

An Application of the Panel Hurdle Model to Experimental Data

Peter G. Moffatt

Using a user-written STATA program developed by Engel and Moffatt (2014), we apply the panel hurdle model to experimental data from a public goods game. Subjects are assumed to be one of two types: free-riders, and potential contributors. We focus on a test of the “house money effect”. This effect is seen to be significant in the first hurdle: specifically, house money makes a subject more likely to be a potential contributor. Hence we find that the effect of house money is more than just an effect on behaviour; it has the effect of changing a subject from one type to another. This result is potentially important in the external validity debate.

A. Introduction

It is standard practice in economic experiments to provide subjects with an initial endowment, for the obvious reason that it is unreasonable to expect them to use their own money in carrying out the tasks that they are required to perform during the experiment. However, an important question is whether the provision of this endowment has any effect on behaviour in the tasks, since this would imply that the endowment is causing behaviour in the experiment to diverge from behaviour that would be expected in real decisions. This effect has come to be known as the “house money effect”.

There is some evidence of a house money effect in risky choice experiments, the effect being to make choices more risk-seeking (Thaler and Johnson, 1990; Keasey and Moon, 1996). There is also evidence from dictator game experiments that subjects are much more willing to give money away when the endowment is unearned rather than earned (Cherry et al., 2002). In a metastudy, Engel (2011) finds that dictators give an average of 29.52% of their endowment if this is “manna from heaven”, while this proportion falls to 12.63% if they had to earn the money in some real effort task.

Here, we are interested in interactive experiments. Clark (2002) has already suggested a number of possible ways in which house money might

have an effect in such settings. First, house money might be used by subjects to “purchase” greater quantities of non-pecuniary goods such as “public-spiritedness”, or “reciprocal fairness”, than they would do in real life. Second, a shift towards more risk seeking may induce subjects to experiment with “strategic giving”, that is, to bet on their co-players being reciprocators (Kreps et al., 1982). What is important about these hypotheses is that they all embody the idea that the provision of house money makes it cheaper, and therefore arguably more likely, that participants trade some pecuniary gain for increasing the payoff of other participants. Critically, (some) subjects are expected to condition their choices on the presence or absence of house money, i.e. on a general feature of the treatment, irrespective of their experiences from repeated interaction with other subjects. This yields a directed hypothesis: if they receive house money, some subjects no longer play Nash.

Now, even subjects who hold social preferences have been shown, in the main, not to be unconditionally cooperative. Rather they condition their willingness to cooperate on their beliefs about the cooperativeness of their interaction partners and, if the game is repeated, on their experiences (Fischbacher and Gächter 2001, 2010). Provided they are pessimistic, or they have experienced little cooperativeness in the current experiment, such participants will not make positive contributions either. This provides a challenge for parametric estimation. One needs to distinguish those participants who do not contribute as a consequence of their type, from other participants who do not contribute for reason of the context being insufficiently favourable.

One statistical model that allows for estimating the relative importance of both causes for seemingly selfish behaviour is the double hurdle model. The double hurdle model allows for two types of zero contribution: zeros which result from the subject being a free-rider, or a “zero-type”; and zeros which result from the circumstances of the task, in spite of the subject being a potential contributor. Hurdle models have a number of advantages. They make it possible to incorporate the “zero-type”, and therefore do not collapse zero-type with zero-censoring. Moreover, the hurdle model allows estimation of the proportion of the population who are of the zero-type. Better still, it allows the probability of a subject being of zero-type to depend on observables.

Since the data set that we shall be working with consists of repeated tasks for each subject, we shall be estimating the panel version of the double hurdle model (Dong, et al. 2004; Dong and Kaiser, 2008; Engel and Moffatt, 2014). The extension of the model to panel data is demanding because

the outcome of the “first hurdle”, that is, the determination of whether a subject is of the zero-type, must apply to that subject for the whole experiment. Switching in and out of the zero-type is ruled out. In contrast, the outcome of the second hurdle, that is, the amount actually contributed in any particular task, is determined at the level of individual observations. In principle, subjects classified as “free-riders” must necessarily contribute zero in *every* task. This may be seen as too strict a definition though; we might like to make it possible for subjects to be classified as free-riders even if they contribute positive amounts on only a small number of occasions. For this purpose, we introduce a “tremble parameter”, which allows such observations to be attributed to lapses of concentration.

We use a data set that has kept experimental economists engaged for more than two decades. The data was originally generated by Clark (2002). In his experiment, there are 10 rounds, and subjects are in groups of 5, with the make-up of groups changing between rounds. In each round, the endowment is 80 tokens, which subjects are required to allocate between an individual account and a group account. There are 150 subjects, of whom 75 are in the “house money” treatment (given 80 tokens at the start of each round), and 75 are in the “own money” treatment (nothing given to them).

Clark (2002) himself did not find a significant house money effect. Yet this non-result might follow from the fact that he used unconditional non-parametric tests with low statistical power. Harrison (2007) re-examined the same data set parametrically and at the individual level, and found evidence of a house money effect. Harrison discussed the important role of zero observations in this data set, and in particular, found that house money impacts on the proportion of free-riding.

The objective of this paper is to continue where Harrison (2007) left off, and to estimate a fully parametric panel double hurdle model with dependence between the two hurdles, and to test for the house money effect within this framework. The dependent variable is the amount contributed to the group account in a particular round. In order to allow house money to impact on type, a dummy for the house money treatment is included in the first hurdle. Key explanatory variables appearing in the second hurdle are the round number, and the mean contribution of the other group members in the previous round.

The remainder of the paper is organized as follows. In section B, we introduce the statistical model. In section C we describe Clark’s data. In section D, we present our results. Section E concludes.

B. The panel hurdle model

Our statistical model assumes that two hurdles need to be passed in order for a positive contribution to be observed. The first hurdle determines whether the subject is a free-rider; if they fall at the first hurdle, they are a free-rider, and their contribution will be zero in all tasks. If the first hurdle is passed, a “second hurdle” is encountered for each task, and whether this is crossed determines whether the contribution in that task is zero or positive. This data generating process consists of repeated observations, which is why we develop a panel estimator. Another feature of the data that we need to deal with is upper censoring, which arises because there is an upper limit to the contribution.

I. Previous Literature

The basic form of the panel hurdle model has been introduced, and applied to data on food demand, by Dong, et al. (2004) and Dong and Kaiser (2008). STATA code to estimate the panel hurdle model has been provided by Engel and Moffatt (2014), and further detail on practical issues of estimation was provided by Moffatt (2020). The model has been applied in many other settings, including public goods games (Cheung, 2014), tax evasion (Guerra and Harrington, 2018), and corporate emissions trading behaviour (Abrell et al., 2021).

Our application is to data on a public goods experiment originally published by Clark (2002). A previous version of the present paper appeared as Engel and Moffatt (2012).

II. Basic Panel hurdle Model with upper censoring

Let i index subject ($i=1, \dots, n$) and t index task ($t=1, \dots, T$). Let y_{it} be the contribution of subject i in task t , and let y_{max} be the maximum possible contribution (80 in the experiment). Let z_i be a vector of characteristics of subject i relevant in explaining whether the first hurdle is passed. Let x_{it} be a vector of characteristics of subject i and/or task t that determine the contribution in that task.

First Hurdle:

$$\begin{aligned}d_i^* &= z_i' \alpha + \varepsilon_{1,i} \\d_i &= 1 \text{ if } d_i^* > 0; d_i = 0 \text{ otherwise.} \\ \varepsilon_{1,i} &\sim N(0, 1)\end{aligned}\tag{1}$$

Second Hurdle:

$$\begin{aligned}y_{it}^{**} &= x_{it}' \beta + u_i + \varepsilon_{2,it} \\y_{it}^* &= 0 \text{ if } y_{it}^{**} \leq 0; \\y_{it}^* &= y_{it}^{**} \text{ if } 0 < y_{it}^{**} < y_{\max}; \\y_{it}^* &= y_{\max} \text{ if } y_{it}^{**} \geq y_{\max}. \\ \varepsilon_{2,it} &\sim N(0, \sigma^2); u_i \sim N(0, \sigma_u^2)\end{aligned}\tag{2}$$

Observed:

$$y_{it} = d_i y_{it}^* \quad i = 1, \dots, n \quad t = 1, \dots, T\tag{3}$$

The central feature of this model is that the first hurdle has only one outcome per subject, and that outcome applies to all observations for that subject. If subject i falls at the first hurdle ($d_i=0$), then all observations on y for subject i must be zero ($y_{it}=0$, $t=1, \dots, T$). If subject i does not fall on the first hurdle, i.e. if she makes a positive contribution in the first round, this suffices to determine the subject's type. She is not a free-rider. If the subject contributes 0 in any later period, even in all later periods, this must be in reaction to observed or expected circumstances, not to type.

Note also that there is a subject-specific random effect term (u_i) appearing in the second hurdle that allows between-subject heterogeneity. Conditional on $d_i=1$, we obtain the random-effects 2-limit Tobit likelihood:

$$\begin{aligned}(L_i | d_i = 1, u_i) &= \prod_{t=1}^T \left[1 - \Phi \left(\frac{x_{it}' \beta + u_i}{\sigma} \right) \right]^{I(y_{it}=0)} \\ &\quad \left[\frac{1}{\sigma} \phi \left(\frac{y_{it} - x_{it}' \beta - u_i}{\sigma} \right) \right]^{I(0 < y_{it} < y_{\max})} \\ &\quad \times \left[1 - \Phi \left(\frac{y_{\max} - x_{it}' \beta - u_i}{\sigma} \right) \right]^{I(y_{it}=y_{\max})}\end{aligned}\tag{4}$$

Note that, at this stage, we are also conditioning on the value of u_i .

Conditional on $d_i=0$, the likelihood is trivial, and depends simply on whether or not all observations are zero for subject i :

$$\begin{aligned}(L_i|d_i = 0) &= 0 \text{ if } \sum_{t=1}^T y_{it} > 0 \\ &= 1 \text{ if } \sum_{t=1}^T y_{it} = 0\end{aligned}\quad (5)$$

The marginal likelihood for subject i is then obtained as:

$$L_i = \int_{-\infty}^{\infty} [\Phi(z_i'\alpha)(L_i|d_i = 1, u) + [1 - \Phi(z_i'\alpha)](L_i|d_i = 0)]f(u; \sigma_u) du \quad (6)$$

where $\Phi(z_i'\alpha)$ is the probability of passing the first hurdle (1), the two conditional likelihoods are as defined in (4) and (5), and $f(u; \sigma_u)$ is the $N(0, \sigma_u^2)$ density function for u . Note that in (6), the i subscript is absent from u ; this is because here u is the random variable over which integration is being performed.

III. Panel hurdle model with tremble

A potential problem with the model described in II above is that it is too rigid. In particular, an individual can only be classified as falling at the first hurdle (i.e. a free-rider, or selfish type) if their observations on y are zero in every task. We might wish to relax this assumption and make it possible to also classify a subject as a free-rider who reports zero in nearly every task, with the rationalisation that the subject may have suffered a lapse of concentration on the small number of tasks for which the observation was positive. Such “lapses of concentration” might be considered to be particularly likely in early rounds, where they might be better interpreted as the result of a misunderstanding of the experiment.

To allow this, we introduce a tremble parameter, ω . This is the probability of a lapse of concentration on any individual task. When this occurs, the distribution of y is assumed to be uniform on $[0, y_{max}]$ where y_{max} is the upper limit of the action space (or, in a setting in which there is no upper limit, the highest observed value of y in the data).

With the introduction of ω , (4) and (5) become:

$$(L_i | d_i = 1, u_i) = \prod_{t=1}^T \left[(1 - \omega) \left(1 - \Phi \left(\frac{x_{it}'\beta + u_i}{\sigma} \right) \right) \right]^{I(y_{it}=0)} \\ \times \left[(1 - \omega) \frac{1}{\sigma} \phi \left(\frac{y_{it} - x_{it}'\beta - u_i}{\sigma} \right) + \frac{\omega}{y_{\max}} \right]^{I(0 < y_{it} < y_{\max})} \\ \times \left[(1 - \omega) \left(1 - \Phi \left(\frac{y_{\max} - x_{it}'\beta - u_i}{\sigma} \right) \right) \right]^{I(y_{it}=y_{\max})} \quad (7)$$

$$(L_i | d_i = 0) = \omega^{N_i(y_{it} > 0)} (1 - \omega)^{N_i(y_{it} = 0)} \quad (8)$$

where $N_i(\cdot)$ is the number of occurrences of the event contained in parentheses, for subject i .

The marginal likelihood is then constructed as in (6), but using (7) and (8) in place of (4) and (5).

A technical problem with the model defined in (7) and (8) is that the log-likelihood function is monotonically increasing in the parameter ω . Intuitively, this can be interpreted in terms of a tendency for all of the positive observations to be attributed to the tremble. This problem does not appear to arise in other classes of model in which tremble parameters are introduced (see Moffatt and Peters, 2001).

One solution to this estimation problem is simply to assume the presence of a tremble, rather than to attempt to estimate the tremble parameter. That is, we propose that a reasonably small value for ω is chosen, say 0.01, and the other parameters of the model are estimated on the assumption that this is the true tremble probability.

IV. Panel hurdle model with dependence

The models developed in 2.2 and 2.3 assume that there is no correlation between the error terms in the two hurdles. In this section, this assumption is relaxed. Subject i 's idiosyncratic propensity to pass the first hurdle is represented by the error term $\varepsilon_{1,i}$; her idiosyncratic propensity to contribute, conditional on passing the first hurdle, is represented by u_i . It is between these two terms that we introduce a correlation:

$$\text{corr}(\varepsilon_1, u) = \rho \quad (9)$$

How is the correlation parameter ρ incorporated in estimation? Let us return to the first hurdle:

$$\begin{aligned}d_i^* &= z_i' \alpha + \varepsilon_{1,i} \\d_i &= 1 \text{ if } d_i^* > 0; d_i = 0 \text{ otherwise.} \\ \varepsilon_{1,i} &\sim N(0, 1)\end{aligned}\tag{10}$$

Since $\text{corr}(\varepsilon_1, u) = \rho$ we may represent ε_i as (suppressing the i subscript):

$$\varepsilon_1 = \rho \frac{u}{\sigma_u} + \sqrt{1 - \rho^2} \xi \tag{11}$$

where $\xi \sim N(0, 1)$ and $\xi \perp u$. The requirement for passing the first hurdle becomes:

$$d_i = 1 \text{ if } \xi > -\frac{z_i' \alpha + \rho \frac{u}{\sigma_u}}{\sqrt{1 - \rho^2}}. \tag{12}$$

whence the probability of passing the first hurdle (conditional on u) becomes:

$$\Phi \left(\frac{z_i' \alpha + \rho \frac{u}{\sigma_u}}{\sqrt{1 - \rho^2}} \right) \tag{13}$$

The subject i log-likelihood function becomes:

$$\begin{aligned}L_i &= \int_{-\infty}^{\infty} \left[\Phi \left(\frac{z_i' \alpha + \rho \frac{u}{\sigma_u}}{\sqrt{1 - \rho^2}} \right) (L_i | d_i = 1, u) \right. \\ &\quad \left. + \left[1 - \Phi \left(\frac{z_i' \alpha + \rho \frac{u}{\sigma_u}}{\sqrt{1 - \rho^2}} \right) \right] (L_i | d_i = 0) \right] f(u; \sigma_u) du\end{aligned}\tag{14}$$

where $(L_i | d_i, u)$ and $(L_i | d_i = 0)$ are defined in either (4) and (5), or (7) and (8), depending on whether a tremble is assumed.

In a double hurdle model with just one observation per subject, there are known to be problems identifying the correlation coefficient ρ (Smith, 2003). By contrast, with panel data, and using the estimation approach outlined in this section, the parameter may be estimated.

V. Estimation

Estimation is performed using the method of Maximum Simulated Likelihood (MSL, Train, 2003). This requires the use of Halton draws, which, when converted to normality, represent simulated realisations of the random effect term u . In the model with dependence, in accordance with (13), the simulated values also appear in the probability of passing the first hurdle. Maximisation of the simulated likelihood function is performed using the `ml` routine in STATA. All programs are available from the authors upon request.

C. Data

The data set we use is from the experiment conducted by Clark (2002). In this experiment, there are 10 rounds, and subjects are in groups of 5, with the make-up of groups changing between rounds. In each round, the endowment is 80 tokens, which subjects are required to allocate between an individual account and a group account. There are 150 subjects, of whom 75 are in the “house money” treatment (given 80 tokens at the start of each round), and 75 are in the “own money” treatment (nothing given to them).




Figure 1: Histogram of Contributions by treatment.

Figure 1 shows that, both with and without house money, the data is left and (less intensely) right censored. The key difference between the two




Figure 2: Plot of contribution against round number.

histograms is that participants appear to be more likely to contribute zero if they are using their own money.

Figure 2 shows a plot of contribution against round number, with a lowess smoother. The smoother reveals that contributions slightly decline as the experiment progresses. Hence the inclusion of round number in the second hurdle is justified.

Figure 3 shows a plot of contribution against the mean of contributions of the other group members in the previous round. An important feature of the experiment is that in each round subjects are made aware of the group mean in the previous round. It is clear from the smoother that participants are sensitive to this measure of recent experience, justifying its inclusion as an explanatory variable in the second hurdle. The reason for using the lagged mean contribution of *other* group members, rather than the group mean, is to guard against the consistency problems inherent in a dynamic panel. At the same time, this variable explains, rather than just neutralizes, the additional dependence resulting from the fact that choices are nested in participants, who are nested in groups. Since groups are randomly and anonymously composed, group effects collapse with the experiences participants encounter within their group.

When using the lagged mean contribution of other group members as an explanatory variable, a slight problem that arises is that this variable contains missing values in the first round. Contributions from the first round




Figure 3: Plot of contribution against mean of contribution of the *other* group members in previous round.

could be excluded from estimation. However, rather than doing this, we follow Bardsley and Moffatt (2007, see their Table II, note 4) by assuming that, in the first round, subjects make a judgement (perhaps based on lifetime experience prior to the experiment) of what a typical contribution of others is likely to be. This value is then used as the value of the lagged mean in the first round. The value we use for this purpose is, in fact, 66 units, and this is obtained, via a grid search, as the value which maximises the sample log-likelihood. This relatively high value implies that subjects tend to start the experiment optimistically. While this assumption may be viewed as somewhat arbitrary, the main result of the paper (relating to the house money effect) is in fact robust to the choice of value made here.

D. Results

Results from various models are shown in Table 1. The first is a binary probit model estimated using the cross section of 150 subjects, with dependent variable taking the value 1 if at least one of the subject's contributions is positive, and zero otherwise. This model gives estimates of the first hurdle, which are used as starting values in the estimation of the corresponding part of the hurdle model. The second model is a panel tobit

Table 1: Results of hurdle and related models.

	probit	Panel tobit	Hurdle ($\rho=0$; $\omega=0$)	Hurdle with dep. ($\omega=0$)	Hurdle dep. ($\omega=0.001$)
First hurdle					
Const	0.941 (0.171)		0.968 (0.177)	0.967 (0.161)	0.673 (0.148)
House money	0.379 (0.263)		0.393 (0.276)	0.256 (0.249)	0.497** (0.236)
Second hurdle					
Const		27.251 (10.453)	42.608 (10.388)	35.027 (10.364)	29.283 (10.010)
round		-2.986*** (0.394)	-3.094*** (0.399)	-3.111*** (0.396)	-2.948*** (0.393)
previous		1.016** (0.392)	0.954** (0.398)	0.941** (0.392)	1.177*** (0.376)
previous ²		-0.013*** (0.004)	-0.013*** (0.004)	-0.013*** (0.004)	-0.015*** (0.004)
σ		25.943 (0.681)	26.090 (0.689)	26.053 (0.688)	25.097 (0.696)
σ_u		43.562 (3.191)	29.986 (2.375)	35.457 (2.950)	34.753 (3.344)
ρ			0	0.901 (0.091)	0.940 (0.065)
n	150	150	150	150	150
T	–	10	10	10	10
LogL	-57.85	-4722.89	-4707.01	-4705.19	-4688.21
AIC (2k-2LogL)		9457.78	9426.38	9428.38	9416.42

* mildly significant ($p<0.10$); ** significant ($p<0.05$); *** strongly significant ($p<0.01$).

model, which gives estimates of the second hurdle; these are also used as starting values. Then, three panel hurdle models are estimated: the basic model; the model with dependence; the model with dependence and a non-zero tremble probability. The last of these is the best model on the basis of the Akaike Information Criterion (AIC).

The effects seen in the second hurdle are as expected. Firstly, In agreement with other public good studies, and with Figure 2 above, the




Figure 4: Posterior probability of passing first hurdle against number of positive contributions (posterior probability obtained from model with no tremble).

negative coefficient of round number provides evidence of a decline in contributions over the course of the experiment. Secondly, the positive coefficient of the mean previous contribution by others, and the negative coefficient of its square, confirm the conclusion drawn from Figure 3, that contributions rise but level off with others' contributions. This amounts to evidence of reciprocity.

The key result is the evidence of the house money effect in the final column: house money increases the probability of passing the first hurdle, that is, it increases the probability of being a "potential contributor". Note that this result is only seen in the final column.

The role of the tremble parameter is crucial in reaching this conclusion. In order to see why, we first plot the posterior probability of passing the first hurdle (obtained using Bayes' Rule) from the model with *no* tremble (whose results are presented in the penultimate column of Table 1). Note that all subjects with positive contributions are certain to pass the first hurdle.

Next, we present the same plot, but for the model in which the tremble probability is set to 0.001 (i.e. the model whose results are reported in the final column of Table 1). Note that some subjects with a small number




Figure 5: Posterior probability of passing first hurdle against number of positive contributions (posterior probability obtained from model with tremble set to 0.001). H=house money treatment; O=own-money treatment.

of positive contributions have high posterior probability of being a free-rider (being low down in the plot). Note further that these subjects tend to be in the “own-money” treatment, and it is this that drives the difference in results between the penultimate and final columns of Table 1. The small number of subjects in the own-money treatment who are being attributed to the “free-rider” type, despite their occasional contributions, are tilting the balance between no house money effect and a house money effect.

E. Conclusion

In experimental practice, it is common to give participants an endowment. This routine makes it easier to run experiments, and goes largely unquestioned. Yet in reality, people are clearly not given money when deciding how much to donate to charity, or when deciding where to invest. It is their own earned income that is at stake.

Clark’s (2002) non-result was therefore a relief for experimentalists. Seemingly, giving participants some money upfront did not endanger ex-

ternal validity. Our paper cautions against such complacency. If only one uses an estimator that matches the complex data generating process, it turns out that, already in Clark's original data, house money moves some individuals out of the free-riding type and into the contributor type.

Given our finding, it seems that the house money debate ended prematurely. More experiments are needed to show under which conditions and, in particular with which experimental designs, the house money effect is negligible, and under which conditions it creates a serious obstacle for extrapolation from the lab to the field.

To enable the thorough investigation of the house money effect, an "own-money" treatment is necessary, and this presents obvious ethical issues. A very useful complement would be field experiments in which subjects are operating in their natural environment where "own money" is the default.

References

- Abrell, J., Cludius, J., Lehmann, S., Schleich, J., & Betz, R. (2022). Corporate emissions-trading behaviour during the first decade of the EU ETS. *Environmental and Resource Economics*, 83(1), 47–83.
- Bardsley N. and Moffatt P.G. (2007), "The experimentics of public goods: inferring motivations from contributions", *Theory and Decision*, 62, 161–193.
- Cherry T.L., Frykblom P. and Shogren J.F., (2002) "Hardnose the Dictator", *American Economic Review*, 92 (4), 1218–1221.
- Cheung, S. L. (2014). New insights into conditional cooperation and punishment from a strategy method experiment. *Experimental Economics*, 17(1), 129–153.
- Clark J.(2002), "House money effects in public good experiments", *Experimental Economics*, 5, 223–231.
- Dong, D., Chung, C., and Kaiser, H.M. (2004), Modeling milk purchasing behavior with a panel data double-hurdle model. *Applied Economics*, 36(8), 769–779.
- Dong D. and H.M. Kaiser (2008), "Studying Household Purchasing and Nonpurchasing Behaviour for a Frequently Consumed Commodity: Two Models", *Applied Economics*, 40(15), 1941–1951.
- Engel, C. (2011): "Dictator Games: A Meta Study", *Experimental Economics*, 14(4), 583–610.
- Engel, C., & Moffatt, P. G. (2012). Estimation of the house money effect using hurdle models. *MPI Collective Goods Preprint*, (2012/13).
- Engel, C., & Moffatt, P. G. (2014). `dhreg`, `xtdhreg`, and `bootdhreg`: Commands to implement double-hurdle regression. *The Stata Journal*, 14(4), 778–797.
- Fischbacher, U. and S. Gächter (2001): "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment", *Economics Letters* 71, 397–404.

- Fischbacher, U. and S. Gächter (2010): “Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments”, *American Economic Review*, 100, 541–566.
- Guerra, A., & Harrington, B. (2018). Attitude–behavior consistency in tax compliance: A cross-national comparison. *Journal of Economic Behavior & Organization*, 156, 184–205.
- Harrison G.W., 2007, “House money effects in public good experiments: comment”, *Experimental Economics*, 10, 429–437.
- Keasey K. and Moon P. (1996) “Gambling with the house money in capital expenditure decisions”, *Economics Letters*, 50, 105–110.
- Kreps, D.M., P. Milgrom et al., 1982, “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma”, *Journal of Economic Theory*, 27, 245–252.
- Moffatt, P.G. (2020). *Experiments: Econometrics for experimental economics*. Bloomsbury Publishing.
- Moffatt, P.G. and Peters S.A., 2001, “Testing for the presence of a tremble in economic experiments”, *Experimental Economics*, 4, 221–228.
- Smith M.D., 2003, “On dependency in double hurdle models”, *Statistical Papers*, 44, 581–595.
- Thaler R. and Johnson E. (1990) “Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice”, *Management Science*, 36, 643–660.
- Train, K.E., 2003, *Discrete Choice Methods with Simulation*, Cambridge.