



The Log of Gravity at 15

J. M. C. Santos Silva¹ · Silvana Tenreyro²

Received: 6 January 2021 / Accepted: 25 August 2021 / Published online: 22 January 2022
© The Author(s) 2021

Abstract

We review the contribution of “The Log of Gravity” (Santos Silva and Tenreyro, *Rev Econ Stat* 88:641–658, 2006), summarize the main results in the ensuing literature, and provide a brief review of the state-of-the-art in the estimation of gravity equations and other constant-elasticity models.

Keywords Gravity equation · Incidental parameters · Jensen’s inequality · Poisson regression · Separation

JEL Classification C13 · F10 · F14 · F21 · F22

1 Introduction

Fifteen years after its publication, this is perhaps a good time to reflect on the influence of our paper “The Log of Gravity” (Santos Silva and Tenreyro 2006).¹ In that paper we challenged the long-established practice of estimating constant-elasticity models in their log-linearized form, and proposed as an alternative the use of an estimator that conveniently coincides with the Poisson pseudo maximum likelihood (PPML) estimator of Gourieroux et al. (1984). Building on early contributions of

¹According to Google Scholar, the paper received more than 750 citations just in 2020.

We are grateful to two anonymous referees and to Michel Beine, José De Sousa, and Tom Zylkin for helpful comments and suggestions. The usual disclaimer applies.

✉ J. M. C. Santos Silva
jmcss@surrey.ac.uk

Silvana Tenreyro
s.tenreyro@lse.ac.uk

¹ School of Economics, University of Surrey, Guildford GU2 7XH, UK

² Department of Economics, London School of Economics and Political Science, London WC2A 2AE, UK

Goldberger (1968), Papke and Wooldridge (1996), and Manning and Mullahy (2001), in Santos Silva and Tenreyro (2006) we presented a clear explanation of why the estimation of log-linearized models could lead to misleading results, provided an unequivocal recommendation for the use of the PPML estimator, and clearly illustrated the advantages of this estimator. In our view, the simple message of the paper and the clarity and relevance of the examples we provided, were the key factors for its popularity.²

In this paper, we consider the reasons for the impact of “The Log of Gravity” and summarize some of the developments that contributed to its enduring relevance. In doing this, we provide a brief review of the state-of-the-art in the estimation of the gravity equation for trade, which may be useful to the less experienced researchers. Many of the methods and developments we discuss are also relevant for the estimation of constant-elasticity (multiplicative) models for other kinds of data, and we also refer to some of these applications.

The remainder of the paper is organized as follows. Section 2 briefly presents the problem with the traditional least squares estimator of gravity equations and Section 3 discusses several aspects of the PPML estimator. Section 4 discusses specification tests for gravity models and Section 5 reviews the simulations and the results of the empirical application we presented in Santos Silva and Tenreyro (2006). In Section 6 we provide examples of the use of the PPML estimator in different fields and, finally, Section 7 contains some brief concluding remarks.

2 The problem

Following Goldberger (1991, p. 5), in Santos Silva and Tenreyro (2006) we interpret non-stochastic economic models such as the gravity equation as the conditional expectation of the variable of interest.³ That is, if economic theory suggests that the non-negative variable y and the vector of explanatory variables x are linked by a constant-elasticity model of the form

$$y = \exp(x\beta), \quad (1)$$

the function $\exp(x\beta)$ is interpreted as the conditional expectation of y given x , denoted $E[y|x]$, where the vector of (semi) elasticities β is the object of interest. An example of a model of this kind is the gravity equation for trade which, in its simplest form, can be written as

$$T = \beta_0 Y^{\beta_1} D^{\beta_2} \quad (2)$$

$$= \exp[\ln(\beta_0) + \beta_1 \ln(Y) + \beta_2 \ln(D)], \quad (3)$$

²Another reason that helps to explain the popularity of the paper is that we replied, and continue to reply, to hundreds of emails with questions about it and always try to provide support to the users. We also created a dedicated website providing data, code, and answers to the most frequently asked questions.

³Alternatively, the model could be interpreted as a different measure of central tendency such as the conditional median or the conditional mode. However, the conditional expectation is a more attractive location measure when the data can have many zeros.

where T denotes the trade flow from an origin to a destination, Y is a measure of the size of the trading partners, D represents the distance between the partners, and β_0 , β_1 and β_2 are unknown parameters.

All econometrics textbooks that we are aware of suggest that the parameters in models such as Eq. 1 can be estimated by the least squares regression of $\ln(y)$ on x . However, this approach may be inappropriate for two reasons. An obvious problem, and our initial motivation to consider alternative estimators, is that this approach is infeasible if y is zero for some observations. The more serious problem is that, due to Jensen's inequality, the least squares regression of $\ln(y)$ on x is generally an inconsistent estimator for the parameters of $E[y|x] = \exp(x\beta)$.⁴

The key insight to understand why the regression in logs is not generally valid is that, although we can go from Eq. 1 to its logarithmic form, and vice-versa, the same is not true for their stochastic counterparts. Indeed, because economic models do not hold exactly, estimation has to be performed using stochastic versions of the equations suggested by economic theory, and that is where Jensen's inequality becomes important.

The stochastic counterpart of Eq. 1 can be written as

$$y = \exp(x\beta) + \varepsilon = \exp(x\beta)\eta, \quad (4)$$

where ε is an additive error term such that $E[\varepsilon|x] = 0$, and $\eta = 1 + \varepsilon/\exp(x\beta)$ is a multiplicative error term with $E[\eta|x] = 1$.⁵ Ignoring for the moment that y can be equal to 0, the model can be made linear in the parameters by taking logarithms of both sides of the equation, leading to

$$\ln(y) = x\beta + \ln(\eta). \quad (5)$$

In Eq. 5, the least squares estimator is consistent for β if $\ln(\eta)$ is uncorrelated with x , but since $\eta = 1 + \varepsilon/\exp(x\beta)$, that condition will be met only under very restrictive conditions on the distributions of ε , and therefore the least squares estimator of the regression defined by Eq. 5 is generally inconsistent for β .⁶ In the next section we consider alternative approaches to estimate β and explain why PPML should be preferred.

⁴This inconsistency is small in many empirical contexts and that explains why the estimation of the log-linear model by least squares is still so popular. However, as we illustrated in Santos Silva and Tenreyro (2006), the inconsistency can be substantial and therefore nowadays it is hard to justify the continued use of this estimator.

⁵We are often asked why we do not write the stochastic version of Eq. 1 as $y = \exp(x\beta + \varepsilon)$. The reason for not doing it is that in this case the conditional expectation of y is not generally given by $\exp(x\beta)$, and therefore this expression is not a proper stochastic version of Eq. 1.

⁶See also Wooldridge (1992). Alternatively, when y is strictly positive, we can interpret the least squares estimator of Eq. 5 as providing consistent estimates of the parameters of the conditional geometric mean; these can be very different from β and can even have different signs (see Reis and Santos Silva (2006), Petersen (2017), Mitnik and Grusky (2020), and Dias and Marques (2021)). However, if y can be zero, the geometric mean is not an interesting measure of central tendency because it is identically zero when $\Pr(y = 0|x) > 0$.

3 The PPML estimator

At first sight, the natural approach to estimate $E[y|x] = \exp(x\beta)$ without transforming the model would be to use non-linear least squares, as done by Frankel and Wei (1993). The problem with this approach is that, as we noted in Santos Silva and Tenreyro (2006), it is based on moment conditions of the form

$$E[\exp(x\beta)(y - \exp(x\beta))x] = 0,$$

which give more weight to the observations with larger variance, and therefore can be inefficient to the point of being useless in empirical applications. This problem has been documented in several simulation studies; see, e.g., Manning and Mullahy (2001) and Santos Silva and Tenreyro (2006, 2011a).

The alternative we proposed in Santos Silva and Tenreyro (2006) is to base the estimator on moment conditions of the form

$$E[(y - \exp(x\beta))x] = 0, \quad (6)$$

which give the same weight to all observations. As will be discussed below, besides being intuitively appealing, this estimator has several other properties that make it particularly attractive in this context.

One of the advantages of the estimator based on Eq. 6 is that it coincides with the Poisson regression estimator and therefore most statistical softwares have commands that make its use very simple. Of course, because in trade-data applications y certainly does not follow a Poisson distribution, this is a pseudo maximum likelihood estimator (see Gourieroux et al. 1984), and a suitably robust estimator of the standard errors should be used.

3.1 Why not use other estimators for count data

The fact that in Santos Silva and Tenreyro (2006) we recommended that gravity equations should be estimated using a method designed for count data generated some misunderstandings.

In count data models, researchers are often interested in estimating the conditional probability of some event, such as $\Pr(y = k|x)$, where k is some non-negative integer. To obtain a consistent estimator of this probability we need to correctly specify the conditional distribution of y , and the Poisson distribution is often seen as too restrictive for this purpose. Therefore, alternative methods based on different distributions have been proposed to estimate count data models, and many of these approaches are more flexible than the basic Poisson regression. This has led some authors to advocate that these estimators would also out-perform the PPML estimator when the objective is to estimate gravity equations. As we explain below, this is wrong.

The first thing to note is that when estimating a gravity equation we want to have an estimator of $E[y|x] = \exp(x\beta)$ that is valid under very mild assumptions, and we do not need to estimate quantities such as $\Pr(y = k|x)$. Therefore, estimators

of β whose consistency depends on incidental distributional assumptions are not as attractive as the PPML estimator, whose consistency depends only on the validity of the assumption that $E[y|x] = \exp(x\beta)$; i.e., that the gravity equation is correctly specified.⁷ Therefore, estimating gravity equations using, for example, the estimator for zero-inflated count data models introduced by Mullahy (1986) is not attractive in this context because the validity of the estimator would depend on very strong assumptions about the distribution of the data.

Another aspect to note is that, in the context of count data models, most of the alternatives to Poisson regression allow for the so-called overdispersion (see, e.g., Cameron and Trivedi, 2013). However, overdispersion is not defined when the dependent variable does not have a natural scale. Indeed, when the dependent variable can be measured in different units, the relation between the conditional mean and the conditional variance will depend on the scale of the data. This implies that estimates obtained using models that allow for overdispersion are sensitive to the scale of the dependent variable and to the units in which it is measured, and therefore are arbitrary. This problem was noted by Bosquet and Boulhol (2014) for the case of the negative binomial estimator, but it affects all estimators that try to accommodate overdispersion, such as the zero-inflated models whose use has been recommended by some authors.

3.2 PPML, fixed effects, and the incidental parameter problem

Since the seminal work of Anderson and van Wincoop (2003), it has become standard to estimate gravity equations accounting for multilateral resistance by including a dummy for each origin and a dummy for each destination, the so-called origin and destination fixed effects (see also Hummels, 1999). In this case, the number of parameters to estimate depends on the number of countries included in the sample, and therefore we need to account for the incidental parameter problem because, in general, it is not possible to obtain consistent estimators for models in which the number of parameters depends on the sample size (see, e.g., Lancaster, 2000).

It is well known that PPML does not suffer from the incidental parameter problem in the traditional panel data case where a single fixed effect is included; see Wooldridge (1999). Because that result does not cover models with two sets of fixed effects, some authors have claimed that PPML suffers from the incidental parameter problem when the model includes origin and destination fixed effects. That claim is, however, incorrect. Indeed, Fernández-Val and Weidner (2016, p. 301) have shown that PPML is immune to the incidental parameter problem in models with two sets of fixed effects, as long as the sizes of the two sets of fixed effects grow at the same rate and the regressors are strictly exogenous or predetermined.

Although PPML is consistent in the two-way gravity model, the usual estimator of the covariance matrix accounting for clustering is invalid due to the incidental param-

⁷This is equivalent to saying that PPML is consistent as long as in Eq. 4 the random disturbances satisfy $E[\varepsilon|x] = 0$, which imply $E[\eta|x] = 1$; no additional assumptions are needed on the distributions of ε and η .

eter problem (see, e.g., Egger and Staub (2015), and Jochmans (2017)). Weidner and Zylkin (2020) provide a solution to this problem.⁸

Following the suggestion of Baier and Bergstrand (2007), researchers sometimes use panel data to estimate three-way gravity models that include origin-time and destination-time fixed effects, as well as pair-fixed effects. The consistency of the PPML estimator in this context does not follow from the results of Fernández-Val and Weidner (2016), but Weidner and Zylkin (2020) have recently shown that PPML is still consistent in this context. Remarkably, Weidner and Zylkin (2020) also show that PPML is the only member of a family of pseudo maximum likelihood estimators that has this property. Weidner and Zylkin (2020) also show that standard significance tests and confidence intervals are invalid in three-way gravity models because the asymptotic bias and the asymptotic standard deviation of the estimator vanish at the same rate. Weidner and Zylkin (2020) propose solutions for this problem and we refer the interested reader to their paper for more details.

Baier and Bergstrand (2007) motivation to introduce a third set of fixed effects is to control for the possible endogeneity of free trade agreements. An alternative way to address this issue would be to use instrumental variable methods, but it is difficult to find convincing instruments that can be used in this context. Additionally, the estimation of gravity models with endogenous regressors is challenging because the instrumental variables counterparts of the PPML estimator (Mullahy 1997; Windmeijer and Santos Silva 1997) suffer from the incidental parameter problem and therefore cannot be used to estimate models that include fixed effects. However, Jochmans's (2017) estimator can be used in this context because it partials-out the origin and destination fixed effects, and therefore does not suffer from the incidental parameter problem. Jochmans and Verardi (2019) present a Stata command that implements Jochmans's (2017) instrumental variables estimator for the case of gravity equations with two-way fixed effects estimated with cross-sectional data.

3.3 Computational aspects

One of the advantages of the PPML estimator is that its objective function is globally concave and therefore it has at most one maximum; the problem is that there are cases where the pseudo loglikelihood function does not have a maximum, and therefore the estimates do not exist. Heuristically, this problem is caused by the presence of regressors that perfectly predict some of the observations for which the dependent variable is zero, implying that the maximum likelihood estimator of their coefficients goes to (minus) infinity.

It is well known that the presence of perfect predictors can lead to the non-existence of the maximum likelihood estimates for binary choice models such as the

⁸Accounting for clustering, something we failed to do in Santos Silva and Tenreyro (2006), requires the researcher to define the relevant clustering structure. The standard practice (see, e.g., Yotov et al. (2016)) is to cluster by the pair identifier, but other approaches have been suggested (see, e.g., Egger and Tarlea 2015). This is an important issue and more research is needed on how to best estimate standard errors in the presence of a potential complex pattern of dependencies in this kind of data.

logit (see, e.g., Albert and Anderson, 1984), but it is much less known that such problem also affects the PPML and other estimators such as the Tobit.

In Santos Silva and Tenreyro (2010), we described the issue and provided a simple method to detect and solve this problem. Subsequently, in Santos Silva and Tenreyro (2011b) we described other numerical issues that can lead to convergence problems and introduced the `ppml` Stata command, which implements the methods discussed in Santos Silva and Tenreyro (2010).

More recently, Correia et al. (2019) revisited the problem and, building on much earlier contributions by Verbeek (1989, 1992) and Wedderburn (1976), presented a refined version of the algorithm to detect the non-existence of the PPML estimates. This method, and the associated solution to the problem of non-existence, are implemented in their `ppmlhdfe` Stata command (Correia et al. 2020). In practice, both `ppml` and `ppmlhdfe` effectively deal with the non-existence problem and therefore nowadays this is not a serious issue in empirical applications.

An interesting result in Correia et al.'s (2019) paper is that Poisson regression is rather special in that the solution to the non-existence of the estimates is simpler in that case than in related estimators such as the gamma and inverse Gaussian pseudo maximum likelihood estimators. This, therefore, is another reason to prefer PPML to other generalized linear models for non-negative data.

The non-existence of the PPML estimates is particularly likely to occur in models with a large number of dummy variables, such as models with origin and destination fixed effects. Although the non-existence in itself is not problematic, estimation of these models is challenging due to the sheer number of parameters that have to be estimated. Correia et al. (2020) address this issue in their `ppmlhdfe` Stata command. Combining earlier results by Guimarães and Portugal (2010) with the Frisch-Waugh-Lovell theorem, Correia et al. (2020) develop an algorithm that greatly simplifies the estimation by PPML of models with multiple sets of fixed effects.⁹

3.4 PPML and structural gravity

The gravity equation provides a reliable way to describe trade flows and to evaluate the partial equilibrium effects of trade policies. To go beyond the partial equilibrium analysis, which ignores the effect of trade policies on third-party countries, we need structural gravity models that take into account the general equilibrium effects of trade policies.

Anderson and van Wincoop (2003) introduced a structural gravity model that permits the general equilibrium analysis of trade policies by considering their effects through multilateral resistance channels. Anderson and van Wincoop (2003) estimate their structural gravity model using a non-linear method, but notice that an alternative is simply to include origin and destination fixed effects in a standard gravity equation, as done by Hummels (1999). However, in general, there is no guarantee that the estimated fixed effects are consistent with the definition of the multilateral resistance indexes and with the equilibrium conditions that they must verify. Remarkably,

⁹For comparable packages in R, see Bergé (2018), Stammann (2018), and Hinz et al. (2019).

Fally (2015) has demonstrated that under reasonable assumptions the estimated fixed effects automatically satisfy these conditions when the gravity equation is estimated by PPML, and therefore the multilateral resistance indexes can be recovered from the estimated fixed effects. Moreover, Fally (2015) also shows that PPML is the only pseudo maximum likelihood estimator with this property.

Building on Fally's (2015) results, Anderson et al. (2018) propose a method to compute general equilibrium effects of trade policies based on a structural gravity model and on the properties of the PPML estimator; see also Yotov et al. (2016).

4 Specification tests

In general, constant-elasticity models can be estimated consistently using any of the pseudo maximum likelihood estimators introduced by Gourieroux et al. (1984). Because all these estimators are consistent under the same mild set of conditions, researchers may use specification tests to choose the best estimator in this family; i.e., the more efficient pseudo maximum likelihood estimator. Manning and Mullahy (2001) suggested that the traditional Park (1966) test could be used for this purpose, but in Santos Silva and Tenreyro (2006) we noted that the test is generally invalid in this context and proposed alternative approaches. However, as we explain below, both the Park (1966) test suggested by Manning and Mullahy (2001) and the tests we suggested in Santos Silva and Tenreyro (2006) are of little use when estimating gravity equations.

As noted in the previous section, the PPML estimator is the only pseudo maximum likelihood estimator for gravity equations that is valid under very mild assumptions, that is valid in models with high-dimensional fixed effects, that is not adversely affected by the possible non-existence of the estimates, and whose results are compatible with structural gravity models. Therefore, there is not really much choice when it comes to selecting a pseudo maximum likelihood estimator for a gravity equation, and the PPML is the only credible option. In other words, PPML is efficient in the class of pseudo maximum likelihood estimators that are valid in models with fixed effects and are compatible with structural gravity models. Therefore, tests to check the relation between the conditional mean and the conditional variance, such as those proposed in Manning and Mullahy (2001) and Santos Silva and Tenreyro (2006), are redundant when the purpose is to estimate gravity equations, and they serve no purpose in this context.¹⁰

In Santos Silva and Tenreyro (2006) we also used a version of Ramsey's (1969) RESET test to check the specification of the models. Although often misinterpreted as a test for omitted variables, the RESET is a very useful general misspecification

¹⁰These tests may, however, be useful when the model being estimated is not a gravity equation. In those cases, we recommend the test based on the estimation by PPML of the regression in equation (12) of Santos Silva and Tenreyro (2006).

test and it can be useful to check the specification of gravity equations (not to choose the estimation method). One thing to keep in mind when performing a RESET-type test in models with fixed effects is that some of the fixed effects may be estimated with a very small number of observations and therefore their estimates will be very noisy. In this case, the fitted values of the linear index whose powers are used in the test should not include the estimates of the fixed effects.

The standard formulation of the gravity equation has been extremely successful in practice and has solid theoretical underpinnings (see, e.g., Anderson and van Wincoop (2003), and the references therein). This standard formulation is a single-index model in which zero and positive observations of trade are treated in the same way. However, authors such as Helpman et al. (2008) have suggested double-index trade models that separate the extensive-margin decision to export from the intensive-margin decision of how much to export.¹¹ In other areas (e.g., health economics) it is also often the case that researchers have to choose between single- and double-index models, and therefore it is interesting to have a method to choose between these competing specifications for models for non-negative data.

Because the standard gravity equation and most single-index models can be estimated by PPML, which does not require the correct specification of the likelihood function, the choice between single- and double-index models for trade cannot be based on information criteria because these are likelihood based.¹² Likewise, the vast majority of tests for non-nested hypotheses also cannot be used for this purpose because they are also likelihood based. However, in Santos Silva et al. (2015) we developed a simple test that can be used for this purpose. The test has not been widely used and that probably reflects the fact that most researchers are comfortable with the traditional gravity equation and do not consider double-index alternatives.

5 Simulations and application

In Santos Silva and Tenreyro (2006) we provided overwhelming simulation evidence that the traditional approach of estimating gravity equations using the least squares regression of $\ln(y)$ on x could lead to very misleading results, and that PPML is generally very well behaved, even when it is not the optimal estimator. However, the dependent variable in the main simulation design considered by Santos Silva and Tenreyro (2006) is strictly positive. The fact that the dependent variable did not include zeros led several researchers to question the validity of our results, and to

¹¹In Santos Silva et al. (2014) we use the framework of Helpman et al. (2008) to present a model for the extensive margin of trade defined as the number of exporting sectors, and propose a suitable estimator.

¹²More generally, information criteria such as the popular AIC or BIC are not useful to compare models estimated by pseudo maximum likelihood and they are not invariant to the scale of the dependent variable. Likewise, goodness-of-fit measures based on the likelihood are also not valid in this context.

unfounded claims that PPML performed poorly in situations where the dependent variable has many zeros.

The reason why we used a strictly positive dependent variable in our main simulations is simple: at the time we did not know how to generate non-negative data with zeros and with an exponential conditional expectation.¹³ We solved this problem in Santos Silva and Tenreyro (2011a) by introducing an attractive data generating process in which the dependent variable can have an arbitrarily-high proportion of zeros and has an exponential expectation.¹⁴ The simulation results presented in Santos Silva and Tenreyro (2011a) confirmed that the performance of PPML is very strong even in the presence of a very high percentage of zeros and, together with the theoretical properties of the PPML estimator established by Gourieroux et al. (1984), should be enough to convince even the more skeptic that the fact that the dependent variable can have a high proportion of zeros does not affect the performance of the PPML estimator.¹⁵

The empirical illustration we presented in Santos Silva and Tenreyro (2006) confirmed that the results obtained with the PPML estimator were substantively different from those obtained with the traditional method and with other methods that are difficult to justify in this context, such as estimators based on the Tobit and estimators based on adding an arbitrary constant to the dependent variable before taking logs.¹⁶

The application also provided an unexpected result that was later confirmed by many other authors: the PPML estimates change very little if the estimation is performed excluding the observations for which the dependent variable is zero. With the benefit of hindsight, we were able to explain why dropping the zeros has little impact on the PPML estimates: observations where the conditional mean is close to zero have low variance and therefore the residuals are close to zero for observations for which the value of trade is small or zero. This implies that observations for which the dependent variable is equal to zero have a very small contribution to the value of the pseudo loglikelihood function, and therefore contribute little to the estimation results. Therefore, what was our initial motivation for using PPML turned out not to be particularly important, but the problems caused by disregarding the implications of Jensen's inequality were more serious than we anticipated.¹⁷

¹³This difficulty also explains why other researchers found that PPML did not perform well when the data has zeros: their data had zeros but did not have an exponential conditional expectation, and therefore PPML is not suitable in that case.

¹⁴See also Eaton et al. (2013).

¹⁵Some researchers are still unconvinced, but hopefully those unfounded worries will be laid to rest soon.

¹⁶Many other studies have confirmed that the PPML estimates are materially different from those obtained using the traditional approaches; De Sousa (2012) is a particularly clear example of this.

¹⁷We estimated models by PPML with and without zeros because we wanted to understand whether it was the different sample that was driving the difference between the PPML estimates and those obtained with the traditional least squares method. We often see that other researchers also estimate models by PPML with and without zeros, but little is gained by doing that now.

6 The PPML estimator in other contexts

The suggestion that the well-established practice of estimating elasticities using log-linear regressions could lead to misleading results was initially met with skepticism;¹⁸ even the referees noted that they “were unconvinced by the practical importance of the issue.” However, the importance of the problem has gradually been recognized and PPML is now widely used for the estimation of gravity equations for trade.

However, as we noted in Santos Silva and Tenreyro (2006), the PPML estimator can be used in a broad range of economic applications where the equations under study are traditionally estimated in their log-linearized form, and PPML is now also gaining acceptance in many other areas. Although it is not possible to provide here a comprehensive review of all the applications that have used the PPML estimator, in this section we refer some interesting examples of its use to estimate gravity equations and other models.

Gravity equations are frequently used in the study of migration flows (see, e.g., Beine et al., 2016). In these studies, the dependent variable is often (but not always) a count, and therefore the use of Poisson regression in this context is even more natural. Indeed, the use of this method was suggested by Flowerdew and Aitkin (1982), but at that time the attractive properties of Poisson regression were not yet known, and therefore this work did not have much impact. More recent work (e.g., Beine and Parsons, 2015) use PPML to estimate models that include a number of fixed effects and where the dependent variable is not a count, very much like in the trade literature.

The study of foreign direct investment (FDI) also relies heavily on the gravity equation, and PPML is now often used in this context (see, e.g., Head and Ries (2008), for an early example). Here, however, there is a possible complication: net FDI flows can be negative. The fact that some observations are negative does not imply that the gravity equation is inadequate and that the PPML estimator should not be applied. Indeed, all that is needed for the validity of the PPML estimator in this context is that the conditional expectation of the net flows is given by the gravity equation, and therefore is always non-negative. If that is the case, the PPML estimator continues to be appropriate even if some net FDI flows are negative.¹⁹

Going beyond the estimation of gravity equations, and reflecting the influence of the pioneering work of Manning and Mullahy (2001), we find many examples of PPML estimation in health economics. For example, Kaiser et al. (2014) use PPML to evaluate the impact of a reform on the retail price of drugs, and Powell and Seabury (2018) use PPML to estimate models for medical expenditures. Models for other kinds of expenditures have also been estimated by PPML. For example, Fisher (2016) uses PPML to estimate models for household expenditures, and Jeong and Siegel (2018) use PPML to estimate models for bribes paid by businesses.

¹⁸In the first few years after the publication of our paper, many authors claimed that our result was incorrect. Those claims are now less frequent.

¹⁹Note, however, that some softwares will not estimate Poisson regressions when the dependent variable has negative observations.

Another early use of the PPML estimator outside of the trade literature relates to the estimation of wage equations, an area we explicitly mentioned in Santos Silva and Tenreyro (2006). Blackburn (2007) estimates wage equations in levels using several pseudo maximum likelihood estimators, including PPML. More recently, Petersen (2017) and Powell and Seabury (2018) also estimate equations for earnings by PPML.

The Cobb–Douglas production function is one of the best-known constant-elasticity models and therefore it is not surprisingly that one of the first uses of PPML outside of the trade literature involved the estimation of production functions. Building on Santos Silva and Tenreyro's (2006), who explicitly mentioned that this is a context in which PPML could be useful, Sun et al. (2011) advocated the estimation of production functions in levels and used PPML in their application. More recently, Dias and Marques (2021) showed that estimates of productivity dynamics based on firm-level data depend on whether logs or levels are used, and argue in favour of using data in levels when the analysis is based on weighted measures of productivity.

More generally, PPML has been employed to estimate models for durations (Abboud et al. 2016 and Call et al. 2018), investment in R&D (Cowan et al. 2015, and Guceri and Liu 2015), debt (Oksanen et al., 2015 and Lee and Mori 2021), losses and returns (Levieuge et al. (2021), and Paniagua et al. (2018)), value of mergers and acquisitions (Todtenhaupt et al. 2020), values of illicit drug sales (Nurmi et al. 2017), wind power capacity (Goetzke and Rave 2016), and to estimate models evaluating the effects of wild fires (Eskelson et al. 2016 and Peterson et al. 2019).

Finally, we note that PPML is also becoming important in the study of intergenerational income mobility. Mitnik and Grusky (2020) make a strong case for the use of PPML in the estimation of models of intergenerational mobility and show that its use makes a material difference; Helsø (2021) also uses PPML in this context.

7 Concluding remarks

The PPML estimator is extraordinarily well suited for the estimation of gravity equations. That was the point we made in Santos Silva and Tenreyro (2006) and, thanks to the follow-up work done by us and many others, that result is today even clearer and widely accepted. Indeed, in the vast majority of cases, there is no reason at all to consider alternative estimators for gravity equations because no other estimator shares all the attractive features of PPML that we discussed in Section 3.

Some years ago, the use of the PPML estimator could be challenging because of computational issues. Indeed, some authors even state that they do not report PPML estimates because of the computational challenges they faced. However, the introduction of the `ppmlhdfc` Stata command by Correia et al. (2020) made it very easy to estimate even complex gravity equations using very large panels. This command represents the state-of-the-art and essentially removed the final obstacles to the generalized use of the PPML estimator.

We often see papers that present results of the estimation of gravity equations using a potpourri of methods, and some authors go as far as recommend that practice. We do not see what can be gained by complementing the PPML estimates of gravity

equations with those obtained by methods that are almost certainly invalid, and suggest that it is better to spend research time making sure that the model is correctly specified and can be used to answer the question of interest.

We conclude with a small anecdote that the readers may find interesting. “The Log of Gravity” started when one of the authors serendipitously emailed the other asking for a copy of a ten-year old working paper (Santos Silva 1991); we did not know each other when we started to work on our paper, our collaboration was entirely done by fax and by email (which was challenging because of the different time zones), and we only met when the paper had already been accepted for publication. No matter how much we plan and how hard we work, luck will always play a big part in our lives and careers, and we have been more fortunate than most.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abbound ME, Band R, Jia J, Pajeroski W, David G, Guo M, Mechem CC, Messé S. R., Carr BG, Mullen MT (2016) Recognition of stroke by EMS is associated with improvement in emergency department quality measures. *Prehosp Emerg Care* 20:729–736
- Albert A, Anderson JA (1984) On the existence of maximum likelihood estimates in logistic models. *Biometrika* 71:1–10
- Anderson JE, Larch M, Yotov YV (2018) GEPPML: General equilibrium analysis with PPML. *The World Economy* 41:2750–2782
- Anderson JE, van Wincoop E (2003) Gravity with gravitas: a solution to the border puzzle. *Am Econ Rev* 93:170–192
- Baier SL, Bergstrand JH (2007) Do free trade agreements actually increase members’ international trade? *J Int Econ* 71:72–95
- Beine M, Bertoli S, Fernández-Huertas Moraga J (2016) A practitioners guide to gravity models of international migration. *The World Economy* 39:496–512
- Beine M, Parsons C (2015) Climatic factors as determinants of international migration. *Scand J Econ*: 723–767
- Bergé L (2018) Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm, DEM Discussion Paper Series 18-13, Department of Economics at the University of Luxembourg
- Blackburn ML (2007) Estimating wage differentials without logarithms. *Labour Econ* 14:73–98
- Bosquet C, Boulhol H (2014) Applying the GLM variance assumption to overcome the scale-dependence of the negative binomial QGPML estimator. *Econ Rev* 33:772–784
- Call AC, Martin GS, Sharp NY, Wilde JH (2018) Whistleblowers and outcomes of financial misrepresentation enforcement actions. *J Account Res* 56:123–171
- Cameron AC, Trivedi PK (2013) *Regression analysis of count data*, 2nd edn. Cambridge University Press, Cambridge
- Correia S, Guimarães P, Zylkin T (2019) Verifying the existence of maximum likelihood estimates for generalized linear models, arXiv:1903.01633
- Correia S, Guimarães P, Zylkin T (2020) Fast Poisson estimation with high-dimensional fixed effects. *STATA Journal* 20:95–115
- Cowan BW, Lee D, Shumway CR (2015) The induced innovation hypothesis and U.S. public agricultural research. *Am J Agric Econ* 97:727–742

- De Sousa J (2012) The currency union effect on trade is decreasing over time. *Econ Lett* 117:917–920
- Dias DA, Marques CR (2021) From micro to macro: a note on the analysis of aggregate productivity dynamics using firm-level data. *J Product Anal*, forthcoming
- Eaton J, Kortum S, Sotelo S (2013) International trade: linking micro and macro. In: Acemoglu D, Arelano M, Dekel E (eds) *Advances in economics and econometrics: tenth world congress*. Cambridge University Press, Cambridge, pp 329–370
- Egger PH, Staub KE (2015) GLM Estimation of trade gravity models with fixed effects. *Empir Econ* 50:137–175
- Egger PH, Tarlea P (2015) Multi-way clustering estimation of standard errors in gravity models. *Econ Lett* 134:144–147
- Eskelson BNI, Monleon VJ, Fried JS (2016) A 6 year longitudinal study of post-fire woody carbon dynamics in California's forests. *Can J Forest Res* 46:610–620
- Fally T (2015) Structural gravity and fixed effects. *J Int Econ* 97:76–85
- Fernández-Val I, Weidner M (2016) Individual and time effects in nonlinear panel models with large N , T . *J Econ* 192:291–312
- Fisher P (2016) British tax credit simplification, the intra-household distribution of income and family consumption. *Oxf Econ Pap* 68:444–464
- Flowerdew R, Aitkin M (1982) A method of fitting the gravity model based on the Poisson distribution. *J Reg Sci* 22:191–202
- Frankel J, Wei S (1993) Trade blocs and currency blocs, NBER Working Paper No. 4335
- Goetzke F, Rave T (2016) Exploring heterogeneous growth of wind energy across germany. *Util Policy* 41:193–205
- Goldberger A (1968) The interpretation and estimation of Cobb-Douglas functions. *Econometrica* 36:464–472
- Goldberger A (1991) *A course in econometrics*. Harvard University Press, Cambridge
- Gourieroux C, Monfort A, Trognon A (1984) Pseudo maximum likelihood methods: applications to Poisson models. *Econometrica* 52:701–720
- Guceri I, Liu L (2019) Effectiveness of fiscal incentives for R&D: quasi-experimental evidence. *Am Econ J: Economic Policy* 11:266–291
- Guimarães P, Portugal P (2010) A simple feasible procedure to fit models with high-dimensional fixed effects. *Stata Journal* 10:628–649
- Head K, Ries J (2008) FDI as an outcome of the market for corporate control: theory and evidence. *J Int Econ* 74:2–20
- Helpman E, Melitz M, Rubinstein Y (2008) Estimating trade flows: trading partners and trading volumes. *Q J Econ* 123:441–487
- Helsø A-L (2021) Intergenerational income mobility in Denmark and the United States. *Scand J Econ* 123:508–531
- Hinz J, Hudlet A, Wanner J (2019) Separating the wheat from the chaff: Fast estimation of GLMs with high-dimensional fixed effects. European University Institute, mimeo
- Hummel D (1999) Toward a geography of trade costs, GTAP Working Papers 1162, Center for Global trade analysis, department of agricultural economics, Purdue University
- Jeong Y, Siegel JI (2018) Threat of falling high status and corporate bribery: Evidence from the revealed accounting records of two South Korean presidents. *Strategic Management* 39:1083–1111
- Jochmans K (2017) Two-way models for gravity. *Rev Econ Stat* 99:478–485
- Jochmans K, Verardi V (2019) IVGRAVITY: Stata module containing method-of-moment IV estimators of exponential-regression models with two-way fixed effects from a cross-section of data on dyadic interactions and endogenous covariates, statistical software components S458698, Boston College Department of Economics
- Kaiser U, Mendez SJ, Rønne T, Ullrich H (2014) Regulation of pharmaceutical prices: evidence from a reference price reform in Denmark. *J Health Econ* 36:174–187
- Lancaster T (2000) The incidental parameter problem since 1948. *J Econ* 95:391–413
- Lee KO, Mori M (2021) Conspicuous consumption and household indebtedness. *Real Estate Economics*, forthcoming
- Levieuge G, Lucotte Y, Pradines-Jobet F (2021) The cost of banking crises: does the policy framework matter? *J Int Money Financ* 110:102–290
- Manning WG, Mullahy J (2001) Estimating log models: to transform or not to transform? *J Health Econ* 20:461–494
- Mitnik P, Grusky D (2020) The intergenerational elasticity of what? The case for redefining the workhorse measure of economic mobility. *Sociol Methodol* 50:47–95

- Mullahy J (1986) Specification and testing of some modified count data models. *J Econ* 33:341–365
- Mullahy J (1997) Instrumental variables estimation of Poisson regression models, applications to models of cigarette smoking behavior. *Rev Econ Stat* 79:586–593
- Nurmi J, Kaskela T, Perälä J, Oksanen A (2017) Seller's reputation and capacity on the illicit drug markets: 11-Month study on the finnish version of the silk road. *Drug Alcohol Depend* 178:201–207
- Oksanen A, Aaltonen M, Rantala K (2015) Social determinants of debt problems in a nordic welfare state: a finnish register-based study. *J Consum Policy* 38:229–246
- Paniagua J, Rivelles R, Sapena J (2018) Corporate governance and financial performance: the role of ownership and board structure. *J Bus Res* 89:229–234
- Papke LE, Wooldridge JM (1996) Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *J Appl Econ* 11:619–632
- Park R (1966) Estimation with heteroskedastic error terms. *Econometrica* 34:888
- Petersen T (2017) Multiplicative models for continuous dependent variables: estimation on unlogged versus logged form. *Sociol Methodol* 47:113–64
- Peterson KF, Eskelson BNI, Monleon VJ, Daniels LD (2019) Surface fuel loads following a coastal–transitional fire of unprecedented severity: boulder creek fire case study. *Can J Forest Res* 49:925–932
- Powell D, Seabury S (2018) Medical care spending and labor market outcomes: evidence from workers' compensation reforms. *Am Econ Rev* 108:2995–3027
- Ramsey JB (1969) Tests for specification errors in classical linear least squares regression analysis. *J Royal Stat Soc B* 31:350–371
- Reis HJ, Santos Silva JMC (2006) Hedonic prices indexes for new passenger cars in Portugal (1997–2001). *Econ Model* 23:890–908
- Santos Silva JMC (1991) Discriminating between the linear and log-linear forms of a regression model: optimal instrumental variables tests, University of Bristol, Department of Economics, Discussion Paper No. 91/301
- Santos Silva JMC, Tenreyro S (2006) The log of gravity. *Rev Econ Stat* 88:641–658
- Santos Silva JMC, Tenreyro S (2010) On the existence of the maximum likelihood estimates in Poisson regression. *Econ Lett* 107:310–312
- Santos Silva JMC, Tenreyro S (2011a) Further simulation evidence on the performance of the Poisson pseudo-maximum likelihood estimator. *Econ Lett* 112:220–222
- Santos Silva JMC, Tenreyro S (2011b) poisson: some convergence issues. *STATA Journal* 11:207–212
- Santos Silva JMC, Tenreyro S, Wei K (2014) Estimating the extensive margin of trade. *J Int Econ* 93:67–75
- Santos Silva JMC, Tenreyro S, Windmeijer F (2015) Testing competing models for non-negative data with many zeros. *J Econometric Methods* 4:29–46
- Stammann A (2018) Fast and feasible estimation of generalized linear models with high-dimensional k-way fixed effects, arXiv:1707.01815v3
- Sun K, Henderson DJ, Kumbhakar SC (2011) Biases in approximating log production. *J Appl Econ* 26:708–714
- Todtenhaupt M, Voget J, Feld LP, Ruf M, Schreiber U (2020) Taxing away M&A: Capital gains taxation and acquisition activity. *Eur Econ Rev* 128:103–505
- Verbeek A (1989) The compactification of generalized linear models. In: Decarli A, Francis BJ, Gilchrist R, Seeber G (eds) *Statistical modelling, proceedings of GLIM 89 and the 4th International workshop on statistical modeling*. Springer, New York, pp 314–327
- Verbeek A (1992) The compactification of generalized linear models. *Statistica Neerlandica* 46:107–142
- Wedderburn RWM (1976) On the existence and uniqueness of the maximum likelihood estimates for certain generalized linear models. *Biometrika* 63:27–32
- Weidner M, Zylkin T (2020) Bias and consistency in three-way gravity models, arXiv:1909.01327v5
- Windmeijer F, Santos Silva JMC (1997) Endogeneity in count data models: an application to demand for health care. *J Appl Econ* 12:281–294
- Wooldridge JM (1992) Some alternatives to the Box-Cox regression model. *Int Econ Rev* 33:935–955
- Wooldridge JM (1999) Distribution-free estimation of some nonlinear panel data models. *J Econ* 90:77–97
- Yotov YV, Piermantiini R, Monteiro JA, Larch M (2016) *An advanced guide to trade policy analysis: the structural gravity model*. Geneva (Switzerland): World Trade Organization