NAKE
Nieuws

Volume 17

December 2005

Netwerk Algemene en Kwantitatieve Economie

Netherlands Network of Economics

NAKE Secretariat

Astrid van Hemert
Faculty of Economics
Tilburg University
PO Box 90153
5000 LE TILBURG
Phone: 013-466 2516 / Fax: 013-466 3042
E-mail: nake@uvt.nl

Homepage

www.nake.nl

INDEX

Preface 2

Workshop report:
"Empirical Strategies"
By Anne Gielen 4

Workshop report:
"Schumpeterian Growth: Theory and Applications"
By Edwin van der Werf 20

Workshop report:
"The Crowding Out of Intrinsic Motivation"
By Fleur Wouterse 31

Upcoming in 2006 42

Preface

The year 2005 has been an unprecedented year of changes for the NAKE board and staff. Occasionally new political developments surfaced, which forced us to communicate with you via e-mail and letters more regularly. As a result, there was a less pressing need to publish *NAKE Nieuws*, which explains why there is only a single issue in 2005.

On October 21, 2005, our Annual Meeting of Members was held at the NAKE Research Day. This day marked a big change in the composition of NAKE's Management Board. After 19 tireless years of service to NAKE, Joan Muysken stepped down from the NAKE Management Board. Since the inception of NAKE in 1986, he has been closely involved in shaping NAKE's structure and activities, first as a member of the Management Board, and later as our Chairman. NAKE will greatly miss him. In addition, Peter Boswijk, Eric van Damme, André Lucas, and Jean-Marie Viaene left the NAKE Board. Please join me in thanking them for all their efforts in making NAKE a successful organization. I also would like to take this opportunity to introduce to you three newcomers. Lex Meijdam (UvT), a former NAKE director, has become our new Chairman. His hands on experience with NAKE will ensure a smooth transition at a time that many changes take place in NAKE's course program. We are happy to have him on board! Eelke de Jong (RU) and Thomas Ziesemer (UM) have become members of the NAKE Management Board.

Since September 2005, we have set up a new program of field courses that is aimed at second-year M.Phil. students. Each year we will offer a stable set of field course of 12 weeks of Friday lectures. The field courses are supplemented by a small segment of topics courses. More information on our course program can be found at the NAKE web site.

In December 2004 and June 2005, we had very successful workshops at Maastricht University and Utrecht University, which attracted roughly 70 and 50 participants, respectively. It has been a long time ago that we had such as a high workshop attendance. In this issue of *NAKE Nieuws* you will find three reports on the December 2004 workshop. Anne Gielen (Tilburg University) has given an excellent overview of the lectures of Joshua Angrist, who managed to involve many of you in his lectures on "Empirical Strategies." The second report is by Edwin van der Werf (Tilburg University), whose report is selected for the second time in a row. Edwin's report provides a reader-friendly coverage of the high-speed lectures of Philippe Aghion on "Schumpeterian Growth Theory." Finally, Fleur Wouterse (Wageningen University) has written a compact and lucid report on one of the lectures of Simon Gächter.

As you may have seen, the *NAKE Bulletin* has received an upgrade in the form of a nicer lay out and more information on NAKE activities. To better serve you with up-to-date information, we are currently thinking of integrating *NAKE Nieuws* into the *NAKE Bulletin.* We welcome any thoughts you may have on this.

I wish you a Happy New Year!

Best regards,

Jenny Ligthart

Empirical Strategies

Report by Anne Gielen*

1 Introduction

This paper provides a report on the lectures on empirical strategies in labor economics given by Joshua Angrist. Empirical strategies are very useful for applied work. Almost 80 percent of recent articles published in labor economics contains some empirical work. Two types of empirical work can be distinguished: descriptive analysis and causal inference. Descriptive analysis can establish facts about the labor market that need to be explained by theoretical reasoning and yield new insights into economic trends. In contrast, causal inference seeks to determine the effects of particular interventions or policies, and to estimate features of the behavioral relationships suggested by economic theory. In this paper the focus is on causal inference, that is, empirical strategies used to estimate features of causal relationships. In particular, estimation of average treatment effects is discussed. This has become important in the program evaluation literature, such as in the evaluation of job training programs. This paper is an attempt to provide a non-technical summary of the lectures, such that it is also a comprehensible reading for non-econometrists. This is done by including alternative explanations from other sources , that is, by other authors than Angrist, on the different topics.[1]

This paper is organized as follows. Section 2 provides an introduction in the basic features of empirical methods. In Section 3 it is shown how causal regression can be used to make inferences about what would have happened with a variable of interest if some of the explanatory variables were to change. This is especially interesting for

*Department of Economics, Tilburg University, P.O. Box 90153, 5000 LE Tilburg. Phone: (+31) 13 466 3216. Fax: (+31) 13 466 3042. E-mail: a.c.gielen@uvt.nl.

[1]Also, the chapter by Angrist and Krueger (1999) in the *Handbook of Labor Economics* has been of great help. If one would like to have more details on some of the results, it is recommended to consult Wooldridge (2002), who provides very comprehensive explanations for several topics in econometric analyses.

economists and politicians studying the effects of changing regulations. The rest of the paper focuses on the effects of a specific treatment. Section 4 discusses the differences between usual regression estimators and matching estimators to estimate the effect of a treatment. Since it is not unlikely that the treatment variable is correlated with one of the other explanatory variables, Section 5 describes some techniques on how to solve the problem of selection bias. Section 6 concludes.

2 Agnostic Regression

Over the years, empirical approaches have developed considerably. However, recent theoretical developments notwithstanding, the most important empirical tool is still old-fashioned regression analysis. Regression analyses are used to study the relationship between a dependent variable, y_i, and a vector of explanatory variables, x_i. Together, these variables make up the sample, which is assumed to be random, that is, independent and identically distributed (i.i.d.). Empirical economists are interested in the relationship between y_i and the explanatory variables. Hence, the goal of most applied econometric studies is to estimate or test hypothesis about the expectation of one variable, y_i, conditional on a set of explanatory variables. To illustrate, an example is used in which y_i will be referred to as wages, whereas x_i contains various individual characteristics, such as education and the number of years of experience or schooling. The relationship between the wage y_i and the explanatory variables x_i can be identified by taking three steps:

- Description: How does y_i vary with x_i?
- Prediction: Can we use x_i to forecast y_i?
- Causality: What is the effect of elements of x_i on y_i?

It is recognized that all relationships can be nonlinear and most are not deterministic. Typically, economists are interested in conditional expectations that allow them to infer causality from one or more explanatory variables to the dependent variable (Wooldridge, 2002). Rather, we look for relationships that hold "on average." This is expressed in the *Conditional Expectation Function* (CEF),

$$E[y_i|x_i] = h(x_i). \tag{1}$$

This function is the summary of the relationship between y_i and x_i and it shows how the average value of y_i changes as elements of x_i change. Note that, by definition, equation (1) implies $y_i = h(x_i) + \varepsilon_i$, where $E[\varepsilon_i|x_i] = 0$. This means that it is always possible to write y_i as its conditional expectation plus an error term, ε_i, that has conditional mean zero.[2] This is a very important property:

Theorem 1 *CEF residuals are mean-independent of conditioning variables, and therefore orthogonal to any function of the conditioning variables.*

This theorem states that the mean of the CEF residuals does not depend on the values and distribution of the explanatory variables. A related property is that the CEF solves a prediction problem:

Theorem 2 $E[y_i|x_i] = \arg\min_h E(y_i - h(x_i))^2$.

This implies that the CEF minimizes the mean-square error (MMSE) in predicting y_i given x_i.[3] Yet, regression is the best linear fit to the CEF. Therefore, the CEF is called the MMSE-predictor for y_i given x_i. A linear projection of y_i on a vector $x_i = (1, x_1, x_2, ..., x_K)$ always exists and is unique:

$$L(y_i|1, x_1, x_2, ..., x_K) = L(y|1, x_i) = \beta_0 + \beta_1 x_1 + ... + \beta_K x_K = \beta_0 + \beta_i x_i. \tag{2}$$

In some cases, then also the CEF is linear, that is, it is linear in the parameters β_i. The two most common sufficient conditions for a linear CEF are ,first, joint normality and, second, a saturated model for discrete regressors. In a saturated model, a discrete dependent variable y_i is explained by discrete explanatory variables only. Because the CEF is the minimum mean square predictor, that is, it gives the smallest mean square error out of all functions, it follows immediately that *if* $E[y_i|x_i]$ is linear in x then the linear projection coincides with the conditional expectation function (Wooldridge, 2002). So, if the CEF is linear, it coincides with the regression of y_i on x_i.

Why would you run a regression? There are three reasons to run a regression. A first reason to do so results from the *regression-CEF Theorem*, which states that if the CEF is linear, it is given by the regression of y_i on x_i. That is, if

$$E[y_i|x_i] = x_i'\pi, \tag{3}$$

[2] It is quite reasonable that $E[\varepsilon_i|x_i] = E[\varepsilon_i] = 0$. The fact that the error term has zero unconditional expectation follows as a special case of the law of iterated expectations. For details, see Wooldridge (2002).

[3] Again, this can be shown by the law of interated expectations (Wooldridge, 2002, p. 33)

then

$$\pi = E[x_i x_i']^{-1} E[x_i y_i], \tag{4}$$

where π is called the coefficient vector for the population regression of y_i on x_i.[4] It means that the estimated coefficient vector can be used to predict values for y_i that are in the population (but need not to be in the sample). By definition, π guarantees that x_i and ε_i are uncorrelated.

If you want to predict y_i, and you want to limit yourself to linear functions, π is the *Best Linear Predictor* (BLP) since it minimizes the mean-square error: $\pi = E[x_i x_i']^{-1} E[x_i y_i] = \arg\min_b E(y_i - b'x_i)^2$.[5] Accordingly, the *BLP Theorem* (or regression prediction theorem) provides another reason to run a regression.

Finally, if you are not interested in individual predictions and outcomes, but you are especially interested in the features of the distribution of y_i, you can use a regression since it provides the best approximation to the CEF in a MMSE-sense. This *Regression-Approximation Theorem* provides the third and most important reason to run a regression. It implies that predicting y_i given x_i is equal to predicting $E(y_i)$ given x_i. A nice thing about this theorem is that it holds independent of whether your regressors are stochastic or fixed, whether the residuals are uncorrelated with the regressors, whether the relationship of interest is truly linear, and whether the dependent variable is continuous or binary.

Ordinary Least Squares (OLS) can be used to estimate the population regression coefficient vector π, such that

$$\widehat{\pi} = [\Sigma x_i x_i']^{-1} \Sigma x_i y_i = \pi + [(1/N)\Sigma x_i x_i']^{-1}(1/N)\Sigma x_i(y_i - x_i'\pi). \tag{5}$$

It is important to check how $\sqrt{N}(\widehat{\pi} - \pi)$ behaves. That is, the distributional characteristics of the model need to be checked on their asymptotic properties.[6] First, according to the *Law of Large Numbers* the sample moments should converge to the population moments as the number of observations, N, goes to infinity: $(\Sigma_{i=1}^N y_i)/N \longrightarrow \mu_{y_i}$, where $\mu_{y_i} = E(y_i)$. Next to that, the *Central Limit Theorem* must be satisfied, that is $(N^{-1/2}\Sigma_{i=1}^N y_i - \mu_{y_i}) \sim N(0, Var(y_i))$. This implies that the sample is centered around the

[4] *Proof:* By the first theorem, it is always true that $E[x_i(y_i - E[y_i|x_i])] = 0$. So, $E[y_i|x_i] = x_i'\pi$ implies $E[x_i y_i] = E[x_i x_i']\pi$.

[5] Proof: $\pi = E[x_i x_i']^{-1} E[x_i y_i] = \arg\min_b E[(E[y_i|x_i] - b'x_i)^2]$. Then the partial derivative with respect to b yields $\frac{\partial \pi}{\partial b} : E[x_i(y_i - x_i'b)] = 0$.

[6] Note that the precise explanation of these properties comprises a rather mathematical calculation. For simplicity and comprehensibility these technical details are neglected in this paper. For details, see Wooldridge (2002), Sections 3.2, 3.4, and 3.5.

plim. Another property is *Slutsky's Theorem*: $p\lim g(y_i) = g(p\lim y_i)$ if $g(\cdot)$ is continuous at plim y_i. This theorem shows that the plim passes through any (non)linear functions of sample moments, provided they are continuous. The final property is the *Delta Method*: $\sqrt{N}(h(N^{-1}\Sigma_{i=1}^{N}y_i) - h(\mu_{yi})) \sim N(0, \Omega)$, where $\Omega \equiv \nabla h'(\mu_{yi})Var(y_i)\nabla h(\mu_{yi})$. These properties can be used to show that $\sqrt{N}(\widehat{\pi} - \pi)$ is asymptotically normal.

3 Causal Regression

The most challenging empirical questions in economics involve "what if" questions about counterfactual outcomes. In the example in the previous section, "what if" questions refer to the effects on the wage of additional schooling, union membership or military service. So, "What would the wage have been if this worker would have been a union member?," whereas in the sample this worker is *not* a union member. These questions are interesting for example from a policy point of view. However, also individual workers may want to know how their wage changes if they take a certain decision. These causal relationships involve comparisons of counterfactual states of the world. Causal regressions can be used to study these relationships. A regression is causal if it can be linked to a model for potential outcomes. Potential outcomes can be defined using an underlying functional relationship that describes what y_i would have been if x_i would have taken different values than they have in the sample. This means that the regression answers "what if" questions about what the distribution of outcomes would be like if the regressor of interest were to change. In the example in the previous section, the relationship between the number of years of schooling (x_i) and a person's wage (y_i) is of interest. This relation may be person-specific, so it can be written as,

$$y_{s,i} \equiv f_i(s), \tag{6}$$

to denote the *potential earnings* that person i would receive after obtaining s years of education, and not just for the realized value s_i. However, only the earnings of a person as a result of the number of years of education that this person has had in his life (s_i) is known. Then, (6) enables us to determine what the earnings of this person would have been if this person would have had less or more years of education (s). Note that the function $f_i(s)$ has an i subscript, while s does not. This indicates the fact that although s is a variable, it is not a random variable. Once this causal relationship has been defined,it can be linked to the observed association between schooling and earnings with a linear model:

$$f_i(s) = \beta_0 + \rho s + \eta_i. \tag{7}$$

This causal model assumes that $f_i(s) - f_i(s-1) = \beta$, that is, person homogeneity, such that for every one year difference in the years of schooling the wage changes with β. In addition, the equation says that β is the same for all individuals, that is, sample homogeneity. Note that this is a causal model and not a regression, because the random part η_i can be correlated with s_i, for example since schooling is a consequence of individual decisions and institutional forces it is likely that there is a correlation between η_i and s_i. Regression strategies try to overcome this problem, because assumptions about the functional form are made and the random part η_i is decomposed into a linear function of the observable characteristics x_i and an error ε_i, such that $\eta_i = x_i'\beta + \varepsilon_i$, where β is a vector of population regression coefficients. Therefore, by construction, ε_i and x_i are uncorrelated. The key assumption is that the observable characteristics x_i are the *only* reason why η_i and s_i are correlated, so $E[s_i\varepsilon_i] = 0$. This is the "selection on observables" assumption (Barnow *et al.*, 1981), in which the regressor of interest is assumed to be independent of potential outcomes after accounting for a set of observable characteristics. If the selection-on-observables assumption is obtained, the multivariate regression becomes

$$Y_i = x_i'\beta_r + \rho_r s_i + e_i, \tag{8}$$

which identifies the causal parameter of interest, ρ.

4 Regression Versus Matching

In the previous section, it was pointed out that it is difficult to answer "what if" questions, because the sample does not contain any information about other possible values for s than s_i. The same problem occurs if you want to estimate the effects of a treatment, say for example being in the military. This has been done by Angrist (1998) who studies the effects of military service on earnings. Earnings for people who have served in the military, y_{1i}, and for people who have not served in the military, y_{0i}, are observed. Denote veteran status by a dummy variable d_i, which has a value 1 if the worker has served in the military and 0 if not. For each person, we observe

$$y_i = y_{0i} + (y_{1i} - y_{0i})d_i. \tag{9}$$

That is, only wages for veterans are observable ($d_i = 1$), but not the wages for these individuals are observable if they would not have been in the army ($d_i = 0$). The same goes for non-veterans, for whom only wages after *not* being in the military ($d_i = 0$) are observed. We are interested in the earnings differences between people who did and did not serve in the military $E(y_{1i} - y_{0i})$, that is, the average causal effect or average treatment effect (ATE). ATE is defined as the expected effect of treatment on a randomly drawn person from the population (Wooldridge, 2002). Since for one person either y_1 or y_0 is observed, the average effect of the treatment (being in the military) on the treated (the ones who did serve in the military), which is denoted by ATE_1, has to be estimated:

$$E[y_{1i} - y_{0i}|d_i = 1] = E[y_{1i}|d_i = 1] - E[y_{0i}|d_i = 1]. \tag{10}$$

ATE_1 is the mean effect on the wage for those who actually served in the military. Note that the last part in this equation is not observed. However, if it is approximated by estimating $E[y_{1i}|d_i = 1] - E[y_{0i}|d_i = 0]$ the result is biased:

$$E[y_{1i}|d_i = 1] - E[y_{0i}|d_i = 0] = E[y_{1i} - y_{0i}|d_i = 1] + (E[y_{0i}|d_i = 1] - E[y_{0i}|d_i = 0]), \tag{11}$$

where the latter part between parentheses denotes the bias. This bias results from the fact that the earnings of non-veterans are not necessarily representative of what veterans would have earned if they had not served in the military. That is, people who serve in the military tend to have personal characteristics that differ, on average, from those of people who did not serve. It seems reasonable that especially very healthy and intelligent men are chosen to serve in the military and the less healthy and intelligent men are not. Since the selected men have better test scores, these men would have earned more than the non-selected even if they did not serve in the army but would have had a regular job. Therefore, we conclude positive selection has taken place. This causes a positive bias if the earnings of non-veterans who have served the army are approximated by the earnings of non-veterans who have not.

A possible solution is matching on characteristics used by the military to screen applicants. So, instead of using a regression, *matching* can be used to estimate the effect of treatment on the treated, because it can control for observed differences between veterans and non-veterans. If veteran status were randomly assigned we would have that the expected earnings for both veterans and non-veterans would be the same. However, this is not very likely. It is more likely to assume that veteran status among applicants is ignorable conditional on a set of observed variables (x_i), that is, age, schooling and test-scores for the application tests. This assumption is called *selection on observables.* According to this assumption, the earnings of both veterans and non-veterans are independent of whether or not they served in the military d_i given the observable characteristics x_i:

$$(y_1, y_0) \amalg d_i | x_i, \tag{12}$$

where $\amalg$ is the notation for statistical independence. This assumption implies that the $E[y_0|d_i, x_i] = E[y_0|x_i]$. Now the effect of being in the army (the treatment) on the veterans (the treated) can be estimated and is, by the law of iterating expectations, equal to

$$E[y_{1i} - y_{0i}|d_i = 1] = E[\{E[y_i|x_i, d_i = 1] - E[y_i|x_i, d_i = 0]\}|d_i = 1]. \tag{13}$$

Since this will result in some cells with no controls, you can try to only estimate $a_c = E[y_{1i} - y_{0i}|d_i = 1, P(d_i = 1|x_i) < 1]$. The resulting estimator is a weighted estimator,

$$\widehat{a}_c = \frac{\Sigma_k \delta_k N_{1k}[\overline{y}_{1k} - \overline{y}_{0k}]}{\Sigma_k \delta_k N_{1k}}, \tag{14}$$

where N_{1k} and N_{0k} is the number of observations in the population of veterans and non-veterans with $X = x_k$, respectively; $\delta_k = 1[n_{1k} > 0, n_{0k} > 0]$, and $\overline{y}_{ik}$ denotes the average earnings of (non-)veterans with $X = x_k$. This estimate corrects for the population distribution function of X among veterans, that is, $\frac{\delta_k N_{1k}}{\Sigma_k \delta_k N_{1k}}$, for which we can define the average earnings differential. The weights are proportional to the probability of veteran status. So the men more likely to serve get the most weight in estimates of the effect of treatment on the treated. Therefore, the effect of the treatment depends on the chance of having served in the army.[7]

Instead of matching, also a regression can be used to estimate the effect of treatment on the treated. A regression produces also a weighted estimator, but it weights each of the underlying treatment effects by the conditional variance of treatment status, which in this case is maximized when $P(d_i = 1|x_i = x) = 1/2$. The empirical consequences of the OLS weighting scheme depend on the distribution of regressors and the amount of heterogeneity in the causal effect of interest. Matching methods provide an alternative estimation strategy that has more control over the weighting scheme used to produce average causal effects. The matching parameter weights each of the underlying treatment

[7]For a derivation of the effect see Chapter 23 by Angrist in the *Handbook of Labor Economics.*

effects by $P[d_i = 1|x_i]P(x_i)$, whereas the regression parameter weights each of the underlying treatment effects by $P[d_i = 1|x_i](1 - P[d_i = 1|x_i])P(x_i)$. Matching is most practical in cases where the causing variable takes on two values. The regression estimator can be compared to the matching estimator to determine the presence of a selection bias and the magnitude of the bias. In most cases the selection on observables assumption will appear to be a key assumption in estimating the effects of a treatment correctly.

5 Instrumental Variables

5.1 Constant-Effects Models

Sometimes when using empirical models, the regression that results is not the regression economists want. For example, they are interested in the regression of wages y_i on schooling s_i and ability a_i, so they want to estimate

$$y_i = x_i'\pi + \varepsilon_i = \alpha + \beta_l s_i + \gamma a_i + \varepsilon_i. \tag{15}$$

For simplicity, define (15) as the "long-regression." However, ability cannot be observed, but it may be correlated with schooling, so regression gives

$$y_i = \alpha + \beta_s s_i + \eta_i, \tag{16}$$

where $\eta_i = \gamma a_i + \varepsilon_i$, so η_i is correlated with s_i if s_i and a_i are correlated. Denote this as the "short-regression." So, when running a regression, only information about the short-regression β_s will appear. Note that this short-regression estimator is composed out of the long-regression estimator, β_l, and a bias term that results out of the correlation between a_i an s_i. Since economists are interested in the long-regression β_l, instead one can use the instrumental variables (IV) approach to consistently estimate it, without actually having data on the control variable a_i. Instrumental variables solve the problem of selection bias by using field variation to approximate a randomized trial.

The IV approach can be illustrated in the example of veterans' earnings (Angrist, 1998). Note that the key is the link with potential outcomes. First, consider the constant-effects model, where $y_{1i} - y_{0i}$ is assumed to be constant for everybody, $y_{1i} - y_{0i} = \beta$. As we know from the previous section, if you want to compare the mean earnings by treatment status (that is, veterans and non-veterans), a bias term arises because the earnings of veterans do not necessarily tell us what veterans would have earned if they had not served.

This is the causal analog of omitted-variables bias in the "short-regression–long-regression problem." IV methods can eliminate this sort of bias if the researcher has access to an instrumental variable z_i that is correlated with the dummy variable indicating whether one is a veteran or not, d_i, but is independent of the mean earnings of non-veterans, y_{0i}. Therefore, y_i and z_i are correlated solely because of the change in d_i induced by z_i. This is similar to a randomized trial, where outcomes differ by treatment status solely because of the effect of the treatment. Examples of instrumental variables can be found, for example, in Card (1995) who uses college proximity as an IV for education, since it is more likely to have a college degree if college was nearby; also the level of education of the parents can be used as an instrument for education level.

Now β_{IV} can be computed as follows:

$$\beta_{IV} = \frac{Cov(y_i, z_i)}{Cov(z_i, d_i)}. \tag{17}$$

When the instrument z_i is binary with value 1 with probability p, the covariances in (17) can be formulated as follows:

$$Cov(y_i, z_i) = (E[y_i|z_i = 1] - E[y_i|z_i = 0])p = (E[y_i|z_i = 1] - E[y_i|z_i = 0])(1 - p). \tag{18}$$

$$Cov(z_i, d_i) = (E[d_i|z_i = 1] - E[dy_i|z_i = 0])p = (E[d_i|z_i = 1] - E[d_i|z_i = 0])(1 - p). \tag{19}$$

As a result the "ratio of difference in means" formula for the population expression β_{IV} becomes

$$\beta_{IV} = \frac{E[y_i|z_i = 1] - E[y_i|z_i = 0]}{E[d_i|z_i = 1] - E[d_i|z_i = 0]}. \tag{20}$$

By estimating $E[y_i|z_i] = \alpha + \beta E[d_i|z_i]$ the Wald (or IV) estimator is obtained, which is the sample analog of the population expression (20).

What does this tell you about the way in which the IV approach solves the selection bias in the example of the (non-)veterans? A natural instrument is draft-eligibility status, since this is determined by a lottery over birthdays, before young men apply for serving in the military. This lottery indicates whether an individual was draft-eligible ($z_i = 1$) or draft-ineligible ($z_i = 0$).[8] Note that this instrument is correlated with whether an individual served the army or not (d_i) but is uncorrelated with the earnings (y_i). In this case, the IV estimator is simply the difference in mean earnings between draft-eligible and ineligible men, divided by the difference in the probability of serving in the military

[8] For details, see Angrist (1990).

between draft-eligible and ineligible men. If you use this instrument in estimating (20), the selection bias is no longer present.[9]

Now suppose the instrument is not binary, but takes on more (say J) values. Then for each value of the instrument you should have

$$E[y_i|z_i = j] = \alpha + \beta P[d_i = 1|z_i = j], \;\; j = 1, ..., J, \tag{21}$$

where $P[d_i = 1|z_i = j]$ is the probability that one served the army for a given value of the instrument z_i. An estimator for β is then obtained by regressing

$$\overline{y}_j = \alpha + \beta \widehat{p}_j + \overline{\varepsilon}_j, \tag{22}$$

where a bar denotes the mean and $\widehat{p}_j$ is a predicted value of serving the army. It can easily been checked that the Wald estimator is just a special case of this when $J = 2$. The efficient estimation approach for (22) is generalized least squares, that is weighted least squares with weights given by $[V(\overline{\varepsilon}_j)]^{-1/2}$, where $V(\overline{\varepsilon}_j) \equiv \frac{\sigma_j^2}{n_j}$ and $\sigma_j^2 = E[\overline{\varepsilon}_j|z_i = j]$ and n_j is the size of the group with $z_i = j$.

5.2 IV Estimation with Heterogeneous Potential Outcomes

The previous section involves a highly stylized description of the world, where causal effects are the same for everyone and the effects are linear. Although these assumptions may provide a useful starting place because they focus attention on basic causality issues, there is no reason to believe they are true in reality. The cost of these simplifying assumptions is that they gloss over the fact that even when a set of estimates has a causal interpretation, they are generated by variation for a particular group of individuals over a limited range of variation in the causing variable (Angrist and Krueger, 1999). Instrumental variables methods exploit identifying information that manipulates treatment for a subpopulation as if by random assignment. Sometimes, the IV approach *is* a random assignment. The modern approach to IV is distinguished by the explicit attention given to the integrity of the experimental design, even if the experiment is not one of "those we should like to make" (Haavelmo, 1944), but rather part of "the stream of experiments that Nature is steadily turning out from her own enormous laboratory, and which we merely watch as passive observers." According to Haavelmo, it makes sense to ask what

[9] One important pitfall of the IV approach is the validity of instruments, that is, the possibility that η_i and z_i are correlated. This should be checked very carefully.

experiment we would *like* to make, and to compare this with the experiment at hand. A good instrument identifies an *internally valid* causal effect. That is, it captures causal effects for a group subject to (quasi-)experimental manipulation. So, if you would like to check the internal validity, you should check whether an empirical relationship has a causal interpretation in the setting where it is observed. Issues of *external validity*, that is, whether a set of internally valid estimates has predictive value for groups or values of the response variable other than those observed in a given study, are not purely empirical. These questions are typically resolved by a combination of indirect evidence and theoretical reasoning. Constant-effect models make it harder to discuss the two types of validity separately, since external validity is automatic in a constant-effects setting. For example in a constant-effects model the effects of military service on earnings is the same for high-school dropouts and college graduates. In this section the interpretation of the IV approach is discussed when the assumption of constant-effects are relaxed.

5.2.1 The LATE Framework

Again the example of (non-)veterans earnings is used, where d_i indicates whether or not individual i has had the treatment. The average treatment effect (ATE), $E[y_{1i} - y_{0i}]$, and the average effect on the treated (ATE$_1$), $E[y_{1i} - y_{0i}|d_i = 1]$, were already discussed. Imbens and Angrist (1994) have defined another treatment effect, which they call local average treatment effects (LATE). LATE has the advantage that it can be estimated using instrumental variables under very weak conditions. However, there are also two potential drawbacks. First, it measures the effect of treatment on a generally unidentifiable subpopulation. Another drawback is that the definition of LATE depends on the particular instrumental variable that we have available. To define LATE, we need to have an instrumental variable z_i. For simplicity we take a binary instrument and we define random variables d_{1i} and d_{0i} where $d_{ji} = d_i$ when $z_i = j$. Hence, d_i can be redefined as

$$d_i \equiv d_{0i} + [d_{1i} - d_{0i}]z_i. \tag{23}$$

Then the local average treatment effects (LATE) are defined as

$$E[y_{1i} - y_{0i}|d_{1i} \neq d_{0i}]. \tag{24}$$

A key assumption is that the vector of potential outcomes and potential treatment assignments is jointly independent of the instrument. That is, all expectations involving functions of $(y_{1i}, y_{0i}, d_{1i}, d_{0i})$, conditional on z_i, do not depend on z_i.

Condition 1 *(Independence and Exclusion): The random variables y_{1i}, y_{0i}, d_{1i}, and d_{0i} are jointly independent of z_i.*

Condition 1 implies that z_i is ignorable or "as good as randomly assigned." Two things become clear from this condition. First, the causal effect of z is identified, that is, $y_i(d_z, z)$ is independent of z. Second, potential outcomes are unchanged in an experiment manipulating z if d_z is unchanged, that is, z satisfies an exclusion restriction.

Using this condition, LATE can be identified by

$$\begin{aligned} E[(y_{1i} - y_{0i})d_{1i}] - E[(y_{1i} - y_{0i})d_{0i}] = \\ E[y_{1i} - y_{0i}|d_{1i} - d_{0i} = 1] \cdot P[d_{1i} - d_{0i} = 1] \\ -E[y_{1i} - y_{0i}|d_{1i} - d_{0i} = -1] \cdot P[d_{1i} - d_{0i} = -1]. \end{aligned} \tag{25}$$

Thus, even if $y_{1i} - y_{0i}$ is positive for everyone, the average difference in outcomes by z_i can be zero or even negative. However, Condition 1 is not sufficient for identifying LATE or any other average causal effect to be identified. In identifying LATE monotonicity is needed.

Condition 2 *(Monotonicity): Either $d_{1i} \geq d_{0i}$ for all i or $d_{1i} \geq d_{0i}$ for all i.*

The latter condition restricts the manner in which the instrument can influence treatment assignment. While the instrument may have no effect on some people, it can push individuals in only one direction. In what follows, we assume $d_{1i} \geq d_{0i}$ for all i. The interpretation of this assumption is relatively simple when z_i is a dummy variable representing assignment to the treatment group: anyone in the population who would be in the treatment group in the absence of assignment (or eligibility) would be in the treatment group if assigned to the treatment group. In case of our example, this simply means that people who would serve the military without being eligible would also serve if they were eligible. Without monotonicity, instrumental variables estimators are not guaranteed to estimate a weighted average of the underlying causal effects, $y_{1i} - y_{0i}$ (Angrist and Krueger, 1999). Both conditions together are sufficient to identify LATE:

$$E[y_{1i} - y_{0i}|d_{1i} > d_{0i}] = \frac{E[y_i|z_i = 1] - E[y_i|z_i = 0]}{E[d_i|z_i = 1] - E[d_i|z_i = 0]}. \tag{26}$$

This is just Wald's (1940) formula for estimating the slope of a line. LATE is the average treatment effect for those who would be induced to serve the military by changing z_i from zero to one.

There are two things that distinguish LATE from other treatment parameters. First, it depends on the instrument z_i. If we use different instruments, LATE generally changes. So, it can be seen that the IV parameters are tied to the instrument. The parameters ATE and ATE_1 are defined without reference to an IV, but only with reference to a population. A second distinguishing aspect is that, because we cannot observe both d_1 and d_0, we cannot identify the subpopulation with $d_1 - d_0 = 1$. By contrast, ATE averages over the entire population, while ATE_1 is the average for those who have actually had the treatment. To conclude, it can be said that the usual IV estimator consistently estimates LATE under weak assumptions.

5.2.2 Heterogeneous Potential Outcomes

Finally, it can also be the case that the treatment of interest, d_i, is not binary, but varies in intensity. This is called a multinomial treatment. In a more general setting with heterogeneous potential outcomes, different instruments estimate different weighted averages of the difference $y_{1i} - y_{0i}$. Examples of heterogeneous potential outcomes can be found in Angrist and Krueger (1991), who use the quarter of birth as an instrument for the number of years of schooling in estimating the wage, and in McClellan *et al.* (1994), where the distance to a cardiac care center relative to the nearest hospital is used as an instrument for whether or not an invasive treatment for elderly heart-attack victims is provided in estimating the chance of survival. Also in this case, the independence and monotonicity assumptions are needed to identify the average causal response.[10]

6 Conclusion

The paper attempts to provide an overview of the empirical strategies most commonly used in labor economics. The discussion of the topics of causal regression, matching, and instrumental variables has highlighted what has to be taken into account in estimating consistently the effects of treatments on a variable of interest. The first step is to define the conditional expectation function (CEF), which is the summary of the relationship between y_i and x_i. Then economists can think of a causal question which makes it

[10]For more details see Angrist and Krueger (1999), Section 2.3.4

interesting to run a regression. If a regression provides an answer to this "what if" question, then inferences on average treatment effects can be done. It is important to check whether the treatment effects of interest can be identified consistently when you exploit assumptions concerning ignorability of the treatment conditional on a set of covariates. If not, one should use an estimator that relies on the availability of one or more instrumental variables that are redundant in the response equations but help determine participation in the program. In this context it is important to keep in mind how the causal effect is identified. In particular, what causes the explanatory variable of interest to vary when other variables are held constant? Also, remember which groups of individuals (with and without treatment) are being compared. Always check the identification strategy by testing it in a situation in which the causal variable is not expected to have an effect.

References

Angrist, J.D. (1990), "Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence From Social Security Administrative Records," *American Economic Review*, 80: 313-335.

——— (1998), "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants," *Econometrica*, 66, 2: 249-288.

——— (2004), Lecture Notes NAKE Workshop, December 2004, Maastricht.

——— and A.B. Krueger (1991), "Does Compulsory School Attendance Affect Schooling and Earnings?," *Quarterly Journal of Economics*, 106: 979-1014.

——— and A.B. Krueger (1999), "Empirical Strategies in Labor Economics," *Handbook of Labor Economics*, Vol. 3, Chapter 23.

Barnow, B.S., G.G. Cain and A. Goldberger (1981), "Selection on Observables," *Evaluation Studies Review Annual*, 5: 43-59.

Haavelmo, T. (1944), "The Probability Approach in Econometrics," *Econometrica*, 12: 1-118, Supplement.

Imbens, G.W. and J.D. Angrist (1994), "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62: 467-476.

Wald, A. (1940), "The Fitting of Straight Lines if Both Variables are Subject to Error," *Analysis of Mathematical Statistics*, 11: 284-300.

Wooldridge, J.M. (2002), *Econometric Analysis of Cross Section and Panel Data*, MIT Press, Cambridge, Massachusetts.

Schumpeterian Growth: Theory and Applications

Report by Edwin van der Werf*

1 Introduction

Schumpeterian growth theory explains long-run economic growth from microeconomic principles. A country's institutional setting and industry structure determine the profits for an innovator and hence the rate of innovation and the country's real GDP growth rate. In a series of five lectures, Philippe Aghion (Harvard University) provided an overview of both the basic Schumpeterian growth theory and its empirical relevance. This report gives an overview of his lectures.[1]

The report is organized as follows. Section 2 presents an underpinning for the need for a 'new' growth theory and sets out the basic Schumpetarian growth model. Section 3 studies convergence in growth rates. Sections 4-6 present the basic insights of a number of papers that take the theory to the data.

2 A Stripped-Down Schumpeterian Growth Model

A major problem of both the Solow-Swan and the Ramsey-Cass-Koopmans growth models is that they fail to explain long-run growth. The models generally exist of a constant returns to scale production function with labor and capital and have a long-run growth rate equal to the exogenous rate of technical change. An increase in the savings rate or a decrease in the discount rate temporarily increases capital accumulation and hence the growth rate of per capita consumption. However, as there are decreasing returns to capital, the economy will converge to a new steady state with a higher level of (per capita) output and capital stock, but with the growth rate still determined by the exogenous

*Department of Economics and CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg. Phone: (+31) 13 466 8730. E-mail: e.h.vanderwerf@uvt.nl.

[1]To a large extent Professor Aghion's lectures followed Aghion and Howitt (2004).

rate of technical progress. As a consequence these growth models do not explain long-run growth but rather assume it. The problem is that the factor determining the level of technology (entrepreneurs, innovators) cannot be paid from output as the model is homogeneous of degree one in capital and labor. As a consequence there is no money left to reward those who cause technical progress and hence no one will invest.

In the late 1980s and early 1990s several models emerged that tried to explain long-run economic growth while innovators would be rewarded (see, for example, Romer (1990), Aghion, and Howitt (1992), Grossman and Helpman (1991)). The papers by Aghion and Howitt are based on the insights by Joseph Schumpeter. Schumpeter claimed that innovation was driven by profit-seeking entrepreneurs. According to Professor Aghion, the Schumpeterian approach to growth theory has the following characteristics:

- productivity growth is driven by innovations;
- innovations result from deliberate investment in research and development (R&D);
- R&D investments are motivated by the prospect of monopoly rents in case of successful innovations; and
- new innovations lead to better products or technologies, which make old products or technologies obsolete (creative destruction).

Where neoclassical growth models could not reward innovators because of constant returns to scale—and thus sustaining a perfectly competitive market structure—new growth models typically contain monopoly power for some producers. The accompanying monopoly rents are used to repay the up-front R&D expenditures.

The Aghion and Howitt (1992) model can be stripped down to the following toy version. Time is discrete and denoted by t. At each point in (discrete) time there is a mass L of individuals who live for one period only. Each individual is endowed with one unit of skilled labor that she supplies inelastically. Each period final output is produced using a Cobb-Douglas technology:

$$y_t = A_t x_t^{\alpha}, \tag{1}$$

where x is the quantity of the intermediate input used in final good production and A is a productivity parameter reflecting the current quality of the intermediate good. Each unit of the intermediate good is produced using one unit of labor so x also denotes the amount of labor in manufacturing. Labor market clearing requires:

$$L = x_t + n_t, \tag{2}$$

where n is the amount of labor employed in research, which generates innovations. Each innovation improves the quality or productivity of the intermediate input by a factor $\gamma > 1$. Hence if there is a successful innovation at time t productivity at this point in time is $A_t = \gamma A_{t-1}$. If a researcher or entrepreneur devotes her time to R&D she innovates, that is discovers an improved version of the intermediate input, with probability λ. Since this probability is constant, the model has constant returns in both production and R&D but increasing returns at the aggregate level as quality A can grow.

By devoting her time to innovation the innovator chooses to forego income from labor in manufacturing. She will only choose to do so if the expected returns to innovation (note that R&D is an up-front investment) are at least as high as the wage rate. When an innovator is successful she obtains monopoly power in the production of the (improved quality) intermediate good. If she were an unconstrained monopolist she would charge the price w_t/α. However, Aghion and Howitt (2004) restrict the monopoly power by assuming a binding competitive fringe that is able to produce the same intermediate but using $\chi > 1$ units of labor. When $\chi < 1/\alpha$ the fringe is binding and the monopolist can charge at most χw_t without being driven out of the market. Hence monopoly profits equal

$$\pi_t = (\chi - 1)\, w_t x_t. \tag{3}$$

Monopoly profits only last one period, as do individuals.

As noted above an individual will only engage in R&D if she expects to earn at least the wage she would get in manufacturing. Since there is free entry in R&D both ways of earning an income will yield exactly the same return:

$$w_t = \lambda \gamma \pi_t. \tag{4}$$

This research arbitrage equation says that the income in manufacturing w_t, or marginal cost of innovation, should equal the expected return to innovation. With probability λ an innovator will get monopoly profits π_t while an innovation multiplies wages and profits by γ.

In steady state the allocation of labor and manufacturing is constant. If we then substitute profits (3) in the research arbitrage equation (4) we can solve for the steady state amount of labor in research:

$$n = L - \frac{1}{\lambda\gamma\left(\chi - 1\right)}. \tag{5}$$

The growth rate of the economy then equals the amount of labor in research times the probability that a researcher innovates times the contribution of an innovation:

$$g = n\lambda\left(\gamma - 1\right) = \left(\lambda L\left(\gamma - 1\right)\right) - \frac{\gamma - 1}{\gamma}\frac{1}{\chi - 1}. \tag{6}$$

The growth rate hence depends on the parameters λ, γ, χ and L. It increases with the supply of skilled labor L and with the productivity of innovation λ, which both point to the importance of education. Furthermore the growth rate increases with the size of innovations γ. The parameter χ represents the inverse of the degree of product market competition and the toy model described in this section has the growth rate decreasing in competition.

In the following sections we will present results from research by Aghion and several collaborators. Starting point is the toy version of the Schumpeterian growth model presented in this section and the following will show the versatility and empirical relevance of the Schumpeterian approach.

3 Convergence of Growth Rates

One of the main advantages of the Solow-Swan model was its ability to explain the (conditional) convergence of countries (Barro and Sala-i-Martin (1995)). The model predicts that if two countries have the same underlying structure but one is less developed in the sense that it has a smaller capital stock, the 'backward' country has an advantage due to decreasing returns to capital accumulation. However, recent papers (see, for example, Quah (1997)) have shown that there appears to be 'club-convergence:' all rich and most middle-income countries belong to the one group of countries with the same long-run growth rate. All other countries seem to have diverse growth rates, all strictly less than that of the convergence club.

The Schumpeterian growth paradigm stresses the role of institutions in economic growth. The paper by Aghion, Howitt, and Mayer-Foulkes (2005) studies the role of financial development on growth and convergence. The underlying theoretical model starts from a multi-country version of the toy model of Section 2. Important new elements are:

(i) If an innovator in a country is successful, its technology level jumps ahead of all other countries, even when it was far behind (that is, there are international knowledge spillovers). However, the cost of innovation increases with the current leading edge technology.

(ii) Instead of introducing a new leading edge technology, a country can try to implement an existing technology. However, this is a costly process, much in the same way as innovation.

(iii) There is an agency problem in financing innovation. An innovator has to borrow the difference between the cost of innovation and her wage income, while the latter depends on the current level of technology in the country. An innovator can defraud her creditors by hiding the results of a successful innovation and hence not having to repay. However, the cost of defraud depends positively on the level of financial development.

In this model there is an advantage of backwardness as long as innovators are credit-unconstrained, but then there is only convergence and not club-convergence. However when there are binding credit constraints there is a disadvantage of backwardness. Lenders will take into account that innovators my defraud. As a consequence incentive-compatibility puts an upper limit on the innovator's investment. This credit limit will be binding if a country is too far behind the frontier technology, as the cost of innovation increases with the leading edge technology while the wage depends on the country's own level of technology. As a consequence in financially developed countries, with a high degree of creditor protection, it is easier to get credit and hence to innovate. More precisely, the effect of financial development on steady-state growth and per capita GDP should be positive up to some critical level and zero thereafter.

Aghion *et al.* (2005) test this theory using a cross country growth regression where a variable representing a country's financial development (using the measure introduced by Levine *et al.* (2000)) interacts with the country's distance to the technology frontier (the difference between the country's income level and US income). They show that this interaction coefficient is indeed significantly negative, while the direct effect of financial intermediation is not significantly different from zero. Hence financial development is an important institution for club-convergence: it explains why some (developed) countries converge while others diverge.

4 Competition and Economic Growth

As noted in Section 2, one feature of the stripped down version of the Aghion and Howitt (1992) model is that it predicts product market competition to be unambiguously bad for growth. More product market competition, a lower value of χ, implies lower monopoly rents to innovators and hence less incentives to innovate. However, empirical work such as Geroski (1995), Nickell (1996), and Blundell *et al.* (1999) has pointed to a positive relation between product market competition and innovation. In reaction to this, Aghion, Harris and Vickers (1997) and Aghion, Harris, Howitt and Vickers (2001) extended the basic Schumpeterian growth model by allowing incumbent firms to innovate. In these models innovation incentives depend more upon the difference between post- and pre-innovation rents than on post-innovation rents per se. As product market competition (PMC) may reduce a firm's pre-innovation rents more than its post-innovation rents, PMC may foster innovations and growth.

Aghion, Bloom, Blundell, Griffith, and Howitt (2002) develop a Schumpeterian growth model that predicts an inverted U relationship between product market competition and innovation. In a multi-sector model, each sector contains of two firms that may differ in unit costs. The output of the two firms is combined to industry output using a CES technology, where the substitution parameter is a measure of product market competition. When a firm innovates it reduces its production costs. Innovation is done step by step, so lagging firms must first catch up with the leader before they can become the leader themselves. In industries with intense competition, so-called neck-and-neck industries, competition may increase the incremental profits from innovating. Hence firms can 'escape competition' through innovation. However in industries where firms are less leveled, competition may also reduce innovation as the laggard's reward to catching up with the technological leader may fall (Schumpeterian effect). As a consequence an increase in competition tends to reduce the number of neck-and-neck industries. Therefore the escape-competition effect tends to dominate for low initial levels of competition, while the Schumpeterian effect tends to dominate at higher levels of competition and the model predicts an inverted U relationship between PMC and innovation.

Aghion *et al.* (2002) confront the model with data from a panel of UK firms. Using a semiparametric approach they show how an exponential quadratic model fits the data extremely well. In addition the result is robust to alternative specifications. This indicates that some kind of escape-competition effect should dominate at lower levels of PMC, whereas the Schumpeterian effect as known in the industrial organization literature and described in Section 2, should dominate at high initial levels of product market competition.

5 Volatility and Economic Growth

Only recently there have been attempts, both theoretical and empirical, to analyze the relation between volatility and growth in depth. Ramey and Ramey (1995) find a negative effect of volatility on growth, using cross-section data from 92 countries. However, when focussing on the sub-group of OECD countries they find no effect. A straightforward explanation for a negative relation between volatility and growth is through investment and risk aversion. More volatility discourages investment and hence growth as it implies more risk. At the same time volatility implies income risk, inducing people to save more through the precautionary principle, which in turn reduces the interest rate and encourages investment. As a consequence the total effect is unclear a priori, which is also shown by the data in Aghion, Angeletos, Benerjee, and Monova (2004) (henceforth AABM): including the level of investment to explain the effect of volatility on growth has only a minor effect.

Instead of focussing on the level of investment, Aghion *et al.* (2004) propose that the main channel through which volatility affects growth is the composition of investment. They develop a Schumpeterian growth model where entrepreneurs may be credit-constrained and face aggregate uncertainty. When an entrepreneur is successful in generating an innovation, she has to incur a positive cost to implement the innovation. Whether she will be able to do so depends on the efficiency of credit markets. From the theoretical model the authors derive several predictions. First, the level of financial development affects the relationship between aggregate volatility and mean growth. Tighter credit constraints have a negative effect on this relationship. Second, when firms face tighter credit constraints, growth becomes more sensitive to (lagged) exogenous shocks. Third, R&D as a fraction of GDP or total investment tends to be countercyclical in the absence of credit constraints, but becomes increasingly procyclical as credit constraints tighten.

Aghion *et al.* (2004) first look at the long-run response of growth to volatility in a cross-section of 70 countries over the period 1960-1995. The model they estimate explains growth from the initial level of income, a vector of country specific controls, volatility, the level of financial development (using the measure by Levine *et al.* (2000)) and an interaction term of the latter two variables. They find that the mean annual growth rate is less sensitive to the standard deviation of growth rates the higher the degree of financial development.

For the second prediction the authors estimate a model, using panel data, with (lagged) shocks, the level of financial development and interaction terms between these amongst the explanatory variables. They find a strong response of growth to terms-of-trade or commodity price shocks lagged one and two years. Besides, the coefficient on the interaction between lagged private credit and the shock is strong and significant.

Finally, the authors look at the response of R&D to shocks. Using annual data on 14 OECD countries between 1973 and 1997, Aghion *et al.* (2004) estimate a model that explains the share of R&D in investment from, among other variables, (lagged) shocks, the level of financial development and interaction terms between these. For higher levels of private credit they find that R&D is countercyclical with respect to (twice-lagged) shocks, confirming their prediction.

6 Education, Distance to Frontier, and Growth

As is empirically well-known for many years, and theoretically supported by the endogenous growth literature, an economy's growth rate is positively affected by (the share of) skilled labor in the economy. However, Krueger and Lindahl (2001) found that "education [is] statistically and positively associated with subsequent growth only for the countries with the lowest level of education." Vandenbussche, Aghion, and Meghir (2004) explain this puzzling result from the fact that the source of technological progress is the result not only of the adoption of existing technologies but also of pure innovation, especially in technologically advanced countries (see Section 3 and Aghion *et al.* (2005)). Since the latter requires a different type of human capital, one should look at both an economy's distance to the technological frontier and the composition of human capital.

In the theoretical part of Vandenbussche *et al.* (2004), technological improvements are the results of a combination of innovation and imitation. Each of these activities in turn requires a combination of high skilled and low skilled labor, where the authors assume that innovation makes relatively more intensive use of skilled labor. Consequently, holding the composition of human capital constant, an increase in its aggregate level is always growth enhancing. Holding the level of human capital constant, however, the growth-enhancing properties of human capital depend on both its composition and the distance to the technological frontier: the growth-enhancing impact of skilled labor decreases with a country's distance to the technological frontier. In advanced economies, imitation contributes only little to growth so unskilled labor contributes little to technological change.

The theoretical model predicts that the growth-enhancing effect of a marginal increase in the stock of skilled labor is stronger the closer the economy is to the technological frontier. Vandenbussche *et al.* (2004) test this prediction using a panel data set covering 19 OECD countries between 1960 and 2000. They estimate several models where growth in total factor productivity is explained by, amongst other variables, the distance to frontier (the ratio of the country's and the US' total factor productivity level), variables representing education and interaction terms between the distance to frontier and the education variables. The results support the theoretical claim that holding unskilled human capital constant, skilled human capital has a higher growth-enhancing effect closer to the technological frontier.

7 Concluding Remarks

Professor's Aghion lectures introduced us into both the theoretical underpinnings of the Schumpeterian growth paradigm and in its empirical relevance. The Schumpeterian growth paradigm starts from the simple insight that innovations come from profit-seeking innovators. As has been shown by the several papers presented in this report, the basic Schumpeterian model presented in Section 2 can be extended in several ways to answer several empirically relevant questions.

Several of the papers presented here show the importance of institutions in explaining a country's rate of development. The institutions studied are education, the degree of product market competition, and the degree of financial development. It would be interesting to see extensions to other institutions like labor market flexibility. In addition it should be noted that in some of the papers links between the institutions studied and the Schumpeterian growth paradigm is rather weak: the regressions presented could come from any model stressing the importance of institutions. Finally it would be interesting to look at open economy models and how the theoretical implications from these would affect the results found with closed economy models.

Notwithstanding these remarks, endogenous growth theory in general, and especially the Schumpeterian growth paradigm, offers important insights into economic growth. The overview presented in this report shows that these theories are not just of theoretical interest, but are able to explain important real life economic facts.

References

Aghion, P., G.-M. Angeletos, A. Banerjee and K. Manova (2004), "Volatility and growth: financial development and the cyclical composition of investment," mimeo available at http://post.economics.harvard.edu/faculty/aghion/papers/Volatility_and_Growth.pdf

Aghion, P., N. Bloom, R. Blundell, R. Griffith and P. Howitt (2002), "Competition and innovation: an inverted U relationship," NBER Working Paper No. 9269.

Agion, P., C. Harris and J. Vickers (1997), "Competition and growth with step-by-step innovation: an example," *European Economic Review, Papers and Proceedings*, pp. 771-782.

Agion, P., C. Harris, P. Howitt and J. Vickers (2001), "Competition, imitation and growth with step-by-step innovation," *Review of Economic Studies*, 68, pp. 467-492.

Aghion, P. and P. Howitt (1992), "A model of growth through creative destruction," *Econometrica*, 60, pp. 323-351.

Aghion, P. and P. Howitt (2004), "Growth with quality-improving innovations: An integrated framework," draft of Chapter 2 in the forthcoming *Handbook of Economic Growth*, P. Aghion and S. Durlauf (eds.).

Aghion. P., P. Howitt and D. Mayer-Folkes (2005), "The effect of financial development on convergence: Theory and evidence," *The Quarterly Journal of Economics*, 120, forthcoming.

Barro, R.J. and X. Sala-i-Martin (1995), *Economic Growth*, McGraw-Hill.

Blundell, R., R. Griffith and J. van Reenen (1999), "Market share, market value, and innovation in a panel of British manufacturing firms," *Review of Economic Studies*, 66, pp. 529-554.

Geroski, P. (1995), *Market Structure, Corporate Performance, and Innovative Activity*, Oxford University Press.

Grosman, G.M. and E. Helpman (1991), *Innovation and growth in the global economy*, Cambridge.

Krueger, A.B. and M. Lindahl (2001), "Education for growth: Why and for whom?," *Journal of Economic Literature*, 39, 1101-1136.

Levine, R., N. Loyaza and T. Beck (2000), "Financial intermediation and growth: Causality and causes," *Journal of Monetary Economics*, 46, pp. 31-77.

Nickell, S. (1996), "Competition and corporate performance," *Journal of Political Economy*, 104, pp. 724-746.

Quah, D. (1997), "Empirics for growth and distribution: Stratification, polarization and convergence clubs," *Journal of Economic Growth*, 2, pp. 27-59.

Ramey, G. and V. Ramey (1995), "Cross-country evidence on the link between volatility and growth," *American Economic Review*, 85, 1138-1151.

Romer, P. (1990), "Endogenous technical change," *Journal of Political Economy*, 98, pp. 71-102.

Vandenbussche, J., P. Aghion and C. Meghir (2004), "Growth, distance to frontier and composition of human capital," IFS Working Papers W04/31.

The Crowding Out of Intrinsic Motivation

Report by Fleur Wouterse*

1 Introduction

According to standard economic reasoning a monotonic and increasing relationship exists between monetary compensation for an activity and the performance level of that activity. This reasoning is based on some basic assumptions according to which performance is positively related to effort, which is unpleasant and needs to be compensated by a monetary reward. It is however, conceivable that factors other than money and effort enter into the decision-making process of the agent (Gneezy and Rustichini, 2000).

One factor that may enter the decision-making process is that of intrinsic motivation. Intrinsic motivation implies that an activity has a motivation of its own, independent of any reward. The offering of a reward (monetary or other) that is different from the intrinsic motivation may undermine this motivation. The result may be that the overall motivation is reduced as well as the activity itself. This effect is termed the motivation crowding-out effect. The existence of the motivation crowding-out effect would reverse one of the fundamental laws in economics namely that raising monetary incentives increase supply (Frey, 2000).

In the report the theory that external intervention undermines intrinsic motivation is tested using a gift-exchange game. In Section 2 the two main concepts—that is, crowding-out theory and intrinsic motivation—are discussed and illustrated using principal-agent theory. Section 3 gives experimental evidence for the undermining effect of external intervention. Section 4 formulates the main findings.

2 Crowding-Out Theory and Intrinsic Motivation

Crowding-out theory implies that all interventions emanating from outside the person may negatively affect intrinsic motivation. The impact of extrinsic interventions upon behavior can be demonstrated in the context of a principal-agent relationship. In this relationship, the principal (for example, an employer) uses rewards and commands in an attempt to raise

* Development Economics Group, Wageningen University, Phone (+31) 317 48 43 60, E-mail: fleur.wouterse@wur.nl

the performance, *P*, of the agent. The agent (for example, an employee of a firm) has been given a task to perform. A representative agent bases its decision-making on his performance by considering the costs and benefits of his action. Both the benefits and costs will increase with effort: $\partial B / \partial P \equiv B_P > 0$, $\partial C / \partial P \equiv C_P > 0$. Higher performance has diminishing marginal returns, $B_{PP} < 0$, and is associated with increasing marginal cost, $C_{PP} > 0$. The benefits and costs of the action performed by the agent are also influenced by the principal's external intervention, *E*:

$$B = B(P, E);\ \ B_P > 0, B_{PP} < 0, \tag{1}$$

$$C = C(P, E);\ \ C_P > 0, C_{PP} > 0. \tag{2}$$

A rational agent would opt to perform at level *P** where net benefits *(B-C)* are maximized, yielding the following first-order condition:

$$B_P = C_P. \tag{3}$$

Differentiating this optimality condition with respect to *E* shows how *P** is affected by changes in the principal's external intervention:

$$B_{PE} + B_{PP}\frac{dP^*}{dE} = C_{PE} + C_{PP}\frac{dP^*}{dE}. \tag{4}$$

The agent's optimal performance may change in a number of different ways as a result of changes in the external intervention of the principal. According to the standard economic principal-agent theory, external intervention in terms of sanctions or rewards, would raise the performance of the agent by imposing a higher marginal cost on non-compliance or equivalently by lowering the marginal cost of performing, $C_{PE} < 0$. This effect is known as the relative price effect of external intervention. If this were the only effect that would result, there would be no crowding-out effect as the external intervention only affects the marginal cost of the performance and leaves the intrinsic motivation, if it is present at all, unchanged. External intervention in this case would raise performance, $dP^* / dE > 0$. Raised performance could also result if the external intervention would increase intrinsic motivation whilst the relative-price effect would also be present.

A contrasting case is also possible. External intervention would then undermine intrinsic motivation and negatively affect the agent's marginal benefit of performing, $B_{PE} < 0$ while there is no effect of the external intervention on the marginal cost of performing, $C_{PE} = 0$. If this were the case, an increase in the external intervention would reduce the agent's performance level, $dP^* / dE < 0$. A last case to consider is in which both effects are active. The relative price effect which imposes higher marginal costs on non-compliance or equivalently reduces the marginal cost of performing so that $C_{PE} < 0$ is active as well as the crowding-out effect which reduces the marginal benefit the agent derives from performing so that $B_{PE} < 0$. It should be clear that when both of these are active, external intervention has two opposing effects on the agent's performance. The success of the intervention from the principal's point of view then depends on the relative size of the two countervailing effects.

The interaction of the price and the crowding-out effect is depicted graphically in Figure 1.

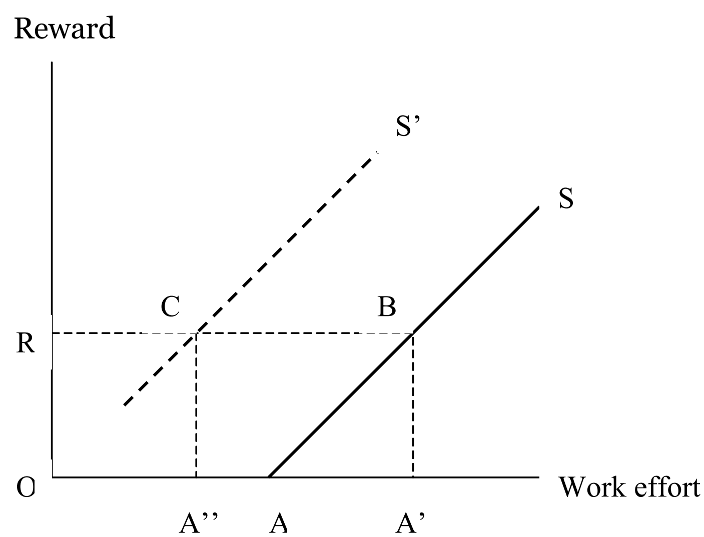

Figure 1: Net-Outcome of the Price and Crowding-Effect (Frey, 2000)

In Figure 1, *S* is the normal performance supply curve where, starting from A, work effort is a monotonic increasing function of the reward. Work effort from *O* to *A* is supplied without a reward, purely on the basis of intrinsic motivation. When the external reward for the effort is raised from *O* to *R*, effort increases from *A* to *A'*. This is the price effect. The crowding-out effect, however, would shift the supply curve from S to S' and would in fact reduce effort from A to A''.

3 Evidence From Experimental Economics

Psychologists as well as economists have investigated the evidence of the existence of a crowding-out effect on motivation. Economists have resorted amongst others to the field of experimental economics. Gächter and Fehr (2002) provide experimental evidence that incentive contracts are likely to undermine intrinsic motivation. Returning to the principal-agent approach, the agent's objective function is assumed to be increasing and strictly concave in income and decreasing in effort. If an agent were not intrinsically motivated, he would determine the optimal effort by maximizing the objective function. If the agent, however, exceeds the effort dictated by the objective function, intrinsic motivation exists. Intrinsic motivation may exist for a number of reasons. An individual may behave irrationally or may have a preference for that particular activity. It is also possible that the agent is seeking social approval of relevant reference agents or because he has certain social preferences.

In the experimental design of Gächter and Fehr (2002) two treatments are compared in the context of a "gift exchange game." The first treatment is the Trust-Treatment (TT), which establishes the extent of voluntary cooperation in the absence of performance incentives for cooperation. The second treatment is the Incentive-Treatment (IT), which introduces performance incentives for contractual compliance. A third treatment termed the Bonus-Treatment (BT) was also included.

3.1 Trust-Treatment Versus Incentive-Treatment

A "gift exchange game" is played consisting of three stages. During the first stage the buyers make a contract offer consisting of a fixed price, p, and a desired quality level, $\hat{q}$. In the second stage, the sellers decide whether or not they accept one of the contract offers made by the buyers. In the third stage, which is entered only when a contract is accepted, the sellers choose their quality level. Sellers are under no obligation to supply the buyers' desired quality level and may choose $q \neq \hat{q}$.

Under TT, where no sanctions or rewards are available to the buyer to ensure that the seller delivers the desired quality level, $\hat{q}$, a buyer's profit is given by:

$$\pi = \begin{cases} vq - p & \text{If the contract is accepted} \\ 0 & \text{If no contract is concluded.} \end{cases} \tag{5}$$

In (5) when the contract is accepted the pay-off to the buyer is *vq*, the return to the buyer as a function of the seller's actual quality, *q,* minus the price paid to the seller. The seller's payoff is given by:

$$u = \begin{cases} p - c(q) & \text{If the contract is accepted} \\ 0 & \text{If no contract is concluded.} \end{cases} \tag{6}$$

If the contract is accepted the seller incurs a cost of *c(q)* in delivering quality *q*, which is a disutility but receives a price *p* for his performance.

Under IT the buyer has the possibility of punishing sellers whose quality delivered is less than $\hat{q}$. The punishment, however, is only possible if the buyer can establish the non-compliance of the agent. Buyer's contracts under IT include a fine *f* that has to be paid by the seller when $q < \hat{q}$ can be verified. Verification of non-compliance can be done with probability $0 < s < 1$. A buyer's expected payoff under IT is therefore given by:

$$\pi = \begin{cases} vq - p & \text{If the contract is accepted and } q \geq \hat{q} \\ vq - p + sf & \text{If the contract is accepted and } q < \hat{q} \\ 0 & \text{If no contract is concluded.} \end{cases} \tag{7}$$

From (7) it can be seen that if the contract is accepted and the seller delivers a quality that is equal to or exceeds the desired quality level, the pay-off to the buyer is the return as a function of that quality minus the price paid to the seller. If the contract is accepted but the quality delivered by the seller turns out to be lower than the desired quality level, the buyer received a fine *f,* which is only paid out by the seller with probability *s* that shirking can be detected.

A seller's expected pay-off is given by:

$$u = \begin{cases} p - c(q) & \text{If the contract is accepted and } q \geq \hat{q} \\ p - c(q) - sf & \text{If the contract is accepted and } q < \hat{q} \\ 0 & \text{If no contract is concluded.} \end{cases} \quad (8)$$

In equation (8) the pay-off to the seller if the contract is accepted and he delivers quality that is equal to or exceeds the desired quality level, is similar to that under TT. However, if the contract is accepted but the seller does not perform at the desired level, he will to pay the fine *f* to the seller only if his non-compliance can be detected.

The game that was played amongst 182 university students had a one-shot nature. Subjects were randomly assigned their role as buyers and sellers. First, buyers had to decide privately on a contract offer determining *p* and $\hat{q}$. In the IT, after buyers had decided their contract offer, the verifiability of non-compliance of the seller was determined by rolling a dice. Sellers, who were in excess supply, could then choose one of the available contract offers. Subsequently, the sellers who had accepted a contract had to determine their actual quality level. In the IT, as mentioned previously, sellers had to actually determine this level before knowing whether their non-compliance was detectable. Payoffs were calculated at the end of the third stage. Because losses could be incurred, participants were given a certain amount of money to cover these losses. Subjects were paid in cash immediately after the experiments.

3.2 Behavior Under TT and IT

Under the TT a rational and selfish seller whose preferences are given by (6) would choose the minimal level of quality $q^{\min}$ because costs increase with higher quality levels. Higher quality levels yield no return because the price paid by the buyer is not a function of quality. The buyer who is aware of the seller's intention would offer the lowest price possible. However, if intrinsic motivation exists amongst sellers, they may deliver a higher quality level in response to generous contracts, that is, $q > q^{\min}$.

One reason for the existence of intrinsic motivation amongst subjects in the experiment may be found in subjects' social preferences. Social preferences are exhibited if the person cares about the material resources allocated to relevant reference agents in addition to material resources allocated to him or her. One particular kind of social preference is the preference for reciprocity. Reciprocity in the principal-agent type framework implies that an agent values the principal's material payoff positively or negatively. If the agent perceives the actions of the principal as kind, the agent values the principal's payoff positively whereas if the actions of the principal are perceived as hostile so is the pay-off. Inequity aversion implies that the agent does not only care about the perceived fairness or unfairness of the intentions of the principal but also about the perceived fairness of the actions' consequences. The perception by the agent of an action as kind or hostile depends on the fairness and unfairness of the consequences and the intention associated with the action, which depends on the equitability of the payoff distribution relative to the set of feasible payoff distributions.

With reference to the game, the existence of reciprocity amongst sellers implies that if the buyer makes a generous offer, the seller will perceive his action as kind and will, therefore, reciprocate and deliver $q > q^{\min}$. Inequity averse sellers will respond similarly to a generous contract offer because by delivering a higher (equalizing quality level) than the desired quality level, $q^{e} > \hat{q}$, these sellers can generate outcomes that they perceive as more equitable. At sufficiently high prices inequity-averse sellers will choose a quality level so that the payoff to the buyer in terms of profit equals their own pay-off in terms of utility.

Under IT when buyers and sellers are selfish, rational preferences of the seller are given by (8). The seller in this case would opt to deliver a quality level q^*, where the benefits of the contract are equal to or will exceed the cost of the contract, i.e.

$$p \geq c(q^*). \tag{9}$$

In (9) q^* symbolizes the level of quality delivered by the seller that, if the contract is accepted, would maximize the material payoff to the seller. Subsequently, when the contract is accepted the seller would only deliver the desired quality if the fine to be incurred exceeds the cost of compliance. This is termed the Non-Shirking Condition:

$$sf \geq c(\hat{q}). \tag{10}$$

Whether shirking occurs or not depends on the parameters specified in the game, determining the probability of detection and the size of the fine. If inequity-averse sellers exist, a payoff equalizing quality level, q^e, exists. If an inequity-averse seller faces a contract where $\hat{q} < q^e$ or the desired quality level of the buyer is below the equalizing quality level, this seller will never shirk and opt to perform at quality level $q^e > \hat{q}$. In this case, the outcome for the inequity-averse seller is similar to that under TT because the seller does not consider performing under $\hat{q}$ to be a viable option. If the inequity-averse seller faces a contract where $\hat{q} > q^e$ then he may decide to shirk. If the seller knew that he would end up having to pay the fine for shirking he would choose quality level so as to minimize payoff differences between himself and the buyer. In other words, he would set:

$$vq - p + f = p - c(q) - f. \qquad (11)$$

In Figure 2, this quality level is depicted by q^f. If the seller knew that the shirking would not be detectable he would set:

$$vq - p = p - c(q), \qquad (12)$$

and choose the effort level q^e.

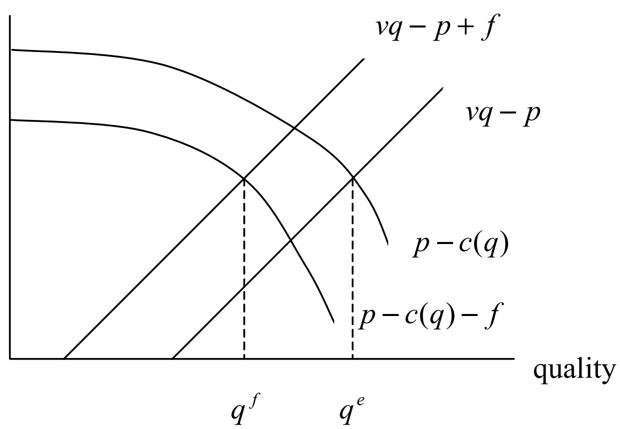

Figure 2: Inequity Aversion and Quality Choice (Fehr and Gächter, 2002)

Figure 2 depicts the cases where the seller knows that he will have to pay the fine for shirking or knows that he does not. In reality, of course, the seller does not know this and his utility depends on his own payoff, which includes the fine with probability s and excludes this fine with probability $1-s$. The non-shirking condition in this case can be written as:

$$\begin{aligned} &s\alpha[vq-p+f-(p-c(q)-f]+(1-s)\beta[p-c(q)-(vq-p)] \\ &\geq c(\hat{q})-c(q)-sf+\alpha[v\hat{q}-p-(p-c(\hat{q}))]. \end{aligned} \tag{13}$$

The left-hand side of equation (13) measures the utility loss or non-pecuniary inequality costs associated with shirking. The non-pecuniary inequality costs the seller incurs with probability s when he is detected are determined by the difference between the buyer's payoff and his own pay-off weighted by α, which is the utility weight for inequality to the seller's disadvantage. He also incurs a utility loss when his shirking is not detected, from the inequality between his payoff and the payoff of the buyer, which is inequality to his own advantage and is weighted by β. For the seller not to shirk, the utility loss incurred by shirking has to exceed the pecuniary and non-pecuniary gains from shirking as described by the right-hand side of equation (13). The pecuniary gains from shirking consists of lower costs involved in delivering lower quality but also include the possibility of having to pay the fine when detected. The non-pecuniary inequality costs associated with the proposed contract are given by the bracketed term on the right-hand side. These costs would be avoided when the seller shirks.

From equation (13) it should be clear that under the IT the fine decreases the monetary gains from shirking and as such deters shirking. But shirking is also determined by the inequality associated with the contract. A more unequal payoff between the seller and the buyer would increase the bracketed term on the right-hand side and thereby increase the gains from shirking. In addition, the desired quality level $\hat{q}$ also determines shirking by increasing the cost to the seller of compliance and by creating a greater inequality between buyer and seller.

It should be clear that in the presence of inequity-aversion the incentive contract may not necessarily produce a better result than a pure trust contract. In addition to inequity aversion, reciprocity motives should also be considered. The introduction of an incentive contract may be perceived as stemming from a hostile intention. Under the IT, a fine has been introduced which in itself may be perceived as hostile. The threat of fining expressed in this contract could be perceived by the seller as an indication of distrust. If the fine is perceived as a hostile act, sellers may no longer be willing to deliver the same quality under IT as they would have under TT.

3.3 Results

The results from the “gift-exchange game” suggest that under TT buyers offer, on average, higher prices and demand higher quality levels than under IT. In addition, it was found that quality and voluntary cooperation were lower under IT compared to TT. One of the reasons for this result to occur is that voluntary cooperation or intrinsic motivation vanished almost completely under IT. The disappearance of intrinsic motivation can be linked to the way the buyer treats the seller. Under TT, the buyer refrains from the threat of fining and this could be perceived by the seller as a fair act, which would induce his cooperation. Under the IT the usage of a fine threat could be perceived as a hostile act inducing the sellers, trustees, to reduce cooperation. A bonus treatment (BT) was also constructed, which is similar to the IT, but under the BT a seller who is caught shirking will not receive a bonus whereas a seller who does not shirk or for whom non-compliance is not detected is rewarded by this bonus. It was found that intrinsic motivation is much lower under IT compared to BT but that under BT intrinsic motivation is still lower than under TT. This finding implies that framing of the incentives is highly relevant. It is likely that framing is relevant because of reciprocity where the perception of actions as hostile or kind is important. The differences in results between the IT and the BT in terms of intrinsic motivation conveys the message to the seller that something will be taken away from him when caught shirking. Under BT the seller may have the impression that he will be rewarded if the desired quality is provided. An incentive contract including a fine may be perceived as hostile by the seller whereas a similar contract including a bonus may be perceived as kind.

4 Conclusion

The finding that incentive contracts undermine intrinsic motivation is at odds with the standard principal-agent model in which agents are assumed to be self-interested and only care about their own payoff. Self-interested agents would under the trust condition deliver the minimum effort as costs increase with effort whereas benefits do not. Under the incentive condition sanctions or rewards would, according to standard economic reasoning, have a disciplining effect by increasing the marginal cost of non-compliance or reducing the marginal cost of performing. If there would be no such thing as intrinsic motivation or voluntary cooperation, the disciplining effect of sanctions or rewards would be positive and a principal may resort to such measures to obtain the desired performance level from the agent.

However, intrinsic motivation may induce agents to supply more than the minimum effort under a trust condition and may reduce effort under an incentive condition. The influence of external interventions such as sanctions and rewards on intrinsic motivation may be attributed to an agent's social preferences, in particular reciprocity and inequity aversion. Reciprocity implies that under the trust condition, the agent perceives the deliberate non-use of the threat of fining as a particularly trusting act and reciprocates this act with his own trustworthy act, delivering the desired quality level. Under the incentive condition such trustworthiness is not expressed by the principal and can, therefore, not be reinforced by the seller. Inequity aversion influences performance depending on the perceived fairness of the consequences of actions. An inequity-averse agent would perform so as to equalize the payoff to the principal to his own payoff.

Both reciprocity and inequity aversion imply that agents do not respond to incentives as dictated by standard economic reasoning. The relative price effect of rewards and sanctions that according to this reasoning would raise performance may in fact not be active or may be counteracted by a crowding-out effect of intrinsic motivation.

References

Fehr, E. and S. Gachter (2002), "Do Incentive Contracts Undermine Voluntary Cooperation?" Institute for Empirical Research Working Paper No. 34, Zurich.

Frey, B. S. and R. Jegen (2000), "Motivation Crowding Theory: A Survey of Empirical Evidence," CESifo Working Paper No. 245, Munich.

Gneezy, U. and A. Rustichini (2000), "Pay Enough or Don't Pay at All," *Quarterly Journal of Economics,* Vol. CXV, pp. 791-810.

Upcoming in 2006…

DATE	***NAKE activity***
13 January	Start of Field courses Blocks 3 and 4, 2005 - 2006
12 – 16 June	Workshop at Erasmus University Rotterdam

Information about these activities can be obtained from Astrid van Hemert, nake@uvt.nl