concentrated on modeling how extrinsic financial incentives might crowd out intrinsic motivation of subjects (e.g. Deci and Ryan, 1985; see Frey and Jegen, 2001, for a survey), and have conducted abundant experimental studies to examine this phenomenon. The resulting evidence concerning

[3] Neuro-biologists attempt to model and detect the neurophysiological processes inside the brain that underlie human decision processes. Gold and Shadlen (2001), for example, suggest that the neural computations – based on information (including incentives) from various sensory stimuli – are organized in a simple likelihood ratio decision framework. Cognitive neuroscience has recently developed methods of brain imaging that allow us to localize neural activities inside the brain associated with various cognitive processes (e.g. positron emission tomography, functional magnetic resonance imaging, etc.). See Anderson (2000) and Glimcher (2003) for an overview.

[4] Gibbons (1998) provides a survey of agency theory which forms a basis of most theoretical work on incentives in organizations. See also Benabou and Tirole (2003) and Acemoglu et al. (2003) for recent theoretical developments. Prendergast (1999) provides a more empirical-oriented assessment of incentive provision in firms.

[5] See Hertwig and Ortmann (2003) and Camerer (2003) for a discussion of experimental practices regarding incentives.

crowding-out effects gathered in several meta-analytical surveys seems controversial and largely dependent on the nature of tasks included (see Deci et al., 1999, and the ensuing comments; see also Eisenberger and Cameron, 1996).

Economists, on the other hand, have focused on modeling the tradeoff between effort and incentives in cognitive production to justify their widespread use of incentives in experiments. The "labor theory of cognition", advanced independently by Smith and Walker (1993) and Wilcox (1993), viewed individuals as choosing how much cognitive effort to exert to optimally balance monetary benefits accruing from cognitive production against the corresponding decision effort costs.[6] The accompanying empirical work has sought to establish a positive link between incentives, effort (mostly measured by decision time) and performance. As a common finding, effort usually does increase with incentive level, but performance improvement follows only in certain decision-making environments.[7]

That incentives indeed work better for some cognitive tasks than others has become apparent from recent meta-studies and empirical surveys which have accumulated evidence on incentive effects from both experimental psychology and economics.[8] What they show is that, although the record is mixed and inconclusive, *the magnitude of incentive effects seems*

[6] But see also Conlisk (1988) for an earlier model with effort decision cost, and Smith and Szidarovszky (2001) for an extension of the labor theory to strategic interactions rather than games against nature. To complete the picture, Harrison (1989 and 1992), in discussing the problems associated with flat maxima of the experimental payoff function, essentially also thinks in terms of a labor-theoretical decision framework.

[7] Econometrically sound tests of the labor theory in the context of lottery choices were performed by Wilcox (1993) and subsequently by Moffatt (2003).

[8] Notice that throughout my work, I mostly turn to evidence on incentive effects from laboratory experiments. This is because analyzing the incentive-performance relationship based on field data may suffer from serious drawbacks. Prendergast (1999) in his critical review of incentive provision in firms reports, among other pitfalls, little control of environmental variables and confusion over what constitutes "compensation". Nevertheless, finding the right balance between laboratory and field evidence will be an ongoing concern: although experimentalists can better control environmental variables, they often do so at the expense of the external validity of their results. Cosmides and Tooby (1994) further argue that the typical decontextualized ways in which experiments are conducted limit participants' access to their "inference machines".

to depend in a complicated fashion on the nature of the assigned cognitive task.[9] This empirical regularity is currently the bone of contention between economists and psychologists. Importantly, it also implies that the labor theory, as it stands, is an unsatisfactory description of cognitive production and incentive effects therein.

The first articulated though informal explanation of why the effect of incentives might depend on the nature of the task was offered by Camerer and Hogarth (1999). The authors' capital-labor-production (KLP) framework extends the labor theory in two dimensions. First, cognitive effort is no longer assumed to depend only on financial stimuli, but also on intrinsic motivation of individuals. And second, the cognitive production function features an additional production input – cognitive capital. Individuals maximize an objective function, consisting of both financial and intrinsic goals, by employing optimal level of effort in cognitive production. The productivity of this effort in turn depends on the availability of cognitive capital relevant to the assigned task, and on the substitutability between the two production inputs.

Camerer and Hogarth do not attempt to model their KLP framework theoretically, but rather describe verbally how it allows average incentive effects to depend on the nature of the task, and cite empirical studies that seem to support their claims. In mundane tasks, the authors argue, incentives are likely to work well since there are no intrinsic motivation problems and only few skills are required. However, a *performance constraint* arises if production is limited from above by the design of the task. Similar situation may arise in more interesting but still skill-wise manageable tasks where intrinsic motivation stimulates

[9] Surveys and meta-analyses were conducted by Bonner et al. (2000, 2002), Camerer and Hogarth (1999), Hertwig and Ortmann (2001), Jenkins et al. (1998), and Prendergast (1999). Building upon Camerer and Hogarth (1999) and their own survey paper, Hertwig and Ortmann (2003) argue that, ironically, financial incentives empirically matter more in tasks that psychologists perform ('judgment and decision' studies) than in those mostly conducted by economists ('game and market' situations).

effort enough to reach the performance constraint. In both cases, the effect of extra incentives on performance may turn out small. More interesting cases arise in complex tasks. These are presumably intrinsically motivating and thus more likely to involve crowding-out motivational problems; at the same time, such tasks require more advanced skills which individuals might not possess and so an *ability constraint* is likely to limit the productivity of effort.[10]

The KLP framework provides a more realistic description of cognitive production than the labor theory, and is able to explain why the size of incentive effects may vary across cognitive tasks. To the best of my knowledge, however, no study has formalized these ideas. Doing so presents one of the goals of my work, and will clearly require nesting several related yet unconnected literatures within a common framework.

In particular, formalizing the KLP framework will extend the labor theory in two important aspects. First, one needs to consider more carefully the relationship between extrinsic (financial) and intrinsic stimulators of effort. To this end, however, the insights offered by psychologists do not seem sufficient. Kreps (1997), for example, points out that the psychological crowding-out theories are trivial in that the predicted effect of financial incentives on intrinsic motivation follows directly from assumptions imposed on the utility function. The author in turn discusses different possible interactions of extrinsic and intrinsic motivational factors associated with people's adherence to social norms. Related to this, Benabou and Tirole (2003) argue that incentives present a weak stimulant if they negatively affect agent's perception of the task or of her own abilities.

---

[10] See Kahneman et al. (1968) and Libby and Lipe (1992) for illustrations of the performance constraint in memory task experiments. See Awasthi and Pratt (1990) for an early empirical investigation of the cognitive ability (or cognitive capital) constraint in the context of problem-solving accounting tasks. Details are available in Section 4 below.

Second, the KLP framework needs to incorporate the literature on the ability-performance relationship which seems relatively under-researched in economics. A rare exception is a study by Devetag and Warglien (2003) who report a significant positive association between individual short-term memory capacity and overall performance in three game-theoretic tasks.[11] Nevertheless, more useful insights about cognitive abilities can be learned from cognitive psychology (e.g. Anderson, 2000, and Cowan, 2001, and Koriat and Goldsmith, 1996a). Further, the emerging neuro-economics field incorporates the existing evidence on neural processes into economic models of individual behavior (e.g. Glimcher, 2003). However, none of these fields considers the impact of incentives on the way that cognitive tools are applied. I will provide this link by addressing both motivational and ability-related concerns in the KLP framework.

As a last point, I will mention a couple of economic literatures that are closely related to the labor theory of cognition. These have examined incentive effects in real-world settings and have implicitly worked with the cognitive production process, but never modeled it explicitly. First, as illustrated above, the industrial organization literature on work compensation has examined the link between incentives and (cognitive) effort and attempted to incorporate intrinsic motivation issues (e.g. Frey, 1997), but apart from considering worker productivity differences in signaling models has largely failed to take into account cognitive abilities. Second, the dominant economic theory of wage determination – the human capital theory (see e.g. Mincer, 1997) – has focused on the relationship between work compensation and cognitive abilities, but has left out intrinsic motivation and cognitive effort issues. In the concluding section, I will discuss what these literatures can learn from extending the labor theory by formalizing the KLP framework.

[11] Admittedly, looking at the three tasks separately, however, the authors find positive association for only one of them. Further, as discussed in Section 4, their memory measure is problematic.

3. Empirical illustration of cognitive production and incentive effects

Whereas a majority of available evidence on incentive effects only reports the *average* effect of incentives on effort and performance, and Camerer and Hogarth (1999) also outline the KLP predictions in this fashion, both the labor theory and the KLP framework are structured in terms of *individual* optimization problem. Indeed, looking at individual-level data is likely to reveal significantly more about the validity of the KLP framework. After all, even if the nature of the task determines the size of incentive effects, it will clearly be *conditional* on the nature of the subject pool. And arguably, the KLP predictions will be driven by individual heterogeneity, especially in terms of cognitive abilities. I illustrate below that this in fact seems to be confirmed empirically in the case of a very general cognitive task, which in turn justifies the modeling of individual incentive effects and their further exploration.[12]

In the figure below, I adapted individual-level data on performance in a psychometric (IQ) test from Gneezy and Rustichini (2000). The authors examined the impact of financial incentives on average IQ score by dividing 160 individuals into four incentive treatments: no-pay, NIS0.1, NIS1 and NIS3.[13] They reported a non-monotonic impact of incentives on performance because average performance was lowest in the NIS0.1 treatment, i.e. even lower than in the no-pay treatment. The authors explained this by arguing that the NIS0.1 subjects were discouraged by the low compensation offered and consequently performed on

[12] Unfortunately, it is not always useful for the purpose of discussing the validity of the KLP framework to retrieve individual-level data from past experiments. Especially in game-theoretic tasks, performance is rarely a continuous variable and is often affected by various chance factors such as interaction with other randomly assigned individuals. See Section 4 where I discuss other deficiencies of past studies reporting incentive effects.

[13] Subjects were also paid a flat participation fee of NIS60. At the time of the experiment, NIS3.5=$1. The fifty IQ-type questions, taken from a test normally used to scan university applicants, involved mainly reasoning and computation skills. The subjects were volunteer male and female undergraduate students at the University of Haifa from all fields of study with an average age of 23 years.

average even worse than the no-pay subjects who apart from the participation fee solved the tasks solely based on their intrinsic motivation. The authors further reported that average performance differed very little between the two high-incentive treatments (NIS1 & NIS3), but was significantly higher than in the two low-incentive treatments (no-pay & NIS0.1).

Figure 1: Illustrating cognitive-capital effects on individual-level data

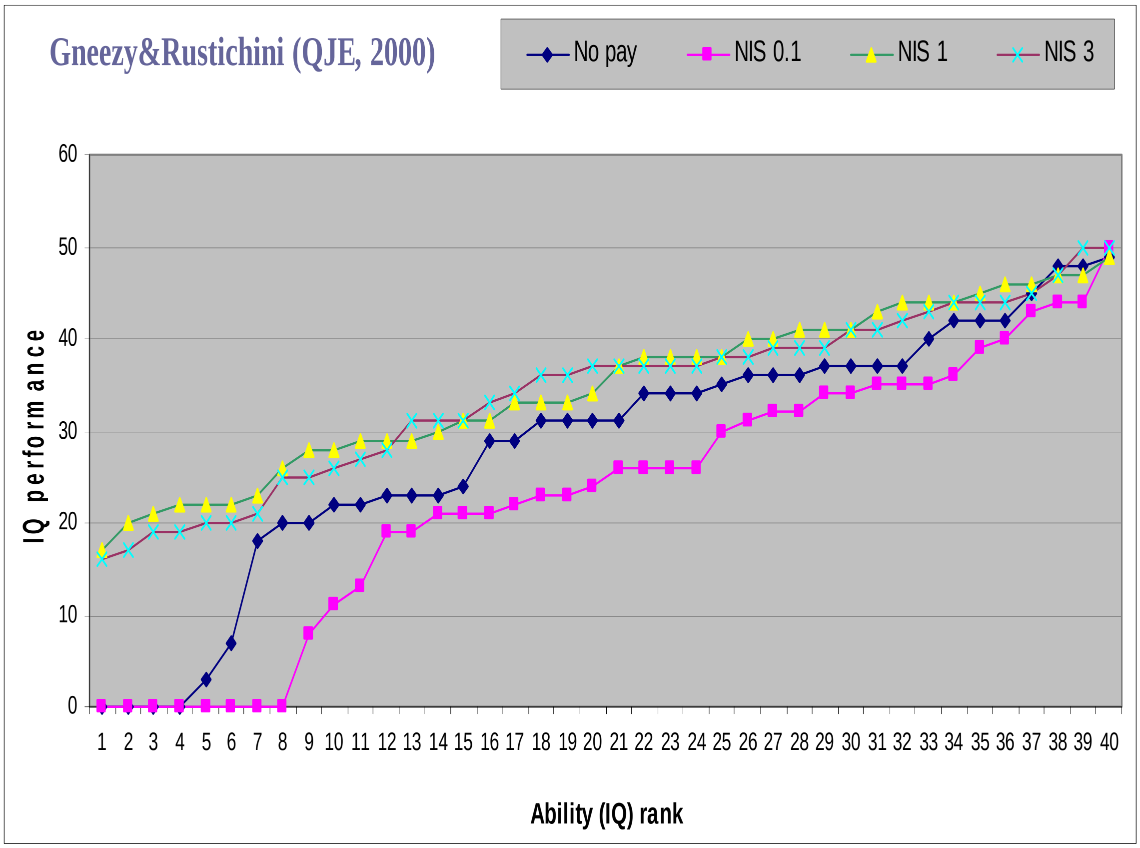

I argue that if one interprets the IQ scores as a broad measure of individual cognitive ability as well as of performance, then plotting the data in ascending order of IQ score for each of the four treatments reveals much more than this. In the figure, connecting all IQ-score observations for a given incentive treatment forms a "performance curve" for that treatment which visually describes the *within-treatment* variation in performance. In contrast, by making comparisons among the four "performance curves", one can inspect the *across-*

treatment variation in performance. First of all, notice that for the two high-incentive treatments (NIS1 & NIS3), the “performance curves” are virtually identical and slope considerably upwards. In other words, there is considerable within-treatment variation in performance but hardly any across-treatment one. I claim that this is most likely to be due to a significant (and similar) within-treatment variation in cognitive abilities. The reason is simple: the huge within-treatment variation in cognitive effort that would be required to generate this result is unlikely, and one would still have to explain why the two “performance curves” are almost identical despite the across-treatment incentive differential. Therefore, predominantly ability differentials seem to determine performance differentials when incentives are high enough.

Next inspect the “performance curves” for the low-incentive treatments (no-pay & NIS0.1). Clearly, Gneezy and Rustichini were correct in arguing that the NIS0.1 subjects overall seem less motivated than the ones in the no-pay pool since the “performance curve” for the NIS0.1 treatment is below that for the no-pay treatment across the whole performance range. It is highly unlikely that this would be caused by across-treatment ability differences, and thus across-treatment differences in motivation and effort must have played the main role. In addition, however, we can see that the gap between the two “performance curves” widens at the low-performance end: the largest number of subjects that answered zero or very few questions correctly is clearly in the NIS0.1 treatment. That such behavior is much less prevalent in the no-pay treatment and entirely absent in the two high-incentive treatments (NIS1 & NIS3) seems to confirm the presence of motivational problems in the NIS0.1 treatment.

Finally and most importantly, focus on the slope of all four “performance curves” and the distance between them. An eyeball test reveals that, leaving aside the motivational problems at the low-performance end, the within-treatment variation in performance is generally much greater then the across-treatment one. To give a meaningful comparison, look at the largest across-treatment performance differential at the middle rank of 20. This turns out to be 13 (i.e., 24 correct answers in the NIS0.1 treatment vs. 37 in the NIS1 treatment), which is equivalent to the performance differential associated with moving up from the first to the third quartile within the NIS1 treatment (28 vs. 41). Note, however, that within-treatment performance differentials can be much larger. For instance, in both of the high-incentive treatments (NIS1 & NIS3), the difference in performance for individuals ranked 1 and 40 is as large as 34. Now, provided that the across-treatment performance variation can be assigned to incentive effects while the within-treatment performance variation to ability differentials, an important and powerful result follows: ability differential among individuals seem to account for a much greater part of performance variation than incentive effects.

This result is further confirmed and strengthened by the fact that all four "performance curves" meet at the high-performance end: that is, the most able subjects reach the maximum possible performance regardless of incentives. In fact, note that if Gneezy and Rustichini happened to have in the subject pool only the most able subjects, say the top half of subjects from each of the four treatments, they would have found very small and probably insignificant average incentive effects, unlike those that would be observed for the subjects from the bottom parts of the "performance curves". Further to the motivational complications, we admittedly do not know whether the disinterested subjects at the low-performance end of the no-pay and NIS0.1 treatments have low abilities: these might as well be high-ability individuals. But this makes the case for cognitive-capital effects even stronger since, were it not for motivational problems of (possibly) high-ability individuals, cognitive-capital effects could be even larger.

The above observations clearly illustrate that neglecting the effect of cognitive abilities is an important shortcoming of most existing empirical studies examining incentive effects.[14] At a theoretical level, the above observations taken together make a strong case for considering cognitive abilities, as well as motivational factors, as major determinants of cognitive performance. This was clearly a major deficiency of the labor theory of cognition, one which I will remedy in my future research by formalizing the KLP framework.

4. Interpreting existing literature in the light of the KLP framework

The powerful example in the previous section has documented the urgent need to reconsider the existing approaches to modeling and estimating incentive effects in cognitive tasks. Accordingly, modeling the impact of incentives on cognitive production processes in a systematic manner, by extending the labor theory of cognition for motivation- and especially ability-related aspects, will form the main goal of my future research. The initial step, however, naturally consists of assimilating and distilling the existing information about how the cognitive production world might look like, and how exactly incentives fit in. To this end, I have organized this section around the KLP framework advanced by Camerer and Hogarth (1999). Admittedly, the KLP framework is a rather imprecise and informal account of cognitive production, and there exist earlier accounts of the role of intrinsic motivation and cognitive abilities (e.g. Awasthi and Pratt, 1990, and Libby and Lipe, 1992). However, once its various parts are properly assembled and organized, the KLP provides useful structure

[14] Awasthi and Pratt (1990) who control for cognitive abilities are an exception. Contrary to what seems to hold for the data of Gneezy and Rustichini, the authors find that incentive effects are significantly stronger for subjects with higher cognitive abilities (measured by a perceptual differentiation test). At a more general level, Hannan et al. (2002) illustrate that experimental results may depend considerably on the choice of the subject pool. The authors find in their gift-exchange experiment that MBA subjects post significantly higher hypothetical effort levels than undergraduate students, and claim that this was due to differences in work experience of the two subject pools. See also Cooper and Kagel (1999) who report similar subject-pool difference in a “ratchet effect” experiment with Chinese students and managers.

around which the effect of incentives in cognitive production can be analyzed. Furthermore, it allows nesting the various literatures on which my work is based within a common framework.

I first provide the basic building blocks of the cognitive production model, noting that the KLP framework and economic literature in general offer very limited information about the properties of the production inputs – cognitive effort and capital (abilities) – and about their substitutability in production. Consequently, I draw mainly from experimental literature to explore the various dimensions of cognitive production. I also refer to the theoretical literature on intrinsic motivation, describe long-run capital formation, and outline general cross-task predictions of the KLP framework. Later sections are empirically oriented, surveying the existing literature on measuring cognitive effort and controlling for cognitive abilities.

4.1 Basic building blocks of the KLP

The capital-labor-production framework introduced by Camerer and Hogarth (1999) is an informal attempt to model production in mental tasks, with particular focus on the effect of financial rewards. As mentioned earlier, the KLP extends the labor theory of cognition (e.g. Smith and Walker, 1993, and Wilcox, 1993) by incorporating a cognitive capital constraint and intrinsic motivation into the individual optimization problem. The labor theory viewed individuals as choosing how much effort to exert to optimally balance monetary benefits accruing from performance against corresponding decision effort costs. The KLP has added that (a) cognitive effort will depend on intrinsic as well as financial goals, and (b) that performance will be affected not only by cognitive effort but also by the availability of cognitive capital. That is, individuals maximize an objective function, consisting of both financial and intrinsic goals, by employing optimal level of effort in cognitive production. The productivity of this effort in turn depends on the availability of cognitive capital relevant to the assigned task, and on the substitutability between the two production inputs.

Camerer and Hogarth attempt to give the KLP an empirical dimension by identifying basic regularities in more than seventy experimental and field studies reporting the impact of incentives on performance.[15] In agreement with other empirical surveys (e.g. Bonner et al., 2000), they report that a majority of experimental studies show little or no effect of incentives on average performance.[16] The authors suggest various explanations for this result, each of them somehow being task-related. The main goal of this section is to organize these explanations in a common framework that will later serve as a basis for formalizing the KLP. Alongside introducing the building blocks of the KLP, I focus on what we do and do not know about particular aspects of the cognitive production process. For illustrations, I will refer only to the most related empirical studies, while general surveys will be reviewed later on.

Similar to physical production, performance in the KLP is generated by two inputs: effort and capital. *Cognitive effort* can have various forms, ranging from simply paying attention to creative thinking. Exerting more effort always brings about higher effort costs but, as explained later, may or may not lead to higher performance. Whether cognitive effort has an upper bound has not been discussed. *Cognitive capital* is a set of

[15] Existing empirical studies have employed a variety of incentive schemes, the details of which will be discussed later on. By "incentives" I mean performance-based financial rewards, as opposed to "flat-rate" schemes meaning hypothetical choices with an initial show-up fee.

[16] Camerer and Hogarth (1999) as well as other empirical surveys mostly refer to results on how varying incentives affects *average* performance in the subject pool, simply because evidence on *individual* incentive effects is sparse or not reported.

mental tools used in production, ranging from memorizing techniques to advanced task-solving skills. Different tasks employ different skills, so capital is task-specific. Employing capital is free (or at least no capital input cost has been mentioned), but its use is limited by an individual-specific capital constraint which is assumed fixed in the short run of an experiment. To what extent the capital constraint varies across individuals has not been addressed.

Resulting from the setup, the short-run production function has only one variable input, effort, whereas capital, being free, is used up to the individual-specific maximum. In the longer run, individuals can acquire additional capital and, if desirable, use more capital-intensive production techniques. Though Camerer and Hogarth do not explicitly discuss the degree of factor substitutability and other properties of the mental production function, the authors seem to assign capital a considerable weight in determining performance.[17] Leaving capital formation effects for later, I will now focus on incentive effects in the short run of an experiment which empirical studies tend to measure.

Conveniently, the effect of incentives on performance can be broken into the *incentive-effort channel* and *effort-performance channel*.[18] As for the former, optimal choice of effort depends on how motivated subjects are to perform well – both intrinsically and financially – and how costly they view effort. These issues will be analyzed once intrinsic motivation is introduced. Turning now to the effort-performance channel, one would expect the choice of effort to map directly to performance. Camerer

---

[17] Though within-subject capital-labor substitution is not possible in the short run when capital is fixed, one still needs to have some clue about this attribute of the production function in order to make between-subject comparisons of production possibilities. Thus capital-labor substitutability should be an integral part of the KLP. This issue will be discussed below in Section 4, namely in the sub-section on capital formation.

[18] The terminology is borrowed from Bonner et al. (2002).

and Hogarth offer two reasons why this may not be so, i.e. why effort may turn out unproductive.

First, a ***performance constraint*** (also termed a “floor effect”) may occur in simple tasks where, in order to reach high levels of performance, subjects require relatively low level of cognitive skills and effort stimulation. If, in addition, experimental design imposes an upper bound on performance (e.g. 100% accuracy of answers), increasing incentives may not stimulate performance if a close-to-maximum performance is already reached by most subjects at a low-incentive level.

For instance, Kahneman et al. (1968) conducted a mental arithmetic experiment in which they required subjects to repeat four-digit strings while adding 0 or 1 to each digit. The authors observed that raising incentives was associated with higher average effort (i.e. greater pupil dilation) but not with higher average performance, presumably due to a performance-constraint effect: average performance was already close to the maximum in the low-incentive treatment (88% accuracy).[19] Another example is a study of Libby and Lipe (1992) who found that, though subjects worker harder (about 3 minutes longer) under higher incentives, only their recall performance improved, among other reasons possibly because their recognition performance reached a performance constraint.

Second, a ***capital constraint*** (also termed a “ceiling effect”) may occur in more complex tasks where successful performance requires a relatively high level of relevant cognitive capital which only very few individuals possess. As mentioned above, such cognitive skills usually cannot be acquired during the experiment itself. Consequently,

[19] One can compare this result with Kahneman and Peavler (1969) who observed in a similar memory-test experiment (remembering noun-digit pairs), that higher incentives were associated not only with greater pupil dilation but also with improved average performance, presumably

increasing incentives may be unproductive in terms of performance if, for most individuals in the subject pool, their relatively low skill level limits performance improvements.

As an illustration of this effect, Awasthi and Pratt (1990) studied the impact of financial incentives and a cognitive skill called “perceptual differentiation” on performance in three problem-solving accounting tasks. The authors first divided subjects into two groups by their perceptual differentiation (low-PD vs. high-PD), and subsequently randomized within these groups by incentive scheme (flat-rate vs. performance-based). The results suggest that whereas average effort (i.e. decision time) was higher for both performance-based incentive groups than for their flat-rate counterparts, average performance responded positively to incentives only for the high-PD group.

4.2 Introducing intrinsic motivation in the KLP

Introducing intrinsic motivation brings about a new dimension into the KLP framework. Without it, the individual optimization problem would be incomplete as one needs to explain the empirical observation that subjects perform relatively well (i.e. their behavior is not purely random) even without performance-based financial rewards (Smith and Walker, 1993). According to Camerer and Hogarth, experimental subjects are intrinsically motivated because they find the task interesting and intellectually challenging or want to appear smart, because they like eliciting behavior that is socially desirable or expected of them, or because they want to amuse others. The authors argue

---

because of no performance-constraint: average performance was only 55% of the maximum for the high-incentive treatment compared to 18% for the low-incentive treatment.

that intrinsic motivation varies across tasks, but implicitly treat it as approximately constant across subjects in a given task.

Intrinsic motivation enters the KLP in two different ways. First, intrinsic motivation may interact with incentives and thus affect the incentive-performance channel. Second, intrinsic motivation may even moderate functioning of the effort-performance channel. Clearly, however, intrinsic motivation can affect neither of the channels in boring or routine tasks where its level is presumably rather low. In such cases, therefore, financial incentives will be the main stimulant of effort, though effort can still turn out unproductive if either the performance or the capital constraint affects the effort-performance channel. Without further information about the nature of the task, the performance constraint is more likely since boring and routine tasks tend to require relatively little capital.

In the tasks with low intrinsic motivation and capital requirements, incentives predominantly play a positive role (Jenkins et al.[20], 1998). Camerer and Hogarth argue that this may be because under low incentives subjects play around or experiment with responses to make the task more interesting which subsequently hurts performance, while under high incentives they are prepared to accept some boredom to perform better and earn more. This story seems to hold in binary probability matching experiments where subjects indeed stop guessing and start moving closer to profit-maximizing predictions following an incentive raise (e.g. Siegel et al., 1964). In such cases, incentives probably stimulate relatively basic forms of effort. Camerer and Hogarth cite as examples memory

[20] Somewhat confusingly, the authors also report that the level of intrinsic motivation did not seem to influence the size of incentive effects. Since intrinsic motivation in the (mundane) task they study is expected to be low anyhow, this finding is unlikely to generalize to other (more

and recall tasks (where higher effort in the form of paying more attention helps improve memorizing abilities) or multi-cue probability learning tasks (where effort in the form of keeping track of past trials improves predictions). Further, Smith (1962) and others report that high-powered incentives seem necessary to achieve fast convergence to theoretically-predicted behavior in repeated-type experiments to overcome fatigue and fading intrinsic motivation.[21]

I now turn to tasks which are sufficiently interesting so that intrinsic motivation is initially high. I say initially because introducing financial incentives may reduce it and thus hurt the incentive-effort channel. Specifically, especially psychologists have argued that incentives may not lead to higher overall effort if they ***crowd out*** intrinsic motivation (see Frey and Jegen, 2001, for a survey). That is, incentives may partly or fully replace intrinsic motivation in its role of effort stimulant, rather than simply strengthening its positive effects. Among other possibilities, psychologists model incentives as increasing the subjective decision effort costs (or alternatively decreasing the intrinsic benefits from performing well) at the margin. If so, individuals may not respond to incentives by increasing effort.[22]

Camerer and Hogarth do not explicitly discuss crowding out effects, though they do recognize that the incentive-effort channel might not function. For example, the authors mention that raising incentives may have negative side-effects such as stress and anxiety (e.g. Camerer, 1998). A study that seems to justify considering crowding out effects as

---

interesting) tasks. This issue, of course, lies at the heart of the intrinsic motivation debate mentioned in Section 2; see also footnote 8 in Hertwig and Ortmann (2001).

[21] To complete the picture, however, other studies report that raising incentives actually *induces* subjects' experimentation with responses, both in easy and complex tasks, as subjects attempt to improve performance by raising effort in the form of creative thinking which turns out unproductive (e.g. Arkes et al., 1986, and Hogarth et al., 1991).

part of the KLP was conducted by Gneezy and Rustichini (2000; see Section 3 for details). In their IQ-test experiment, average performance was higher in a flat-rate than in a low-incentive treatment, yet clearly highest in a high-incentive one. This non-monotonic relationship between incentives and performance could hardly be explained by considering financial incentives alone.[23]

Even if intrinsic motivation is not crowded out and the incentive-performance channel functions, high level of intrinsic motivation may still hinder the effort-performance channel because it makes performance or capital constraints more imminent. Namely, if subjects' intrinsic motivation is itself sufficient to stimulate effort high enough to come close to either of the constraints, then extra effort stimulated by incentives will be mostly unproductive. Hence back to the above description of performance and capital constraints.

To support this claim one needs to find evidence for performance and capital constraints in tasks where intrinsic motivation is high and which have a flat-rate control treatment. Bonner and Sprinkle (2002) report that incentives generally have no effect in problem-solving tasks where intrinsic motivation is presumably high yet capital constraint bites. Camerer and Hogarth observe that incentive effects are generally stronger in studies comparing flat-rate and performance-based rather than low-incentive and high-incentive treatments. This would seem to suggest that, first, introducing incentives complements rather than crowds out intrinsic motivation, and second, that

---

[22] See Section 2 for details and for references concerning the crowding-out debate.

[23] In Gneezy and Rustichini, individual performance varied in all three treatments, presumably due to varying cognitive abilities in the subject pool. The greatest dispersion of performance occurred in the flat-rate treatment, confirming the claim of Smith and Walker (1993) that incentives are usually associated with variance reduction. Camerer and Hogarth argue that this is mainly due to "cleaning" the data from outliers.

increasing incentives may indeed activate either of the constraints. However, to draw stronger conclusions from this observation, one would need to classify the included studies by their levels of intrinsic motivation.

4.3 General 'cross-task' predictions of the KLP

Focusing again on the short run of an experiment, several predictions can be made from the basic KLP setup outlined above. Such predictions will be closely interconnected with the nature of the task and of the subject pool. The two will determine the level of intrinsic motivation and presence of performance or capital constraints, and hence in turn whether incentive-effort and effort-performance channels function or not. I will first suggest some general patterns concerning the incentive-effort-performance relationship, and then predict how the size of incentive effects is likely to vary across tasks.

The first general observation concerns the (occurrence of) the performance and capital constraints. On average, both of them will be a function of task complexity. That is, ceteris paribus, the harder the task the more likely the capital constraint and the less likely the performance constraint.[24] Furthermore, the likelihood and power of the performance and capital constraints is likely to depend on the distribution of cognitive capital in the subject pool. Ceteris paribus, the wider is the variability of cognitive capital across individuals, the less imminent will be the constraints, because the variability ensures that some individuals in the capital distribution remain unaffected by the constraints (unless these are very tight in the sense that they reach most subjects).

The second general point concerns the behavior of effort when either of the constraints becomes binding. In such situations, Camerer and Hogarth argue, the marginal return to effort is low. Accordingly, individuals should not find it worthwhile to increase effort. In contrast, most studies that measured effort report that increasing incentives generally leads to higher effort regardless of whether performance improvement follows. This observation must be taken into account when formalizing the KLP model. Namely, it brings up the question of whether subjects actually observe the constraints they might face. The presence of a performance constraint should be easier detected in within-subject design where subjects can judge from their past experience (i.e. from the previous incentive treatment) whether increasing performance is possible and more effort worthwhile. Clearly, such a direct measure of performance is not available to subjects in a between-subject design. Whether the capital constraint is at all a priori observable by subjects is questionable. The possibility that subjects may have limited information about their own cognitive abilities needs to be addressed (e.g. Benabou and Tirole, 2003).

From what has been said so far, can one make more specific predictions from what is known about the nature of the task? The answer is yes, provided one makes a rough assumption that easier tasks carry with them less intrinsic motivation and lower capital requirements, and further that a tasks' intrinsic motivation does not vary across subjects. Then it follows that in simple tasks where both intrinsic motivation and capital requirements are low (i.e. capital constraint is unlikely to occur), incentives are expected to stimulate both effort and performance unless performance constraint is present. The performance constraint is relatively well observable by the experimenter: one simply compares actual to maximum performance at the baseline incentive treatment. In case

[24] It should also be noted that the performance constraint is a function of experimental design and can to some extent be avoided by appropriate choice of performance scale.

that performance stagnates, effort may still rise if subjects cannot observe the performance constraint. With greater task-complexity, subjects' intrinsic motivation rises, the probability of capital constraint increases, and the probability of performance constraint falls. If the task is such that intrinsic motivation gets crowded out by incentives than both effort and performance should increase little. If effort rises but performance stagnates, this cannot be due to crowding out, but most likely due to the capital constraint being present but unobserved, though one should again pre-check the presence of the performance constraint.

Overall, this creates some structure for categorizing tasks according to their complexity and intrinsic motivation expected to be involved in their completion. The assumptions that complexity and intrinsic motivation are positively related may not be appropriate, but can be tested. The performance constraint will interact whenever present, but at least is directly observable. Given that task can be ordered by their intrinsic motivation, one can look at the size of incentive effects and, controlling for the observed presence of performance constraints, test the KLP framework at an across-study level.

Such a broad across-study test has been attempted by Bonner et al. (2000). The authors examine the effect of task complexity and incentive-scheme type on performance, hypothesizing that the functioning of the effort-performance channel depends on the former, while the functioning of the incentive-effort channel on the latter.[25] In a logit-

[25] More specifically, the authors assume that average skills in a subject pool remain roughly constant across studies (mainly college students), so task complexity determines the complexity-skill gap. Tasks are broadly divided into 5 categories according to their average information processing requirements. Performance is expected to be less responsive to effort increases in more complex tasks for which the gap is bigger. Incentive scheme type describes how rewards are linked to performance; schemes differ in financial and non-financial aspects (e.g. quota vs. tournament schemes). Size of incentive effects is discretely categorized as positive, none and negative.

type estimation using 131 studies, they find that incentive effects are stronger where the complexity-skill gap is smaller, and strongest in studies using quota incentive schemes and flat-rate as a control. The authors note, however, that their broad categorization of tasks hides other correlates such as the level of intrinsic motivation. Further, interpreting the results is difficult due to an apparent across-study correlation between the choice of the control treatment and the employed incentive scheme, and also due to the different perception of incentives by subjects in the various schemes.

4.4 Capital formation in the longer run

Subjects come to the experiment with some "procedural knowledge", i.e. cognitive capital, which helps them in one way or another in solving the assigned task.[26] I have argued above that acquiring extra capital during the experiment itself is limited. This at least seems to be the stance of cognitive psychology literature suggesting that capital (in the form of expertise) gets acquired slowly through learning by doing rather than learning by thinking, and even without being aware of it (e.g. Ericcson and Smith, 1991, and Reber, 1989). As discussed below, however, this does not preclude the acquisition of capital in repeated-type experiments through building up experience.

Camerer and Hogarth argue that individuals acquire new capital only if given a chance and if they find it profitable. That is, individuals "invest" into additional capital – by foregoing current performance and devoting effort to learning – if the resulting net expected gain in terms of improved future performance is positive. This mechanism seems quite consistent with the learning-by-doing hypothesis, and again brings in

[26] According to Camerer and Hogarth, such procedural knowledge should be contrasted with "declarative knowledge", i.e. descriptive information about the task obtainable from the

incentive effects since these will determine the performance-investment tradeoff. The relationship between learning and incentives was independently examined in the context of a repeated-type experiment by Merlo and Schotter (1999) who report that subjects learn differently in a learn-before-you-earn and learn-why-you-earn incentive treatment. Judged in terms of subsequent performance, they seem to "learn better" in the former treatment, whereas in the latter they seem unprepared to sacrifice short-run performance for longer-run gains.[27]

Exactly how the newly-acquired capital is incorporated into production can be observed only indirectly. Camerer and Hogarth conjecture that capital-labor substitution takes much the same form as in physical production: individuals switch to more capital-intensive production techniques. For example, in the "stagecoach" task consisting of connecting an initial node to a destination, individuals who acquire relevant capital (which here amounts to knowing that the problem is best solved backwards) substitute such capital for mental effort and thus avoid the effort-intensive solution of starting at the initial node and evaluating every possible connection. As a game-theoretic example, searching for asymmetric mixed-strategy equilibria requires experience and skillful judgment rather than calculating all possible probability mixes.

Repeated-type experiments, such that of Merlo and Schotter, allow for capital formation by building up experience during a single experimental session. The interaction of experience and incentive effects will be further discussed below in Section 4. There are other forms of capital formation such as between-session learning and between-

instructions (e.g. information about the goals of the task, about its players and payoffs, etc.), and other descriptive pieces of information obtained in the past.

subject communication that might occur outside the lab. However, these effects are usually neither allowed nor studied in experimental settings because experimenters fear that they might lose control over the design. Yet the celebrated cleanness of lab results may come at the expense of their greater applicability (see also Section 2).

Consistent with this claim, Camerer and Hogarth argue that between-session learning and especially communication (learning from others or even "teaching" others) are important forms of capital acquisition in real world settings. Accordingly, these should be considered in laboratory settings to make lab experimental results more robust and broadly applicable. Doing so could help reconcile the debate about whether lab incentive effects are underestimated (e.g. Jenkins et al., 1998). To this end, Bonner et al. (1999) conducted a 12-week mental-multiplication experiment where they controlled for both incentives and individual skills. Interestingly, though in the short run subjects responded to higher incentives only by increasing effort, they were able to gradually improve performance through learning. Introducing communication among team members presents another interesting type of experiment, though it is also a perfect example of the resulting confounds: how could one distinguish between the effects of team capital formation and team division of labor?

4.5 Measuring cognitive effort

As noted above, measuring effort is important for identifying which part of the incentive-effort-performance relationship might not function. This is because inspecting only the relationship between incentives and performance without paying attention to

[27] This finding cannot be easily generalized to other learning situations, however, since Merlo and Schotter permitted experimentation with responses (in the learn-before-you-earn treatment) rather than pure experience building.

effort blurs empirical identification of what exactly causes incentives to work or not to work. In particular, is it motivation crowding out, or either of the performance and capital constraints, that drive the effectiveness of incentives? In addition, measuring effort could also offer some hint as to whether individuals face an upper bound on effort.

Unfortunately, measuring effort is a complicated and rather controversial issue. Smith and Walker (1993), for example, claim that cognitive effort (in the form of concentration, attention, thinking, monitoring, reporting, acting, etc.) is unobservable, or observable only indirectly through its effect on performance. In contrast, Camerer and Hogarth claim that data on effort are essential for explaining the decision processes that take place during cognitive production, and they point out to several studies that have attempted to measure decision effort directly.

The most frequent measure of cognitive effort is ***response time***. As observed in the studies mentioned above, higher incentives are usually associated with longer response times. This result was confirmed econometrically by Wilcox (1993) and recently by Moffatt (2003) who both examined the relationship between response times, incentives and other covariates (such as task complexity and experience) in the context of lottery choices. Unobtrusively measured response times may not be ideal for correctly detecting effort (though it avoids worrying about effort upper bound). In particular, Bonner et al. (2002) and Libby and Lipe (1992) argue that effort intensity as well as effort duration matters.

Libby and Lipe (1992) provided a preliminary attempt to control for effort intensity by measuring the number of times subjects cycled (scrolled) through a computer screen listing items to remember. More advanced measurement of ***looking-up patterns*** was conducted by Johnson et al. (2001) in a sequential bargaining experiment where the authors recorded the number and the pattern of computer-mouse movements subjects performed in searching for task solution. Interestingly, the recorded mouse movements suggest that subjects looked forward rather than backward (despite backward induction being a more effective solution concept). Further, the number of mouse movements turned out to be a U-shaped function of response times.

Researchers have also come up with physical measures of effort intensity. As reported in the above mentioned studies of Kahneman and his colleagues, ***pupil dilation*** increases with incentives. Dickhaut et al. (1997) measured how ***heart rate*** and ***galvanic skin response*** varied during their English auction experiment. Consequently, it seems a natural question to ask to what extent these physical measures correspond to the neural processes in the human brain. That is, to what extent they manifest the rate of firing of neurons and which parts of the brain are most affected. These and other questions have been considered in the uprising neuro-economics literature (e.g. Glimcher, 2003). There even exist attempts to model neural activity in economic style (e.g. Gold and Shadlen, 2001).

Camerer and Hogarth offer other potential measures of effort that can be obtained from verbal protocols filled out by subjects after the experiment. For example, the authors suggest that a ***measure of recall*** of certain aspects of the experiment could serve as a “proxy for the amount of decision effort expended in the first place” (p.29). As an even more indirect measure of effort, seemingly contradicting the variance-reduction function of incentives advocated by Smith and Walker (1993), the authors suggest that an increase in the ***variance of responses*** can often be interpreted as higher overall effort. As an example, they cite Samuelson and Bazerman (1985) who found that when bidding for real rather than hypothetical money, average performance in their difficult acquire-the-company problem did not improve but dispersion increased dramatically. Libby and Lipe (1992) document that, in their memory experiment, both effort (response time) and performance (recall or recognition) have a larger dispersion under their piece-rate incentive treatments.

Since measuring effort in the laboratory has received even less attention than measuring incentive effects, it seems too early to judge the validity of the various techniques. One would probably be better off using several of them simultaneously. I will illustrate in the next section that a similar story holds for measuring cognitive capital and disentangling its effects from those of other variables.

4.6 Interacting cognitive capital with incentives

As noted above, the KLP describes cognitive production and incentive effects therein as an individual optimization problem, and consequently lends itself to deriving individual-level predictions for the effects of incentives and capital on mental performance. Some preliminary across-task predictions were informally sketched out in Section 4.3 above. Below I discuss existing empirical approaches to controlling for cognitive capital and other environmental variables in identifying the 'true' incentive effects. Throughout, it is implicitly assumed that the (artificial) performance constraint can be avoided by appropriately adjusting the experimental design. My focus will be on why existing approaches have partly failed to estimate the true size of incentive effects, and how they could be improved upon and combined to get a clearer picture thereof.

Measuring cognitive capital directly, for each individual in the subject pool, is the first possible approach to controlling for it. In a between-subject design, one can subsequently randomize subjects into respective incentive treatments, obtain a measure of performance (and possibly effort) from the experimental task(s), and estimate the separate effects of incentives and cognitive capital on performance.

This procedure has several potential drawbacks. First, it requires that the task-specific cognitive capital is a priori identified and measured by a suitable proxy. Otherwise the resulting explanatory power will be low and estimates at best inefficient. Without a priori information about their validity, one should probably combine quite general proxies, such the Wechsler short-term memory (STM) test used by Devetag and Warglien (2003), with more specific proxies, such as the EFT test of perceptual differentiation used by Awasthi and Pratt (1990). There are even more general proxies relating to educational background such as the GPA score used by Eckel (1999).[28]

The second problem also relates to the measurement of cognitive capital proxies: should they be measured under incentive or hypothetical conditions? Several studies report that the results of various STM tests seem to depend on design of the test (e.g. Koriat and Goldsmith, 1996a and 1996b; see also Section 2 above). But this would invalidate the assumption underlying the KLP that the capital constraint is fixed and thus also invariant to incentives, leading to circularity of reasoning about whether STM

[28] Of the three studies, only Awasthi and Pratt (1990) study incentive effects explicitly, though they cannot do full justice to the estimation due to their binary performance measure. As mentioned earlier, Devetag and Warglien (2003) find a significant positive correlation between STM and a performance in three game-theoretic tasks. Eckel (1999) uses natural framing (of lottery choices), which she argues is a substitute for increased incentives, rather than explicit incentives. As to the choice of appropriate proxies, the general ones are likely to encompass others while more specific ones could significantly add to explanatory power. Apart from the

measures capital or rather some form of performance. Cooper and Kagel (1999) note that the responsiveness of STM to incentives is probably due to improved chunking (i.e. remembering items in bigger groups) rather than enlarged capacity of working memory itself. This would suggest that more rudimentary working-memory tests should be used as cognitive capital proxies.

Third, an additional complication arises in repeated-type experiments (or in within-subject design) where the effect of incentives and cognitive capital on performance will be confounded with experience effects across rounds (or treatments).[29] To the extent that acquiring experience in a particular task is negligible – as argued by cognitive psychologists – one can ignore such confounding effects and consequently even control for cognitive capital differences by studying within-subject incentive effects. The measured capital proxies could then be use to assess the size of incentive effects across the cognitive capital distribution. However, capital formation through building up experience during experiments (as well as real-world tasks) seems substantial.

Hence ***measuring experience***, or controlling for it, may prove a more efficient and realistic way of interacting incentives with capital. Assuming that experience is the major piece of capital which overwhelms the effects of any capital built up prior to the experiment, one could simply ignore all other forms of capital, including those discussed above, and focus on controlling for experience alone. Namely, providing sub-pools of subjects with different levels of experience and then randomizing them into the chosen incentive treatments permits estimating average incentive effects while controlling for experience.

---

huge literature on memory capacity, the literature on the identification and measurement of other proxies is sparse.

This approach was to some extent adopted by Jamal and Sunder (1991) in the context of a repeated auction experiment. The authors report that incentive effects were significant in the (smaller-sized) market with inexperienced traders but not with experienced ones, and that experience actually had a statistically stronger effect on performance than incentives.[30] Sprinkle (2000) analyzed experience and incentive effects in the context of repeated output choices based on a random realization of a state of nature. Learning by requesting feedback was allowed, but costly in terms of foregone time. Comparing flat-rate and incentive treatments, incentives induced higher average effort (measured as time) over the whole course of the experiment, but corresponding average performance differences only occurred after several repetitions and further increased with them.[31]

The above approach of controlling for experience or learning has a potential drawback in that we do not know much about how the learning process evolves and consequently how to interpret the results. Going back to Merlo and Schotter (1999) cited above, subjects under flat-rate and incentive schemes seem to learn differently (and also different things). A related question concerns how long a period of learning to allow and based on which results to measure performance. These complications further suggest that controlling for experience effects by simply inserting dummies for rounds (e.g. Cooper and Kagel, 1999, and Moffatt, 2003) should be treated with caution.

---

[29] For an important illustration, see Harrison et al.'s (2003) critique of Holt and Laury (2002).

[30] This result seems rather counter-intuitive: one would expect incentive effects to be stronger for experienced subjects. However, inspecting the results closet reveals that experienced subjects sharply converged to the theoretical optimum even without incentives and so reached the performance constraint once incentives where introduced.

[31] This seems to support the argument that results obtained from one-shot experiments with no learning possibilities and without appropriate incentive stimulation may be misleading. However,

To avoid these complications, one can look into other ways of controlling for experience. One possibility is to "activate" cognitive capital (experience) by ***framing the task*** in a way that makes it familiar to a sub-pool of experimental subjects.[32] Then one can again randomize subjects into incentive treatments and measure the incentive-effect differential across the sub-pools.

A modification of this procedure was used by Cooper and Kagel (1999) in the context of a repeated signaling ratchet-effect experiment with Chinese students and managers. The authors report that a five-fold increase in incentives was associated with a better initial performance (more strategic play), but that this differential disappeared with repetition. Moreover, managers (experienced with ratchet effect from practice) performed considerably better under high incentives than students, yet only in the business-context treatment.[33] In a different approach to examining framing effects, Hannan et al. (2002) performed a gift-exchange experiment where they required subjects to post hypothetical effort level in response to firm wage offers. They found that MBA subjects posted significantly higher effort levels than undergraduate students, and claim that this was due to differences in work experience of the two subject pools.

Camerer and Hogarth argue that natural labeling has been largely unexplored in economics, and that the resistance to using natural labels comes from fears of creating

---

it is probably too early to make a stronger conclusion: Cooper and Kagel (1999) report a contrasting result that incentive effects matter only initially and disappear with repetition.

[32] Giving problems familiar contextual labels allows subjects to "activate" task-specific cognitive capital, such as domain-specific heuristics or choice rules, which would otherwise be left unused because people often seem unable to transfer capital across tasks (see evidence on capital transfer below). The importance of social context for activating reasoning skills was discussed for example by Ortmann and Gigerenzer (1997).

[33] This was the case with the exception of older managers who, regardless of framing, performed even worse than students. The authors argue that this was due to their low formal educational level.

non-monetary stimulation in the experiment, thus loosening control over incentives (see also Ortmann and Gigerenzer, 1997). Although being closer to reality, this is indeed a potential drawback of using framing to disentangle the incentives and cognitive capital effects. In the words of the KLP, does framing operate only on capital, or does it also alter the level of intrinsic motivation that individuals possess? A further complication of framing has to do with its implementation, i.e. to what extent it is it possible to reconstruct the natural environment of subjects in the laboratory. This motivates the search more a less confounding yet still flexible way of controlling for capital effects.

Supplying relevant capital through instructions potentially presents the most flexible way of controlling for cognitive capital differences in the subject pool. In contrast to above methods, cognitive capital is "imposed" on subjects rather than acquired or activated by them. To estimate average effects of incentives, one simply needs to randomize subjects into instructions and incentive treatments, respectively. Similar to controlling for experience, this presumes that cognitive capital acquired prior to the experiment is unimportant relative to the task-specific capital supplied through the instructions.

This approach was used in a physical setting by Baker and Kirsch (1991) who in their pain-endurance experiment studied the effects of incentives as well as of supplying coping-skill instructions which explained how to deal with pain arising from emerging hands into cold water. The authors reported that compared to the control condition performance (measured by time hands remained under the water) increased in both the incentives and instructions treatments. Unfortunately, the authors did not combine the two treatments. Johnson et al. (2001) found in their ultimatum bargaining experiment that though subjects could not learn by repetition to make close-to-equilibrium offers, brief instructions about the backward induction solution concept caused them quickly converge to the theoretical optimum. Further, these "taught" subjects were partly able to teach other non-taught subjects when mixed with them. The authors again did not combine the instructions treatment with varying incentives.

Supplying capital as part of instructions may face the same problems as framing: if better instructions make the task more interesting, intrinsic motivation effects will be confounded with those of incentives and capital. As a methodological note, Camerer and Hogarth call for overall simplifying instructions to reduce decision costs and help subjects concentrate on production rather than understanding the task. The authors have in mind for example clearly conveying the production requirements, minimizing attention and memory requirements, and helping subjects produce the mapping from actions to payoffs by supplying payoff tables. With respect to incentive effects, clearer and unified instructions are likely to clean the data of unnecessary noise caused by subjects unable to comprehend the task requirements.[34]

Lastly, ***approximating capital based on past performance*** is the most general yet admittedly also the noisiest way of controlling for cognitive abilities. It encompasses all the above measures, taking over their merits and pitfalls. If one gives all subjects equal opportunity to learn (be it under hypothetical or real stakes), the resulting performance can be taken as a very rough all-encompassing ex post measure of cognitive capital. Provided that subjects are subsequently randomized into incentive treatments and

[34] A similar point was also made in a methodological debate by Binmore (1999).

perform the same or similar task (anticipated or unanticipated one), one can separately estimate the effects of capital and incentives.

The past-performance measure of capital was implicitly used by Vandegrift and Brown (2003) in the context of a multi-cue probability learning tournament. The authors report that in their complex task (where the performance constraint was not present), predominantly only winners from the previous rounds were able to react to raising incentives by improving their performance. The experiment of Merlo and Schotter (1999) also resembles this type of setting, in addition allowing different types of learning (under real and hypothetical rewards). However, the authors do not establish any link between individual performance after several rounds of learning (i.e. "capital" formation) and the corresponding performance in the final (surprise) round, and as mentioned earlier their argument is not so much about capital formation but more about experimentation.

As already mentioned, the measure of capital based on past performance is rather noisy, encompassing differences in initial cognitive abilities prior to the experiment, and learning abilities and exerted effort level during the experiment itself (thus involving the circularity of reasoning problem mentioned in the context of the STM test). Despite these disadvantages, there still may be a good reason for using such a measure, considering how little is known about what constitutes individual cognitive capital and what governs its acquisition. The correct choice of capital measure seems to depend predominantly on the type of task and on the purpose of measurement.

First, I claimed above that experience effects may overwhelm the impact of other forms of capital, and that allowing building up of expertise could thus be the appropriate way of distilling capital and incentive effects. Yet one may not be sure how fast expertise builds up, and several field experiments even suggest that real-world expertise may be virtually useless. For example, Ball and Cech (1996) in their study of subject pool choice and treatment effects in experiments report that well-educated subjects perform not much differently from less-educated ones using only simple formulas. Dawes et al. (1989) arrive at similar conclusions in their review of nearly hundred field studies of judgment tasks.

Second, if one really needs an individual measure of cognitive capital for a specific task – which would be hard to obtain by using the experience, instructions or framing measures discussed above – past performance may be the only choice. Even if the experimenter is able to identify a set of proxies and succeeds in measuring them, this still does not ensure that subjects actually comprehended them and were able to use them.[35] Awasthi and Pratt (1990) reported that a significant proportion of their subjects were not able to apply abstract cognitive tools to concrete settings despite comprehending them. In general, limited transferability of cognitive tools to even slightly different settings seems well established in cognitive psychology literature (e.g. Anderson, 2000) as well as other experimental contexts.[36] Consequently, identifying the very specific pieces of cognitive capital useful in particular situations may be an unrewarding task.

5. The road ahead: formalizing and testing the KLP framework

Section 4 above provides the insights for formalizing and testing the KLP framework. For exposition of the forthcoming theoretical work, I present elsewhere a preliminary attempt to

[35] Indeed, with respect to the discussion immediately above, one should ensure comprehension by supplying clear instructions.

[36] For example, Kagel and Levin (1986) found that even simple heuristics (avoiding winner's curse) are not transferable to a very similar task (with merely changed number of bidders).

formalize the KLP.[37] To this end, the labor-theory formulation of the individual optimization problem does not lend itself to a clear extension. Consequently, since there is little a priori guidance regarding alternative modeling techniques, I start with a deterministic neoclassical framework that resembles firm production decisions. Focusing on the short-run case of an experiment, I model the cognitive capital constraint as fixed and variable across individuals; the trivial performance constraint and its potential interaction with the level of intrinsic motivation in the task (see Section 4) is for now ignored. I introduce intrinsic motivation into the model in a simplistic way along the lines of Frey (1997), but into the subjective cost function rather than the production function. This does not add any major insights regarding incentive effects, and further work will be needed along the lines of Benabou and Tirole (2003) to make the interaction between motivational factors more realistic.

On the other hand, introducing an individual-specific cognitive ability constraint can clearly drive the results, yet the analysis only seems tractable once we assume a particular form of the cognitive production function. Accordingly, I have been experimenting with the unknown degree of factor substitutability, starting with the relatively general CES formulation but subsequently turning to more specific cases such as Cobb-Douglas and Leontief production functions. I have further explored the implications of corner solutions in effort, as well as the variability of effort upper bound across individuals. For the Leontieff case, a "non-optimizing" version of the model are considered in an imperfect attempt to explain why – in the case of a capital constraint – effort might increase despite subsequent performance stagnation. It turns out, however, that this result may be obtained by making the individual cognitive ability constraint partly unobserved. Although my theoretical framework

[37] For details of the proposed theoretical and empirical work, see Appendix B and C, respectively, to my Dissertation Proposal at http://home.cerge-ei.cz/rydval.

is clearly in need of further refinement and should in the future have a stochastic element, it provides a basic idea of how to discipline the empirical analysis.

One of the ultimate goals of the proposed research is to test econometrically the theoretical predictions derived for the effect of incentives on effort and performance. However, it seems more natural to start with estimating the cognitive production function itself. To this end, one would need reliable individual-level data on cognitive effort, capital and performance, of which only the last item is usually available. As outlined in Section 4, measuring cognitive effort is mostly done by means of measuring decision time (e.g. Wilcox, 1993). This may, however, be an imprecise measure of effort duration and further does not take into account effort intensity (e.g. Libby and Lipe, 1992). I discussed various other measures of effort (e.g. pupil dilation used by Kahneman et al., 1968) which nevertheless do not seem suited for estimation purposes. As illustrated in Section 4, one can also have limited confidence in the currently available measures of cognitive capital. The reliability of these measures and their applicability to particular task settings is questionable. As an example, I discussed the potential estimation and methodological problems associated with the measurement of memory capacity. Although the measurement issues will clearly require further research, I am currently investigating additional problems associated with estimating the cognitive production function (e.g. Levinsohn and Petrin, 2003).

The econometric option more directly related to the theoretical KLP modeling consists of estimating the relationship between incentives and cognitive performance while controlling for cognitive capital. Although doing so circumvents the problems associated with measuring effort, those associated with measuring cognitive capital prevail. In fact, since capital seems incredibly task-specific in its nature,[38] the best available technique of controlling for it seems

[38] For example, Kagel and Levin (1986) report that even simple heuristics (i.e. avoiding winner's curse) are not transferable to a very similar task with merely a changed number of bidders.

“supplying” it in different quantities to subjects as part of the experimental instructions (e.g. Johnson et al., 2001) and subsequently randomizing subjects across incentive treatments. This also partly circumvents the problem of controlling for prior expertise in the task which would otherwise make the estimates inconsistent due to an omitted variable bias. In Section 4, I discussed other potential remedies for this problem which have often been employed at the expense of introducing the confounding effects of experience building (e.g. Sprinkle, 2000). Experience formation, apart from presenting a capital-formation confound, also brings with it additional concerns about the validity of the employed randomization procedures since subjects may be influenced by the accompanying between-round wealth effects (e.g. Ham et al., 2001).

Even if cognitive capital can be adequately controlled for, further estimation issues concern especially the manner in which individuals are matched according to their abilities and other observable characteristics (e.g. Imbens, 2003, and Heckman and Navarro-Lozano, 2003), notwithstanding the potential complications arising from (unobserved) individual-specific variability in intrinsic motivation. However, particular estimation problems will be solved more effectively when considered in conjunction with the design of the experiment(s) necessary to obtain adequate data. In this respect, arguably the best initial strategy is to replicate a well-known existing study, such as the one by Gneezy and Rustichini in Section 3, which reported mixed (aggregate) evidence on incentive effects but did not further examine why this was so. I also have in mind an adaptation of a multi-cue probability learning task (e.g. Vandegrift and Brown, 2003, and Hogarth et al., 1991) which would in fact permit incorporating a stochastic element into the individual optimization problem and thus merge closer the theoretical and econometric representations. I further noted in Section 2 that replicating more complicated game-theoretic tasks runs several risks associated with capital measurement as well as the interpretation of individual performance.

6. Contribution to economic research

The proposed research focuses on examining cognitive production in a more thorough manner which in turn will help enhance our understanding of how mental production processes work. Consequently, economists will be better equipped for drawing insights from other fields that also study cognitive processes and incentive effects, such as neuro-biology and cognitive psychology, and for incorporating these insights into their models of human behavior. Apart from this general contribution, I sketch out below three practical applications of my research.

The first application concerns unifying experimental methodology regarding incentives. As outlined in Section 2, there has been a heated debate among experimentalists about the use of financial incentives in experiments. Apparently, the use of incentives has also acted as a sharp but maybe unjustified discriminator of which studies get accepted for publication.[39] The research on cognitive production should help us inform the experimental methodology debate about when and how to use incentives. Furthermore, it should also contribute to our understanding of whether rationality anomalies that are repeatedly reported in experiments arise due to inadequate incentive stimulation, to inappropriate cognitive abilities in the subject pool, or other design and implementation issues.[40] And perhaps most importantly, the study of cognitive processes and especially their systematic empirical exploration should shed light on the issue of robustness of laboratory results and their applicability to real settings (see Carpenter, Harrison and List, forthcoming).

The second application of my research concerns economic modeling and real-world design of work compensation schemes. As discussed in Section 2, the widely used assumption that incentives necessarily stimulate effort and performance does not seem to be

[39] See Hertwig and Ortmann (2001) and Camerer and Hogarth (1999) for details.

fully supported by evidence. I argued in several places above that the various attempts to incorporate intrinsic motivation issues into agency models have so far been unsatisfactory, and that no study has attempted to explicitly incorporate the crucial cognitive ability constraint. That the existing theory may be inadequate seems reflected in the fact that firms in practice rarely offer contracts predicted by the theory (Prendergast, 1999). Despite this, I argue that firms might actually want to learn from the proposed cognitive production research how to redesign tasks and incentive schemes to suit human capital of their workers.[41]

Third, the knowledge of individual cognitive abilities and their role in the incentive-performance channel is often crucial for designing public economic policies. Consider the debate surrounding the measurement of the "true" return to education. The fact that cognitive ability is at least partly unobservable while at the same time highly correlated with schooling choices makes it inherently difficult for labor economists to separate individual return to schooling from the return to ability itself (e.g. Heckman & Vytlacil, 2001, and Taber, 2001). But knowing whether further education creates valuable skills or serves merely as a screening device for employers is important for deciding where to direct resources within the educational system. The proposed research on incentives could also broaden our knowledge of the optimal design of taxation or anti-corruption schemes, issues that are of particular interest to transition economies. In short, a closer look at the incentive-performance channel is likely to reveal invaluable information.

---

[40] Gigerenzer et al. (forthcoming), for example, discuss the issue of cognitive illusions.

[41] A related literature is that on (rank-order) tournaments where payments are determined in advance and only relative performance matters (e.g. Lazear and Rosen, 1981, and Prendergast, 1999). The designers of tournament-like schemes might arguably want to learn even more from the research on the role of cognitive capital in production. As the above re-analysis of Gneezy and Rustichini data suggests, such incentive schemes are likely to be largely ineffective if ability turns out to be a major determinant of performance.

Bibliography

Acemoglu, D., M. Kremer and A. Mian (2003), "Incentives in Markets, Firms and Governments," *NBER Working Paper 9802*.

Anderson, J. R. (2000), *Cognitive psychology and its implications*, W H Freeman (5th edition).

Awasthi, V. and J. Pratt (1990), "The Effects of Monetary Incentives on Effort and Decision Performance: The Role of Cognitive Characteristics," *Accounting Review*.

Baker, S. L. and I Kirsch (1991), "Cognitive Mediators of Pain Perception and Tolerance," *Journal of Personality and Psychology*.

Ball, S. B. and P.-A. Cech (1996), "Subject Pool Choice and Treatment Effects in Economic Laboratory Research," In R. M. Isaac (ed.), *Research in Experimental Economics*, JAI Press.

Benabou, R. and J. Tirole (2003), "Intrinsic and Extrinsic Motivation," *Review of Economic Studies.*

Binmore, K. (1999), "Why Experiment in Economics?" *The Economic Journal.*

Bonner, S. E. and G. B. Sprinkle (2002), "The effects of Monetary Incentives on Effort and Task Performance: Theories, Evidence, and a Framework for Research," *Accounting, Organizations & Society*.

Bonner, S. E., R. Hastie, G. B. Sprinkle and S. M. Young (2000), "A Review of the Effects of Financial Incentives on Performance in Laboratory Tasks: Implications for Management Accounting. *Journal of Management Accounting Research.*

Bonner, S. E., R. Hastie, S. M. Young, J. W. Hesford, and D. Gigone (1999), "Effects of Monetary Incentives on the Performance of a Cognitive Task: The Moderating Role of Skill," *University of Southern California Working Paper.*

Camerer, C. F. (1998), "Behavioral Economics and Nonrational Decision Making in Organizations." In J. Halpern and Sutton (eds.), *Decision Making in Organizations*, Cornell University Press.

Camerer, C. F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press.

Camerer, C. and R. Hogarth (1999), "The Effects of Financial Incentives in Experiements: A Review and Capital-Labor-Production Framework," *Journal of Risk and Uncertainty*.

Carpenter, J., G. Harrison and J. List (eds), "Field Experiments," *Research in Experimental Economics series*, JAI Press (forthcoming).

Conlisk, J. (1988), "Optimization Cost," *Journal of Economic Behavior and Organization*.

Cooper, D. J. and J. H. Kagel (1999), "Gaming Against Managers in Incentive Systems: Experimental Results with Chinese Students and Chinese Managers," *American Economic Review*.

Cosmides, L. and J. Tooby (1994), "Better than Rational: Evolutionary Psychology and the Invisible Hand," *American Economic Review*.

Cowan, N. (2001), "The Magical Number 4 in Short-term Memory: A Reconsideration of Mental Storage Capacity," *Behavioral and Brain Sciences.*

Dawes, R. M., D. Faust, and P. E. Mehl (1989), "Clinical versus Actuarial Judgement," *Science.*

Deci, E., R. Koestner and R. Ryan (1999), "Meta-analytic Review of Experiments Examining the Effects of Extrinsic Rewards on Intrinsic Motivation," *Psychological Bulletin.*

Devetag, G. and M. Warglien (2003), "Games and Phone Numbers: Do Short term Memory Bounds Affect Strategic Behavior?" *Journal of Economic Psychology.*

Dickhaut, J., K. Smith, K. McCabe, N. Peck, and V. Rajan (1997), " The Emotional and Mental Effort Dynamics of the English Auction," *University of Minnesota Working Paper*.

Eckel, C. (1999), "Comment on 'The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework'," *Journal of Risk and Uncertainty.*

Eisenberger, R. and J. Cameron (1996), "Detrimental Effects of Reward: Reality or Myth?" *American Psychologist.*

Ericsson, K. A. and J. Smith, (eds.), (1991), *Toward a General Theory of Expertise: Prospects and Limits,* Cambridge UK: Cambridge University Press.

Frey, B. S. (1997), "On the Relationship between Intrinsic and Extrinsic Work Motivation," *International Journal of Industrial Organization*.

Frey, B. S. and R. Jegen (2001), "Motivation Crowding Theory," *Journal of Economic Surveys*.

Gibbons, R. (1998), "Incentives in Organizations," *Journal of Economic Perspectives.*

Gigerenzer, G., R. Hertwig, U. Hoffrage, and P. Sedlmeier, "Cognitive illusions reconsidered," in Plott, C. R. and V. L. Smith (eds.), *Handbook of experimental economics results*, North Holland/Elsevier Press (forthcoming).

Glimcher, P. (2003), *Decisions, Uncertainty, and the Brain*, Cambridge: MIT Press.

Gneezy, U. and A. Rustichini (2000), "Pay Enough or Don't Pay At All," *Quarterly Journal of Economics.*

Gold, J. I. and M. N. Shadlen (2001), "Neural Computations that Underlie Decisions about Sensory Stimuli," *Trends in Cognitive Sciences*.

Goldsmith, M., A. Koriat and A. Weinberg-Eliezer (2002), "Strategic Regulation of Grain Size in Memory Reporting," *Journal of Experimental Psychology – General*.

Ham, J., J. H. Kagel and S. Lehrer (2001) "Randomization, Endogeneity and Laboratory Experiments in Economics," *Mimeo*, Ohio State University (accepted subject to editorial revision, *Journal of Econometrics*).

Hannan, R. L., J. H. Kagel and D. V. Moser (2002), "Partial Gift Exchange in an Experimental Labor Market: Impact of Subject Population Differences, Productivity Differences, and Effort Requests on Behavior," *Journal of Labor Economics*.

Harrison, G. W. (1989), "Theory of Misbehavior of First-Price Auctions," *American Economic Review*.

Harrison, G. W. (1992), "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review*.

Harrison, G. W., E. Johnson, M. McInnes, and E. E. Rutstrom (2003), "Risk Aversion and Incentive Effects: Comment," under review at *American Economic Review*.

Heckman, J. J. and S. Navarro-Lozano (2003), "Using Matching, Instrumental Variables and Control Functions to Estimate Economic Choice Models," *IZA Discussion Paper* No. 768.

Heckman, J.J. and E. Vytlacil (2001), "Identifying the Role of Cognitive Ability in Explaining the Level of and Change in the Return to Schooling," *Review of Economics and Statistics*.

Hertwig, R. and A. Ortmann (2001), "Experimental Practices in Economics: A Methodological Challenge for Psychologists?" *Behavioral and Brain Sciences*.

Hertwig, R. and A. Ortmann (2003), "Economists' and Psychologists' Experimental Practices: How They Differ, Why They Differ, and How They Could Converge," in Brocas, I. and J. Carrillo (eds.), *Economics and Psychology*. New York: OUP.

Hogarth, R. M., B. J. Gibbs, C. R. M. McKenzie and M. A. Marquis (1991), "Learning from Feedback: Exactingness and Incentives," *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

Holt, C. A. and S. K. Laury (2002), "Risk Aversion and Incentive Effects," *American Economic Review*.

Imbens, G. W. (2003), "Semiparametric Estimation of Average Treatment Effects under Exogeneity: A Review," forthcoming in *Review of Economics and Statistics*.

Jamal, K. and S. Sunder (1991), "Money vs. Gaming: Effects of Salient Monetary Payments in Double Oral Auctions," *Organizational Behavior and Human Decision Processes.*

Jenkins, D., Jr., A. Mitra, N. Gupta and J. Shaw (1998), "Are Financial Incentives Related to Performance? A Meta-Analytic Review of Empirical Research," *Journal of Applied Psychology.*

Johnson, E. J., C. F. Camerer, S. Sen, and T. Rymon (2001), "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining," *J. of Economic Theory*.

Kagel, J. H. and D. Levin (1986), "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*.

Kahneman, D., S. W. Peavler and L. Onuska (1968), "Effect of Verbalization and Incentive on the Pupil Response to Mental Activity," *Canadian Journal of Psychology*.

Kahneman, D. and S. W. Peavler (1969), "Incentive Effects and Pupillary Changes in Association Learning," *Journal of Experimental Psychology*.

Koriat, A. and M. Goldsmith (1996a), "Memory Metaphors and the Real-life/Laboratory Controversy: Correspondence Versus Storehouse Conceptions of Memory," *Behavioral and Brain Sciences*.

Koriat, A. and Goldsmith, M. (1996b). Monitoring and Control Processes in the Strategic Regulation of Memory Accuracy," *Psychological Review*.

Kreps, D. M. (1997), "Intrinsic Motivation and Extrinsic Incentives," *American Economic Review*.

Lazear, E. P. and S. Rosen (1981), "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*.

Levinsohn, J. and A. Petrin (2003), "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*.

Libby, R. L. and M. G. Lipe (1992), “Incentives, Effort, and the Cognitive Processes Involved in Accounting-Related Judgments,” *Journal of Accounting Research.*

Merlo, A. and A. Schotter (1999), “A Surprize-Quiz View of Learning in Economic Experiments,” *Games and Economic Behavior.*

Milgrom, P. and J. Roberts (1992), Economics, Organization and Management, Englewood Cliffs, N.J.: Prentice Hall.

Mincer, J. (1997), “The Production of Human Capital and the Life Cycle of Earnings: Variations on a Theme,” *Journal of Labor Economics.*

Moffatt, P. G. (2003), “Econometric Analysis of Effort Applied by Subjects in Economic Experiments,” *unpublished manuscript*.

Ortmann, A. and G. Gigerenzer (1997), “Reasoning in Economics and Psychology: Why Social Context Matters,” *Journal of Institutional and Theoretical Economics.*

Prendergast, C. (1999), “The Provision of Incentives in Firms,” *Journal of Economic Literature.*

Reber, A. S. (1989), “Implicit Learning and Tacit Knowledge,” *Journal of Experimental Psychology: General.*

Samuelson, W. F. and M. H. Bazerman (1985), “The Winner’s Curse in Bilateral Negotiations,” *Research in Experimental Economics.*

Siegel S., A. Siegel and J. Andrews (1964), *Choice, Strategy and Utility*, McGraw-Hill.

Smith, V. L. (1962), “An Experimental Study of Competitive Market Behavior, *Journal of Political Economy*.

Smith, V. L. and F. Szidarovszky (2001), “Monetary Rewards and Decision Cost in Strategic Interactions,” *unpublished manuscript*.

Smith, V. L. and J. Walker (1993), “Rewards, Experience and Decision Costs in First Price Auctions,” *Economic Inquiry*.

Sprinkle, G. B. (2000), “The Effect of Incentive Contracts on Learning and Performance,” *The Accounting Review*.

Taber, Ch. R. (2001), “The Rising College Premium in the Eighties: Return to College or Return to Unobserved Ability?” *Review of Economic Studies.*

Vandegrift, D., and P. Brown (2003), “Task Difficulty, Incentive Effects, and the Selection of High-Variance Strategies: an Experimental Examination of Tournament Behavior,” *Labour Economics.*

Wilcox, N. (1993), “Lottery Choice: Incentives, Complexity, and Decision Time,” *The Economic Journal.*