

The Agent's Ethics in the Principal-Agent Model

Øyvind Bøhren

ABSTRACT. This paper evaluates the current use of the Principal Agent Model (PAM) in accounting and finance, focusing on the agent's use of private information. The agent's behavioral norms in the the PAM deviate from commonly held ethical values in society, from models of man in conventional economic theory, and also from behavioral foundations of related business school fields like corporate strategy, business ethics, and human resource management. Still, it would be unwise to reject the PAM solely because of its distasteful ethical assumptions. The model does appear to have predictive power, but its descriptive or normative qualities remain unexplored. The popularity of the PAM, with its extreme model of man, raises fundamental questions about the impact of this model on business school stakeholders and society at large.

The principal-agent model (PAM) of Jensen and Meckling (1976) clarifies, extends, and formalizes earlier ideas on conflicts of interest between organizational stakeholders and the mechanisms for solving such conflicts.¹ The paper initiated an active and increasingly influential research program, and it was soon predicted that due to agency-based research, "the foundations are being put into place for a revolution in the science of organizations" (Jensen, 1983, p. 319).

Just four years after this prediction, Ross (1987) called agency theory the central approach

to the theory of managerial behavior. Today, the PAM is widely used in economic disciplines like accounting, micro-economics, and finance. The principal-agent framework has also entered marketing (Bergen et al., 1992) and non-economic social sciences like sociology, organizational behavior, and political science (Eisenhardt, 1989).

Despite the widespread adoption of the principal-agent approach in the social sciences, there have been rather few attempts at evaluating the PAM from a methodological point of view. Kaen et al. (1988) relate the intellectual roots of agency theory to political ideologies in the U.S. Walker (1989) finds that regardless of application area, most versions of the PAM will be hard to falsify, mainly because the validity of required auxiliary hypotheses cannot be checked. Comparing the model to competing theories in sociology and organizational behavior, Perrow (1986) is highly skeptical, primarily because the PAM ignores the cooperative aspects of social life. Brennan (1994), who argues strongly against including self-interest in the definition of rationality in economic theory, criticizes the PAM for being a too cynical model of human nature.

On the other hand, Eisenhardt (1989) concludes her survey of the principal agent literature rather enthusiastically, feeling that the PAM offers organization theory "a unique, realistic, and testable perspective on problems of cooperative effort".

This paper extends the methodological literature on the PAM by evaluating the model as currently used in accounting and finance. The focus is exclusively on the PAM's implicit assumption about the agent's indifference to dishonest use of private information, which has so far only been indirectly addressed in the

Øyvind Bøhren is a professor of finance at the Norwegian School of Management. His current research is on dividend policy, intercorporate shareholdings, and equity flotation methods. A paper on the philosophy of science in stochastic choice theory was recently published in the Journal of Economic Psychology.

economic literature (Noreen, 1988; Hausman and McPherson, 1993; Brennan, 1994).

The rest of the paper is organized as follows. Section 1 presents the logic of the PAM, emphasizing the role of the agent and clarifying the implicit ethical assumptions. The model is evaluated from a positive, normative, and educational perspective in section 2, followed by the concluding third section.

1. The agent in the principal-agent model

The principal in the PAM delegates decisions to an agent who has private information about his abilities, efforts, and the outcomes of a stochastic process. The two parties have conflicting goals and possibly different risk attitudes, and the economic problem is to design contracts which maximize the principal's expected utility of wealth, given that the agent will maximize his own expected utility and also never accept a contract which offers less than his best alternative.

According to the PAM, the agent's expected utility EU is determined by his wealth and by the effort he must expend to attain this wealth:²

$$EU = f(\text{Wealth}, \text{Effort}) \quad (1)$$

In this model, EU always increases with wealth (pecuniary nonsatiation) and decreases with effort (disutility of effort). For instance, if the PAM reflects a labor contract, the agent (employee) is supposed to prefer more wage to less at any compensation level, and he feels additional work is unpleasant regardless of hours worked. Expending effort only makes sense because it increases material consumption opportunities through the wage. Moreover, as both the principal and the agent derive utility only from their private efforts and pecuniary payoffs, the distribution of wealth and effort *per se* is irrelevant. This lack of concern for equity in the PAM means, for instance, that the principal never feels bad because he thinks the agent works too long hours or gets too low pay.

The PAM's behavioral assumptions of effort aversion, pecuniary non-satiation, and no distributional concerns are fully consistent with

standard modelling practice in micro-economics, where disutility of effort is implied by assuming that leisure utility dominates time spent at work. These three assumptions, which we will call materialistic egoism, have often been criticized for being extreme, amoral and unethical, particularly by non-economists.³ However, this paper is not still another analysis of materialistic egoism in economic models. Rather, we will focus on the PAM's implicit assumption that the agent has no moral barriers against dishonest behavior. This aspect of the agent's ethics is implied by his utility function as modelled by (1).

To see this more clearly, consider the classic micro-economic model, where each participant in a transaction has *symmetric* (i.e., no private) information, whether deterministic or probabilistic. Effort is the agent's only decision variable, and his preferences are supposed to be captured by the expected utility formulation in (1), where the decision variable enters both directly as effort and indirectly by the impact of effort on wealth. The agent's decision problem is to choose the effort level which optimally trades off the effort's positive wealth effect against its negative impact on effort aversion. At the optimal effort level, the marginal utility of wealth will just offset the marginal disutility of effort.

If we move on from the classic micro-economic model by introducing *asymmetric* (private) information and also assume that the agent knows more than the principal, the number of decision variables is increased from one to two: Besides effort, the agent can also use private information to influence his well-being. If we were to apply the same preference setup to the second decision variable (use of private information) as the classic symmetry model does for the first (effort), the agent's expected utility EU* would be:

$$EU^* = g(\text{Wealth}, \text{Effort}, \text{Information use}) \quad (2)$$

In (2), the agent's use of information may influence his well-being both indirectly through its impact on wealth and effort and also *directly* as a separate consequence variable. Thus, the third term in (2) may capture consequences of information use which the agent does care about, but

which are not picked up as a personal wealth effect or an effort effect in the two first terms.

As already noted, the PAM does not model the agent's preferences by (2), but rather by (1). Thus, the PAM really starts out with the classic micro-economic model, introduces private information for the agent as an additional decision variable, but retains the classic model's utility function, which only allows for a direct utility effect of effort. This is equivalent to assuming that the agent derives no discomfort *per se* from utilizing his private information in dishonest ways. In order to optimize the mix between high wealth and low effort, the agent uses private information without any ethical restrictions. For instance, the agent will gladly lie about his product's quality (Akerlof, 1970), the firm's output (Diamond, 1984), the value of the shares he offers for sale to new stockholders (Myers and Majluf, 1984), or about his accident proneness in insurance contexts (Rasmusen, 1989, p. 143). The agent will unscrupulously embezzle (Rasmusen, 1989, p. 147), and he tries to shirk and steal by excessive perk consumption (Jensen and Meckling, 1976).

Even though the agent in the PAM feels no discomfort with dishonest use of private information, his concern with wealth and effort may still make an inherently dishonest agent act honestly. This happens because the agent's preferences are well known to the principal, who may design a contract providing the agent with sufficient monetary and effort-based incentives to refrain from dishonesty. If such an honesty-ensuring contract is not optimal for the principal, he will take the predictable dishonesty (e.g., overstated production figures) into account when remunerating the agent for his effort. As the agent competes with other agents who have the same type of preferences, he must personally carry the anticipated costs of his dishonesty. This fact gives the agent incentives to convince the principal that he will indeed behave honestly.

In this sense, nobody is fooled or even disappointed in the equilibrium solution of the PAM, and the agent's materialistic egoism and the contract design may indeed collectively make him end up acting honestly. From an ethical perspective, however, the important point is that

the agent does not refrain from dishonest information use because he feels dishonesty is morally wrong. The only reason an agent happens to behave honestly is that such behavior pays off better than dishonesty in terms of the wealth-effort mix. In fact, in multiperiod versions of the PAM, the agent may be investing in a reputation for trust by acting honestly for many periods (Dasgupta, 1988). Again, this solution is forced upon him by market forces, as such behavior maximizes the present value of future wealth-effort combinations. It is not due to a feeling that honest use of private information has a value by itself, unrelated to wealth and effort.

This lack of moral restrictions on feasible actions makes the agent in the PAM differ radically from the agent in the classic micro-economic model, where each market participant behaves honestly: Every consumer and producer meet their budget constraints, and everybody delivers or picks up the goods they promised to trade when the deal was made. However, what may look like an ethical standard is really no independent behavioral assumption in the classic model. It is simply implied by the fact that when everybody is equally informed, everything is transparent to everybody, and nobody can be fooled. Therefore, it makes little sense to discuss the ethics of information use under symmetric information.

However, this implication from information structure to information ethics does not hold when knowledge is asymmetrically distributed. After the asymmetric distribution of information between principal and agent has been determined, the modeller must still choose between alternative ethical standards for the agent's use of private information. In the PAM, the modeller has implicitly chosen the rather radical assumption that the agent does not care whether a given wealth-effort mix results from honest or dishonest use of his private information.

Finally, note that the dishonesty property of the PAM is not just the well-known materialistic egoism in disguise. Materialistic egoism follows from (1) by assuming pecuniary non-satiation and no distributional concerns for wealth or effort. As demonstrated by the third element of (2), however, the assumption of

materialistic egoism in (1) may very well be transferred to a principal-agent context *without* simultaneously introducing unrestricted dishonesty. That is, the agent may be modelled as being egoistic in terms of wealth and effort, but he does not have to be indifferent to honesty issues as well.

2. Evaluating the principal-agent model

Most of us would probably feel that the agent's behavior in the PAM is morally questionable, and that we would dislike to interact frequently with persons whose preferences comply with (1). Without much hesitation, we might also assert that we certainly do not act this way ourselves, and that our preferences are better described by (2), with a negative direct utility effect of lying, stealing or cheating. However, we would certainly not deny having experienced lying agents, and we would undoubtedly admit that theft, cheating, and misreporting are very real problems in many organizations. Also, we all personally experience social contexts where the temptation to be dishonest is prevalent.

With this initial response to the agent's ethics in mind, the issue is whether the PAM is a good model or not. This question cannot be meaningfully answered until we specify what *purpose* the PAM is supposed to serve. This section focuses on positive, normative, and educational objectives of the PAM.

2.1. The PAM as a positive theory

Positive models aim at predicting, describing, or explaining real-world phenomena. With such a purpose, the question is whether the predictions, descriptions or explanations of the PAM stand up better against empirical evidence than competing theories do.

Prediction. With a predictive objective, it makes less sense to criticize a theory's assumptions. What matters is the nature of the theory's propositions and their performance in empirical tests; not the descriptive realism of the axioms from which the propositions are derived (Friedman,

1953). Hence, the critical issue is not if the PAM produces a good photograph or an X-ray of real-world principals and agents, including the agent's ethical standard implicit in (1). Rather, the question is whether the relationship can be successfully modelled *as if* contracting parties in principal-agent contexts typically act in a PAM-consistent way.⁴

From a predictive point of view, the PAM has generated remarkably innovative hypotheses. The model offers a rich set of testable predictions in accounting and finance, often in areas where conventional theories are silent. In corporate finance, for instance, the PAM predicts that a firm's financial leverage is positively associated with firm value, the default probability, the extent of regulation, the free cashflow, and its liquidation value. The model predicts that the stock price will fall when the firm announces a new equity issue, and that the stock price will be unaffected by changes in riskless debt. The stock price is hypothesized to increase when golden parachutes are adopted, or if restrictive voting rights are eliminated.⁵ In most of these cases, alternative finance theories have little or nothing to offer in terms of identical or competing predictions.

The empirical track record of the PAM is based on numerous tests from the last ten years. Many studies find statistically significant supporting evidence on the model's predictions about financial leverage, dividend policy, and share prices.⁶ Contrary to the PAM's predictions, however, there seems to be only a weak or possibly no association between corporate performance and managerial compensation, and managers seem to get a disproportionately high fraction of multiple-vote supershares (Jensen and Murphy, 1990; Loderer and Martin, 1994).

Description and explanation. When the purpose of a positive theory changes from prediction to description and explanation, the criteria for success change as well. If one studies an employer-employee relation in order to evaluate, copy, restructure or regulate it, the major concern is to characterize the parties as precisely as possible. Clearly, when predictions do not hold up empirically in properly designed tests, the

model lacks descriptive power as well. For instance, according to Mueller (1992), the reason the PAM's hypotheses about performance-dependent contracts get unconvincing empirical support is not that agency costs are unimportant in practice. Rather, this phenomenon is observed because in most firms, shareholders neither hire and fire managers nor write their contracts (Vancil, 1987). Thus, according to Mueller, a basic premise of the PAM is descriptively wrong, as managers tend to select themselves and to design their own contracts.

The less obvious case occurs when the PAM accurately predicts the equilibrium solution to a real-world conflict of interest, but still fails at describing basic properties of the parties. The model may simply give *the wrong reason* why a specific phenomenon is observed. For instance, if the manager in a firm with high free cashflow gets a substantial part of his pay in terms of call options on the firm's stock, the PAM would argue that if the manager were instead given performance-independent pay, he would shirk, buy another corporate jet and remodel his office every month. However, the real reason options are used may be that such an employment contract reduces the manager's tax bill compared to the tax on a fixed wage. The PAM is predictively right, but for the wrong reason. Its predictive performance is high, but the explanatory and descriptive power is low.

In general, as long as there is no one to one (if and only if) relationship between axioms and hypothesis, predictive success does not imply high descriptive and explanatory power. Moreover, because the PAM is virtually the only formal model in accounting and finance dealing with conflicts of interest under asymmetric information, it becomes particularly questionable to grant the PAM descriptive and explanatory qualities based solely on observed predictive power. We need competing models for conflicts of interest which, from different assumptions about principals and agents, make independent predictions about cost-effective ways of solving agency problems.

The PAM and existing theories of ethics. To get a better feeling for alternative ways of modelling

principal agent contexts, it may be useful to relate (1) and (2) to existing theories of ethics. The model of man as stated by the PAM in (1) is a special case of the more general hypothesis that your actions are only motivated by your self-interest, i.e., by materialistic and non-materialistic consequences which affect yourself. This theory of human nature goes back to the Sophists in ancient Greece and was further developed two millennia later by Hobbes (1651). Such behavior is called psychological egoism or just egoism, and 20th century theories along these lines are termed models of ethical egoism (Parfit, 1984).

Although (2) contains a direct utility effect of information use which is unrelated to personal wealth and effort, it may still reflect ethical egoism as defined above. For instance, if you feel embarrassed if your dishonesty is uncovered, it is clearly not in your best self-interest to disregard the possibility of such personal non-materialistic consequences when your decisions are made. Thus, although (2) breaks away from the materialistic egoism of (1) as defined in economic models, (2) may still be consistent with the more general notion of ethical egoism, which allows for both materialistic and non-materialistic egoism. Within such a framework, (1) may even be considered a special case of (2): Unlike the self-interested agent in (2), the self-interested agent in (1) disregards personal non-materialistic consequences.

There are still two reasons why (2) complies with a much wider set of ethical theories than just ethical egoism. First, the model's third argument allows for a *tradeoff* between materialistic and non-materialistic consequences. Second, the third argument may capture the impact of the agent's actions on others, even when these effects do not eventually fire back on the agent as personal consequences later on (like the impact of present actions on future generations). If such non-egoistic effects matter to the agent, there are elements of *altruism* in his preferences.

The extreme version of (2) reflects Kant's (1785) categorical imperative: Certain rules of social conduct (like truth-telling) are absolute duties to society and apply to all its members in any context. This means the preference ordering is lexicographic with respect to categorical con-

sequences: Wealth and effort cannot be traded off against dishonesty under any circumstances by any agent.⁷

In between Hobbe's psychological egoism and Kant's categorical imperatives, there are several ethical theories which allow for a tradeoff between self-interest and altruism. For instance, the moral sense school of Shaftesbury (1711) hypothesized a natural drive for harmony between virtue and selfishness, whereas Butler (1726) theorized about conscience as a restriction on self-interested behavior. Mill (1863) and Sidgwick (1874), the founders of Utilitarianism, argued along similar lines.

Honesty and tradeoff. Tradeoff theories imply that the more negative the direct utility effect of dishonest information use (be they personal or altruistic), the more the agent is willing to sacrifice in terms of higher effort and less wealth in order to act more honestly. Thus, unlike (1), the preference ordering in (2) may reflect an underlying hypothesis that every man has his price, and it may accomodate Reder's (1979) suggestion of measuring your honesty by the amount required to make you lie. It also describes why a basic desire to behave decently may be harder to transform into decent actions the more it costs you in terms of reduced wealth and more effort.⁸ Moreover, (2) accomodates the fact that the willingness to trade off (i.e., the implicit price on honesty) may differ from person to person, depending on factors like economic status, gender, type of education, profession, and internalization of social norms.

These ethical tradeoff theories seem directly relevant for generating alternative hypotheses about behavior and contract types in principal agent settings. Notice, however, that even though they differ fundamentally from the PAM in terms of describing the agent, the tradeoff theories may still make predictions which are close to those made by the PAM. Thus, despite the fact that the agent considers both non-materialistic personal effects and also altruistic consequences, the predicted contract types may still be PAM-like. For instance, performance-based contracts may be optimal even if the agent feels that cheating is asocial. Truly enough, the optimal fixed

portion of total wages may be higher with (2) than in the best PAM contract. However, as long as the agent trades off ethical concerns against the personal wage-effort mix, an optimal solution under ethical tradeoff theories may still involve some dishonesty. Unlike in the PAM, however, this behavior is not observed *only* because lying influences personal wealth or effort. Therefore, the higher the tradeoff cost for the third element in (2), the smaller the need for honesty-enforcing contracts, i.e., the higher the ratio of fixed to total wages.

Along the same lines one may argue that the significance of the honesty variable in (2) varies systematically across transaction types. In particular, the behavioral assumption of the PAM may be descriptively successful when the impact of your actions on others is impersonal or even negligible, like in the neoclassical world of atomistic agents and perfect competition. However, in less anonymous settings involving face-to-face, repetitive interactions, social norms and basic personal values become more important for individual behavior.

Summarizing, the descriptive and explanatory properties of the PAM are practically unexplored. This issue cannot be settled until we have competing principal-agent models which, from different ethical assumptions make predictions about the same observable phenomena, such as performance pay vs. a fixed wage, debt vs. equity financing, high vs. low dividends, and risky vs. less risky projects. Competing principal-agent theories of this kind do not yet exist.

2.2. Normative use of the PAM

We have argued that the PAM has predictive power. Does that imply the PAM is also a viable normative model, i.e., does it provide agents and principals with good advice for designing labor contracts, bond covenants, dividend policy, or accounting systems?

In general, the answer is no; predictive fit does not imply normative power. One reason is the unexplored descriptive validity of the PAM discussed above (right prediction, but maybe wrong cause). Moreover, the PAM contains a

model of 'man' which runs counter to most agents' moral values. Because dishonest information use may have a negative direct utility effect which is not reflected in (1), the PAM prescribes behavior which is not in the best interest of agents whose preferences are not well described by materialistic egoism. Therefore, the PAM is not the best normative theory for an agent whose ethics is better described by the tradeoff theories of (2) rather than by (1).

This argument does not directly apply to the *principal*, as he has no private information. Still, a real-world principal may dislike behaving like the PAM's principal, who thinks of his employees as inherently unreliable individuals. Even if the principal were only concerned about maximizing personal wealth, he may feel it would have a demotivating and unproductive effect on the agent if he were to follow the advice of the PAM and set up a full battery of carrots and sticks. Rather than building on individualism, suspicion, and negative thinking, he would establish a principal agent relationship based on a team approach, mutual trust, and positive attitudes. His basic approach to dealing with colleagues, which he clarified in detail years ago in the business school's human resource course, makes him reluctant to implement the incentive contracts suggested by the PAM, which he was repeatedly exposed to in accounting, micro-economics, and finance classes. Thus, like we just argued for the honesty-concerned agent, predictive success does not imply normative power when real-world principals are poorly described by the PAM.

Although this logic may appear very reasonable and convincing, it may easily provide the principal with very poor advice. As we have seen, the empirical evidence suggests that the principal faces a real risk of getting dishonest agents. If these findings also reflect a high descriptive validity of the PAM (which is still unclear), it can be prohibitively costly to just stick to the idealistic hope that a principal's blind trust produces honest agents in the long run. Rather, it would be in the principal's best interest and maybe the only sustainable solution to design employment contracts according to the PAM's prescriptions, even if the principal's wellbeing were partially determined by altruism. In such a perspective, a

strong predictive track record would make the PAM a viable normative model as well. Not because the real-world principal likes or identifies himself with the model's principal, but simply because he cannot afford the alternative of not protecting himself against the agent's dishonesty. Once more, we notice the importance of empirically verifying the descriptive validity of the PAM.

2.3. Educational impact of the PAM

Teaching the PAM. According to the empirical findings surveyed by Frank et al. (1993), certain student personalities are attracted to certain academic fields, and the design of a particular course influences subsequent student behavior. In particular, economics programs seem to recruit a disproportionately high number of self-interested students, and economics courses based on mainstream economic models tend to discourage students' tendency to cooperate.

Noreen (1988) argues that behavioral norms are primarily passed on by instruction and probably reinforced by the absence of counter-examples: "If everyone thinks it is normal behavior to cheat and deceive, then people will cheat and deceive without feeling guilty. . . . [I]f people are instructed that they should not steal, and they do not observe anyone stealing, they will tend not to steal".

On this background, the question is whether academia's exposition of agency theory influences the tendency of real-world agents and principals to expect and even accept dishonesty as the prevailing moral norm outside the class room.

There are plausible reasons why this concern is warranted. Modern theories of accounting and finance are heavily influenced by agency-based arguments. Students read time and again about unscrupulous agents surrounded by carrots and sticks which keep them in line; about honesty which does not come from the heart, but is forced upon them by economic market forces. Eventually, students (and maybe even professors) may come to feel that the agent's personality as modelled by the PAM is not that extreme after all. Or, being reminded that the PAM is just a

model, they may at least stop reacting to its ethical content, particularly if the professor does not highlight the implicit ethical assumptions and what alternative models may look like.

After graduation, the former student may think, like the finance class seemed to suggest, that the real-life job environment is already set up for agents who cannot be trusted because they tend to overreport performance. In such a world, it appears economically optimal for you as an agent to overreport, as the principal is likely to reduce the performance measure you give him anyway before using it to calculate your wage. Eventually, an organizational design which presupposes dishonest agents may end up producing such individuals.

There are several counterarguments to this critique. First, the PAM was never meant to be normative and maybe not even descriptive by its founders (Jensen, 1983). Second, it can always be questioned whether students change ethical values by going to business school, reading finance textbooks, and by listening to accounting professors. Moreover, models are simplifications, designed to represent certain parts of reality at the expense of others. In that perspective, it can validly be argued that the important contribution of the PAM is that it breaks away from the black box approach to the theory of the firm and also from the harmony model of frictionless delegation. Unlike other models, the PAM makes testable predictions of how contracts between organizational stakeholders minimize the costs of production in a world of conflicting objectives, private information, and inherently dishonest agents.

The costs of dishonest information use. Given our focus on finance and accounting, maybe the most uncontroversial way of clarifying our thinking about educational impacts of the PAM is by appealing to standard economic arguments. With unconditional honesty (like if truth-telling were a categorical imperative), no resources would be wasted by dishonest information use or by efforts to control it. The principal could simply trust the agent on his word, and the first-best solution would be attained.

Enter the inherent information asymmetry and

the ethics of the parties in the PAM, it is equally clear that given this setting, both the agent and the principal sleep well at night and feel no regret or disappointment whatsoever when observing the other party's behavior. They know they have the best contract, given their situation. Nevertheless, as long as the principal cannot fully observe the agent's effort, the agent's work ethics prevents the first best solution from being realized, i.e., the allocation that would be obtained if the principal could blindly trust the agent. The difference between these two solutions is a deadweight loss to the contracting parties, and the agent must always carry parts of it. This deadweight loss consists of opportunity costs (e.g., the manager chooses a capital structure which does not maximize firm value) and out-of-pocket costs for bonding, certification and monitoring (e.g., a detailed system for verifying hours worked). Clearly, the net value gain would be positive whenever the agent's preferences could be modified at a cost which is smaller than the deadweight loss under optimal wealth- and effort-based incentives.

Hence, regardless of what one feels about the agent's ethical standard in (1), attaining the cost-effective, optimal contract in the PAM is certainly not costless to the contracting parties and generally more costly than if the agent experienced a negative direct utility effect of dishonesty, like in (2). Moreover, on an aggregate level, Arrow (1975) argues that a society's behavioral norms is a valuable economic asset in any private transaction because shared, internalized rules of conduct (like truth, trust, loyalty, and justice) reduce contracting costs, and thus facilitate the process of exchange in a market economy. Hausman and McPherson (1993) even assert that without such virtues, economic life would come to a halt.

The positive effect of an increased trust between asymmetrically informed parties is a welfare-economic justification for taking educational impacts of the PAM seriously. Without running the risk of pushing one's own ethics upon others, it can at least safely be concluded that accounting and finance professors should be more explicit to their students on the ethical assumptions of the PAM. Also, the relevant

ethical alternatives should be highlighted, and it should be pointed out that the agent's insensitivity to honesty issues in the PAM may have negative economic externalities by increasing the future transaction costs for society at large.

Clearly, discussions along these lines is not current teaching practice in accounting and finance. In fact, Noreen (1988) in accounting and Brennan (1994) in finance seem to be notable exceptions to the overall impression that well-known researchers in these two fields have not yet seriously addressed educational effects of the PAM. According to Noreen (1988): "to a large extent, we conduct our research by grubbing around in the underbrush of business looking for evidence of opportunistic behavior. We need to keep in mind that while our models presume that businessmen are opportunistic and that we can find (statistically significant) examples of such opportunism, it does not follow that businessmen generally are or should be opportunistic." Brennan (1994), who criticizes the PAM's reliance on unrestricted self-interest, feels that if we go on teaching principal-agent issues the way we currently do, we may eventually find that the theory becomes descriptively true, but at the cost of producing a society that will not function.

3. Concluding remarks

Behavioral models are not value-neutral, and theories may have social power. In particular, popular theories are probably more easily internalized and put to practical use than less popular ones by professors, students, and managers alike.⁹

The Principal-Agent model (PAM) is currently a very productive vehicle in the theory development process of accounting and finance. The model involves a non-standard, pessimistic, and rather extreme model of man which has been known in the theory of ethics for centuries. Given the popularity of the PAM, the lack of methodological discussions in the existing PAM literature, and also the increasing attention paid to ethics in both business schools and industry, three issues seem particularly important to address.

First, the descriptive validity of the PAM

should be explored in a much more extensive way. Second, principal-agent models should be developed which do not presuppose that agents are indifferent between being honest and dishonest, and which allow for the possibility that honesty makes sense even if it does not increase your wealth or reduce your effort. Third, accounting and finance professors should address the educational impact of the PAM in a more explicit and systematic way by pointing out the model's implicit ethical assumptions, by addressing the externality costs to society caused by private dishonesty, and by clarifying alternative ways of modelling ethics in principal agent settings.

Acknowledgements

I acknowledge helpful comments to earlier versions of this paper from Morten Berg, Michael J. Brennan, Lars Fallan, Frøystein Gjesdal, James Horrigan, Jon Lundesgaard, Eric Noreen, Helge Rynning and Fred Wenstøp. This research was initiated during my 1991–92 sabbatical at the University of North Carolina, Chapel Hill.

Notes

¹ Some frequently cited examples are Adam Smith (1776, p. 700), Berle and Means (1932), Baumol (1959), Williamson (1964), Alchian and Demsetz (1972), and Fama and Miller (1972). The basic assumptions and the modelling tools of the PAM come from the moral hazard and adverse selection literature in information economics. See Akerlof (1970) and Spence (1973) for early examples and Rasmusen (1989) for an introduction and overview.

² The expected utility theorem is general in the sense that under certain regularity conditions, preferences may be defined over any number of consequences, which may be of any nature, like non-economic, personal consequences or any externality for society at large (Keeney and Raiffa, 1976). In the PAM as applied to accounting and finance, however, personal wealth and personal effort are the only consequences modelled.

³ Perrow (1986), Hirsch et al. (1987), Horrigan (1987) and Etzioni (1988) are recent examples.

⁴ A corresponding case is predictive use of the

expected utility theorem. The issue is not whether decision makers actually perform expected utility calculations, but if their choice between risky actions can be well approximated by such a model, i.e., as if real-world decision makers computed the expected utility numbers and picked the action with the highest number.

⁵ See Harris and Raviv (1991) for references and more examples.

⁶ Harris and Raviv (1991) and Eisenhardt (1989) provide partial surveys.

⁷ As a lexicographic preference ordering does not have an expected utility representation (Keeney and Raiffa, 1976), it is formally incorrect to list this as a special case of (2), which is stated in expected utility terms.

⁸ Principal-Agent models dealing with pooling vs. separating equilibria have two exogenously given types of agents; honest ones and dishonest ones. As honesty is not a separate argument in either type's utility function, neither agent makes a tradeoff between ethics and the wealth-effort mix. Thus, the idea that every man has his price cannot be properly modelled in this setting, as the price of honesty is either assumed to be zero (dishonest) or infinite (honest). The dishonest type never stops lying because he feels morally bad, and the honest type will always stick to honesty regardless of its materialistic costs.

⁹ According to Perrow (1986, p. 235), "theories shape our world; they encourage us to see it a certain way, and then we exclude other visions that could direct our actions". Perrow criticizes the PAM on this background because he feels it models principals at the mercy of agents, that monitoring is only available to the principal, that organized actions by workers (like unionization) are absent, that the PAM ignores how job-training and socialization reduces real-life information asymmetries, and that it implies a political ideology which is "both blatant and conservative" (Perrow 1986, p. 220).

References

- Akerlof, G.: 1970, 'The Market for Lemons: Quality Uncertainty and the Market Mechanism', *Quarterly Journal of Economics* **84**, 488–500.
- Alchian, A. A. and H. Demsetz: 1972, 'Production, Information Costs, and Economic Organization', *American Economic Review* **62**, 777–795.
- Arrow, K.: 1975, 'Gifts and Exchanges', in E. S. Phelps (ed.), *Altruism, Morality, and Economic Theory* (Russell Sage Foundation, NY), pp. 13–28.
- Baumol, W. J.: 1959, *Business Behavior, Value and Growth* (Macmillan, NY).
- Bergen, M., S. Dutta, and O. C. Walker, Jr.: 1992, 'Agency Relationships in Marketing: A Review of the Implications and Applications of Agency and Related Topics', *Journal of Marketing* **58**, 1–24.
- Berle, A. and G. Means: 1932, *The Modern Corporation and Private Property* (Macmillan, NY).
- Brennan, M. J.: 1994, 'Incentives, Rationality and Society', *Journal of Applied Corporate Finance* **7**, 31–39.
- Butler, J.: 1726, *Fifteen Sermons*.
- Dasgupta, P.: 1988, 'Trust as a Commodity', in D. Gambetta (ed.), *Trust: Making and Breaking Cooperative Relations* (Basic Blackwell, NY).
- Diamond, D. W.: 1984, 'Financial Intermediation and Delegated Monitoring', *Review of Economic Studies* **51**, 393–414.
- Eisenhardt, K. M.: 1989, 'Agency Theory: An Assessment and Review', *Academy of Management Review* **14**, 57–74.
- Etzioni, A.: 1988, *The Moral Dimension-Toward a New Economics* (The Free Press, NY).
- Fama, E. F. and M. H. Miller: 1972, *The Theory of Finance* (Holt, Rinehart, and Winston, NY).
- Frank, R. H., T. Gilovich and D. T. Regan: 1993, 'Does Studying Economics Inhibit Cooperation?', *Journal of Economic Perspectives* **7**, 159–171.
- Friedman, M.: 1953, 'The Methodology of Positive Economics', in M. Friedman (ed.), *Essays in Positive Economics* (University of Chicago Press, Chicago, IL).
- Harris, M. and A. Raviv: 1991, 'The Theory of Capital Structure', *Journal of Finance* **46**, 297–355.
- Hausman, D. M. and M. S. McPherson: 1993, 'Taking Ethics Seriously: Economics and Contemporary Moral Philosophy', *Journal of Economic Literature* **31**, 671–731.
- Hirsch, P., S. Michaels and R. Friedman: 1987, 'Dirty Hands versus Clean Models: Is Sociology in Danger of being Seduced by Economists?', *Theory and Society* **16**, 317–336.
- Hobbes, T.: 1651, *Leviathan*, ed. C. B. Macpherson (Penguin Group, London, 1968).
- Horrihan, J. O.: 1987, 'The Ethics of the New Finance', *Journal of Business Ethics* **6**, 97–110.
- Jensen, M.: 1983, 'Organization Theory and Methodology', *Accounting Review* **58**, 319–338.
- Jensen, M. and W. Meckling: 1976, 'Theory of the Firm: Managerial Behavior, Agency Costs, and Ownership Structure', *Journal of Financial Economics* **3**, 305–360.
- Jensen, M. and K. J. Murphy: 1990, 'Performance Pay

- and Top-Management Incentives', *Journal of Political Economy* **98**, 225–264.
- Kaen, F. R., A. Kaufman and L. Zacharias: 1988, 'American Political Values and Agency Theory – A Perspective', *Journal of Business Ethics* **7**, 805–820.
- Kant, I.: 1785, *Grundlegung zur Metaphysik der Sitten*, trans. H. J. Paton as *Groundwork of the Metaphysics of Morals* (Harper & Row, NY, 1964).
- Keeney, R. L. and H. Raiffa: 1976, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (John Wiley: NY).
- Loderer, C. and K. Martin: 1994, 'Executive Stock Ownership and Performance: Tracking Faint Traces', unpublished working paper, University of Bern and University of Iowa.
- Mill, J. S.: 1863, 'Utilitarianism', in M. Warnok (ed.), *Mill: Utilitarianism and other Writings* (Collins, Glasgow, 1962).
- Mueller, D. C.: 1992, 'The Corporation and the Economist', *International Journal of Industrial Organization* **10**, 147–169.
- Myers, S. C. and N. S. Majluf: 1984, 'Corporate Financing and Investment Decisions when Firms have Information that Investors do not Have', *Journal of Financial Economics* **13**, 187–221.
- Noreen, E.: 1988, 'The Economics of Ethics: A New Perspective on Agency Theory', *Accounting, Organizations and Society* **13**, 359–369.
- Parfit, D.: 1984, *Reasons and Persons*.
- Perrow, C.: 1986, *Complex Organizations: A Critical Essay* (Random House, NY).
- Rasmusen, E.: 1989, *Games and Information* (Basil Blackwell, NY).
- Reder, M.: 1979, 'The Place of Ethics in the Theory of Production', in M. Boskin (ed.), *Economics and Human Welfare: Essays in Honor of Boris Scitovsky* (Academic Press, NY), pp. 133–146.
- Ross, S. A.: 1987, 'The Interrelations of Finance and Economics: Theoretical Perspectives', *American Economic Review* **77**, 29–34.
- Shaftesbury, E.: 1711, *Characteristics of Men, Manners, Opinions, Times*.
- Sidgwick, H.: 1874, *The Methods of Ethics*.
- Smith, A.: 1776, *The Wealth of Nations*, Cannan edition (Modern Library, NY, 1937).
- Spence, A. M.: 1973, 'Job Market Signalling', *Quarterly Journal of Economics* **87**, 355–374.
- Vancil, R. F.: 1987, *Passing the Baton* (Harvard Business School, MA).
- Walker, M.: 1989, 'Agency Theory: A Falsificationist Perspective', *Accounting, Organizations and Society* **14**, 433–453.
- Williamson, O. E.: 1964, *The Economics of Discretionary Behavior: Managerial Objectives in a Theory of the Firm* (Prentice-Hall, Englewood Cliffs, NJ).

Norwegian School of Management,
Elias Smiths vei 15,
N-1300 Sandvika,
Norway.
E-mail: Oyvind.Bohren@bi.no