Does Output Influence Productivity? - A Meta-Regression Analysis*

Ludwig List
PhD Candidate
CEPN, Université Paris 13

October 6, 2017

Abstract

The goal of this paper is to conduct a meta-regression analysis (called MRA hereafter) regarding the effects of the 'Kaldor-Verdoorn effect' – the relation between output/demand and productivity. The Kaldor-Verdoorn effect was subject of many econometric studies and while the overwhelming majority finds a positive overall effect, there is no consensus on its size – the results are "all over the place".

MRA estimates a "true value' of the Kaldor-Verdoorn effect without interference from potential publication selection bias (also known as the "filedrawer problem") via the use of multivariate MRA. It is argued that "moderator variables" can be used to get a deeper understanding of the "political economy of publication", as there might be specific factors which influence the direction and overall strength of potential publication selection bias. A series of moderator variables is being used to check for their effect of excess variation, including amongst others the year of publication, academic journal, sources of funding, estimation method and the sectors and the countries studied.

This MRA study uses available data from 12 published studies with 118 estimations of the Kaldor-Verdoorn effect. While the literature regarding the Kaldor-Verdoorn effect covers well more than 120

*Many people helped me creating this study. First of all I have to thank Sophie-Dorothee Rotermund for help with the database creation. I also want to thank Dany Lang, Tom Stanley and Marc Lavoie for helpful feedback, as well as the participants of the EAEPE Summer School in Rome 2017 and the Meta-Regression Summer School in Crete 2017.

publications the biggest part of them which was published between the 1950s and 1970s fails to present key statistical variables and could thus not be incorporated into this work. A key reason for this is the missing of econometric standards in said time period and the absence of powerful hard- and software that enabled researchers to create standardised regressions.

1 Introduction

In 1949, P.J. Verdoorn published an article pointing at a potential link between the long-run rate of growth of labour productivity and the rate of growth of output, while indicating that the causation runs from the latter to the former. This finding was first referred to byArrow (1962) and later on coined “Verdoorn’s Law” by Kaldor (1966). Kaldor (1966) also presented the findings by Verdoorn (1949) as a ’stylised fact’ in order to explain the by then slow rate of growth of the UK economy – the reason for which Verdoorn’s law is also known as “Kaldor’s second law”.[1] [2].

The reactions to Verdoorn (1949) then split up into three interpretations of which two represent the currently most widely used ones.

The first interpretation of Verdoorn’s law emphasises the importance of “learning by doing” and the complementarity of capital and labour, which applies to both Verdoorn’s early interpretation and the work by Arrow (1962).

The second interpretation was put forward by Kaldor (1966) and stressed the macroeconomic impact of aggregate demand on growth. Here higher aggregate demand creates the need for rationalisation and more efficient production schemes, which in turn affects technical progress and productivity growth. This interpretation on microeconomic terms also emphasises the existence of increasing returns to scale, which challenged the at that time dominating hypothesis of constant returns to scale in neo-classical economic theory.[3]

[1] “Kaldor’s first law” characterises a positive relationship between the growth rate of manufacturing output and aggregate output.

[2] While Kaldor (1966) and his use of Verdoorn’s law created an intensive debate on its theoretical implications and empirical regularities, it is important to point out that Verdoorn himself never referred to his findings as a law-like relation. Indeed, in an exchange of letters with A.P.Thirlwall, Verdoorn explicitly forbid the publication by Thirlwall of an english translation of Verdoon’s work, stating that unlike at the time of the original’s publication he was now convinced that this relationship was only stable at the steady state. Verdoorn restated in 1980 that “(t)he“law”that has been given my name appears therefore to be much less generally valid than I was led to believe in 1949”

[3] Parts of this interpretation can be traced back up to Adam Smith’s and Allyn Young’s notions that the division of labour increases with the activity of the market.

The third interpretation used perfect substitutability of labour and capital via the use of a Cobb-Douglas production function. Neo-Classical theory in the 1950s and 60s had no complete growth model explaining technological change, instead assuming exogenous technological growth, full employment and steady state growth. Additionally the dissatisfaction with the idea of constant returns to scale, the new availability of international consistent data led the pursuit of different models. Finally the presence of very different growth patterns around the world were at odds with the idea of the theory of the tendency of growth rates to converge stated by the Solow-Swan model. Arrow (1962) provided a stepping stone which later inspired the creation of the endogenous growth literature (Romer (1988); Lucas, 1988) with increasing returns to scale and endogenous technological growth. This increased the need for a proper explanation concerning the relation between production and technology (and vice-versa) and gave rise to a series of empirical studies on the presence and overall effect of Verdoorn's law.

The validity of the Verdoorn effect would imply more than just the importance of aggregate demand in the development economic system. Its link to the potential existence of increasing returns to scale would lead to far-reaching conclusions for policy especially for developing countries and infrastructure decisions. Similarly in a country with sustained higher growth, the respective country would not necessarily lose in competitiveness since the increase in output would induce higher productivity growth. Growth itself might be reinforcing. On the other hand a sustained period of slow growth might put a country into a descending spiral of ever-deteriorating competitiveness.

Since Verdoorn's first publication, a lot of empirical studies have been published on the overall validity of his 'law'. The biggest part of those studies were published in the 1980s, with the 1990s showing less interest in the possible existence of the Verdoorn effect. Only in the recent decade there seems to be a renewed interest in this topic, which might have to do with the idea of a period of "secular stagnation". In such a scenario innovation would have to play a keyrole and thus the determinants of technological change become increasingly important again. An attempt at summarising the available econometric literature on the Verdoorn effect has been made by McCombie et al. (2002), who surveyed the literature dating from Verdoorn's first publication on this topic in 1949 until 2001. The authors conclude that the a one percentage increase in output (or demand, depending on the reader's interpretation) growth raises productivity growth between 0.3 and 0.6 percentage points. This relationship generally holds over aggregate and sector-wide, single country or regional studies, and different forms of estimations like for example cross-section estimations or time series. Similarly, Hein (2014) extends this survey by several recent studies and finds similar conclusions. Even though most of the available studies hint at the presence of such an effect as Verdoorn (1949) described, the size of the reported overall effect is "all over the place".

"*On the whole, the law appears to be largely substantiated in these studies, although, as is the case for most statistical economic relationships, the estimates sometimes need to be qualified. Indeed, in certain circumstances, the law still needs further work to solve a number of econometric problems. However, it is fair to say that Verdoorn's Law should be regarded as something more than just a 'stylised fact'*" (McCombie et al., 2002, p.1).

2 Bringing Order into the Chaos: Meta-Regression Analysis (MRA) as a quantitative literature survey

2.1 The Specification Problem

Meta-Regression Analysis builds upon a technique commonly known as 'meta analysis' in other fields like psychological and educational research, medicine and the social sciences (Stanley and Jarrell, 2005; Stanley and Doucouli agos, 2012). Meta analysis tries to summarise and integrate the existing empirical literature about a common parameter. As such it aims to present a systematic review of all scientific knowledge currently available and explain the given findings in all its vast variety in a comprehensive way.

Traditional literature surveys are often not able to present an all-encompassing survey of already existing studies, one of the obvious reasons being the word limit imposed by academic journals. But there is more to it. As Stanley and Jarrel (2005) argue,

> "*The reviewer often impressionistically chooses which studies to include in his review, what weights to attach to the results of these studies, how to interpret those results, and which factors are responsible for the differences among those results. Traditionally, economists have not formally adopted any systematic or objective policy for dealing with the critical issues which surround literature surveys As a result, reviews are rarely persuasive to those who do not already number among the converted.*"

Meta analysis can thus be of great help when it comes to examining a certain effect on which a lot of empirical studies have been published – it enables the researcher "to see the bigger picture". Additionally, MRA offers the tools to estimate the effect of different model specifications on the overall results and their significance. This way the researcher can distinguish true economic effects from disturbances caused by wrong model specification more easily. Another reason for the use of MRA is the idea of the "filedrawer problem". As the standard deviation of the estimated correlations are becoming smaller with an increase in the number of observations, studies using a dataset with a comparatively small amount of observations face higher difficulties to obtain significant results. This might become important insofar as that peer-reviewed journals may prefer publishing only studies that offer significant results, even though from a methodological point of view the publication of not significant results would be equally important for the progress of economics as a science. In the worst case, such strict publication policies might incentivise researchers to alter their estimation model successively until significant results have been obtained ("publication selection bias").

Meta analysis has been successively used in the economics literature during the past decades. The most commonly quoted studies include Rose and Stanley (2005) on the effect of common currencies on international trade, Doucouliagos (2005) on the link between freedom and economic growth, Groot and Florax (2005) on the existence of world-wide growth rate convergence, Nijkamp and Poort (2005) on the unemployment elasticity of wages, Weichselbaumer and Winter-Ebmer (2005) on the gender wage gap and Knell and Stix (2005) on the income elasticity of money demand. While MRA became a well-accepted approach in other scientific fields, its appearance in economics is as of yet a comparatively rare sight. Nevertheless a guideline for a more standardised use of MRA in economics has been proposed by the "Meta-Analysis of Economics Research Network" (MAER) in order to improve both the transparency and the quality of future Meta analysis.

2.2 Meta-Regression Analysis

The basic idea of MRA is following several steps. First, the researcher is collecting all available studies on a specific effect he or she wants to study. Whether those studies are published in peer-reviewed journals or not should a priori not play any role. Indeed, one is even encouraged to include non-published studies, as those studies not being published might not necessarily indicate unscientific methods or a lower quality with respect to the used methods, but rather point at potential publication selection bias, as was explained in the former section.

In a second step, the estimated results in the studies gathered are being treated as individual entries in a new dataset. For a study to be able to be included in the dataset, the researcher has to be able to infer at least the estimate, the corresponding t-value (or its standard deviation) and the number of observations in the used dataset (or the degrees of freedom). If those three variables can be obtained then the estimate can be included in the dataset. Similarly the researcher might be interested in adding several characteristic elements of the specific study that might be worth considering such as for example the sources of the datasets being used, the year of publication (if published), the method of estimation, the country or sector examined *etc.* This possibility is also one of the big advantages of MRA. Not only is it possible to infer a more precise estimate for any given variable but MRA enables the researcher to find out which socio-economic circumstances might skew the estimated results and lead to under- or overestimation of the effect in question. Finding and explaining this differences via MRA is based on statistical, not economic theory and can thus try to shed light into topics of constant discussion between different schools of thought.

The third step then consists of a two-step regression in which the first regression points at the presence or absence of publication selection bias – which in the existing MRA literature has almost always been found – while the second regression tries to estimate this very publication selection bias and the "true value" of the parameter in question.

3 A Meta-Regression Analysis on the Kaldor-Verdoorn Effect

3.1 The basic model

Following Stanley and Jarrell (1989), the most common approach to do meta-analysis in economics consists in using effect sizes in reported econometric studies. The following section builds mostly on Stanley and Doucouliagos (2012) as well as the guidelines published by the MAER network (Stanley et al., 2013). The notation is orienting itself on Paldam (2015).

In order for being able to do MRA, studies that estimate the Kaldor-Verdoorn effect are being collected only when they meet two conditions. First, the studies collected must be estimating comparable effects (Becker and Wu (2007)). In order to make them comparable MRA studies are using so-called effect sizes , secondly, the studies are transparent in that the researcher is able to gather at least the estimated coefficient, its corresponding standard deviation or t-value and the number of observations used in the study the results were taken from.

MRA uses the relation between an estimates' t-value and its 'precision' to draw its conclusions. Consider a sample of perfectly homogeneneous estimated studies with an average β_0.

$$b_i = \beta_0 + \beta_1 se_i + u_i \tag{1}$$

In this case the reported estimats should all be randomly distributed around the 'true value', β_0 which should correspond to the average of the perfectly homogeneneous estimated studies.

The idea of publication selection assumes that researchers with a smaller sample and thus higher standard errors are forced to search longer for statistically significant results than their colleagues with bigger samples who will be statisfied with their potentially smaller, but significant estimates (for example via searching for additional data or for reasons to eliminate 'potential outliers'). Hence in the case of publication selection the estimate will be positively correlated with the standard error se_i while if the null hypothesis is true and no publication selection bias is present the estimate will be equal to the true value ($b_i = \beta_0 + u_i$ if $se_i \to 0$).

Clearly it is usually the case that studies will have varying sampling errors, se_i, i.e. there will be heterogeneity which will have to be corrected for. Dividing equation 1 by se_i will give us a WLS version which is in fact a basic MRA of the estimate's t-value against its precision, $p = \frac{1}{se_i}$. In case of homogeneity, the former error divided by the measured sampling error must be equal to 1.

$$t_i = \beta_M p + \beta_F + v_i \tag{2}$$

t_i refers to the estimate's t-value and β_M is the 'meta-average' with p being the 'precision score' $\frac{1}{se_i}$. Equation 2 can equally be rewritten as $t_i = \beta_M \frac{1}{se_i} + \beta_F + v_i$. Both parts of equation 2 are being used for testing. Testing β_F for the null hypothesis that $\beta_F = 0$ is called the 'funnel-asymmetry test' (FAT-test) and checks for heterogeneity. A rejection of the null hypothesis points at the existence of publication selection bias. β_M represents the 'precision estimate test' (PET-test) and is used to estimate the meta-average in case of publication selection bias. Again both tests rely on the assumption that the reported estimates are normally distributed around the underlying effect. Thus estimates far away from the underlying effect should have low precision while estimates closer to the 'true effect' should have high precision.

The FAT-PET test thus enables MRA to not only find out about the possible existence of publication selection bias, but to also correct the ordinary average of the reported estimates to be corrected for the estimated publication selection bias and to get a 'cleaner' estimate closer to the overall effect.

3.2 The Dataset

The dataset was created completely anew and will be published in the meta-data repository shortly after this study was accepted for publication. It consists of 12 published studies starting in 1975 and ending in 2016 with a total amount of 118 estimates. The platforms used for finding those studies cover the biggest array possible in order to account for as many studies as possible and consist of Econlit, JStore, Google Scholar and Google Search. Keywords for the searching process were 'Verdoorn effect' and similar terms such as 'Kaldor-Verdoorn effect' and 'Kaldor's second law', 'productivity', 'productivity-growth nexus'.

The following step included extracting t-values, standard deviations, the estimated coefficient, the number of observations and/or the degrees of freedom from every specific regression for every single study. Additionally several other variables of possible interest were recorded as dummies for further analysis. It has to be made clear that treating every single regression as a single entry might introduce some bias, as papers with several control estimation will be overrepresented. How to tackle this issue is not agreed upon in the economics MRA community. In this case no weights were being attached to the individual estimates since doing so might introduce an even bigger bias than the one it is supposed to prevent.

Another problem consists in the fact that most of the studies found were not reporting all the key variables needed. For an MRA to be done the database needs to contain at least any two of the estimate, the corresponding standard error or its t-value (since the t-value is the ratio of the estimate divided by the standard error, assuming that the null hypothesis is $b = 0$). For robustness checks other variables like the number of observed variables/the degrees of freedom should be included, since they can be used to compute partial correlations which are another effect size apart from the t-value commonly used in MRA. This proved more difficult than expected and as a result of the over 120 studies present at the beginning of this study only 12 studies were in an acceptable state for inclusion into the database. This might be explained by their date of publication. Econometric standards were not as etablished in economics in the 1970s as they are today. Similarly without the use of powerful statistical programs as we have available today doing econometric estimations was far more challenging. One might simply be less motivated to calculate important test statistics if one has to do so by hand. This naturally diminishes the explanatory power of this MRA since less studies can be investigated upon. On the other side an econometric study without such basic statistic measures by today's standards would not be able to get accepted for publication in peer-reviewed journals.

4 A Meta-Regression Analysis on the Kaldor-Verdoorn Effect

Since the variance of the estimated effect and hence u will vary from reported estimate to estimate estimation model 1 has to deal with heteroskedasticity, which is why weighted least square (WLS) estimation as in equation 2 is being used. As every estimate gets weighted with its corresponding standard error, estimates with large standard errors are given smaller weight while more precise estimators are given more weight. This makes WLS-MRA more resistent to outliers. Before the presentation of the MRA results one implausible estimate by Ghosh and Mizuno (1985) was removed. The removal of this enfluential was necessary to prevent complete domination of one single entry in the database and was justified using the DFBETA statistic ($\mid DFBETA \mid= 10.62$).

Figure 1 presents the combined 117 observations from 12 studies on the Kaldor-Verdoorn effect with year of publication ranging from 1975 to 1999 in a scatter plot called the funnel plot. It plots the estimates against their precision (the inverse of their corresponding standard error). In case of homogeneity the reported estimates should be normally distributed around the underlying effect, with estimates decreasing in precision the farther away they are moving from it. A skewed distribution of reported estimates hints at the possible existence of publication selection bias. The funnel plot, even though only being descriptive in nature is still useful for an initial overview of the existing literature. The average estimated Kaldor-Verdoorn effect amounts to 0.78 and on a first look the reported estimates seems to be rather symmetric. The estimates get increasingly more precise the closer they get to the average while losing precision with more distance.

Now this is of course only a descriptive analysis. For more a more thorough investigation more detailed tools are needed. MRA uses a regression between a reported estimate and its standard error as a more objective method of finding and measuring potential publication selection bias. In the absence of publication selection bias there should be no significant correlation between the estimate and its standard error while the opposite would be true in case of publication selection bias.

Table 1 shows the regression results for the FAT-PET test. The constant term gives insight about the nature of β_F - is there a sign of publication selection bias? Here the null hypothesis cannot be rejected at any of the significance levels commonly used in economics - thus the result of the FAT test supports the impression given by the funnel plot - the reported estimates seem to be close to the underlying economic effect, without any obvious over-

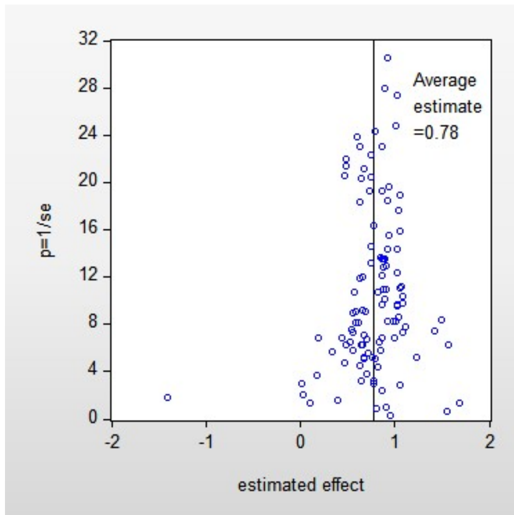

Figure 1: Funnel Plot with average estimated Kaldor-Verdoorn effect

Dependent Variable: TSTATISTIC
Included observations: 116

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	0.485347	0.814988	0.595527	0.5527
P	0.415415	0.062200	6.678712	0.0000

R-squared	0.281234	Mean dependent var	5.090832
Adjusted R-squared	0.274929	S.D. dependent var	5.494291
S.E. of regression	4.678449	Akaike info criterion	5.940901
Sum squared resid	2495.219	Schwarz criterion	5.988377
Log likelihood	−342.5723	Hannan-Quinn criter.	5.960174
F-statistic	44.60520	Durbin-Watson stat	0.544655
Prob(F-statistic)	0.000000		

Table 1: FAT-PET Test

or underestimation.

However it has to be taken into account that the FAT-test is seen as a relatively weak test by the MRA-community. Thus even though the null hypothesis could not be rejected the PET-test should nevertheless be used. The FAT-test might not be able to include all possible forms of publication selection bias and in the case of absence the meta-average which comes as a result of the PET-test should be very close to the unweighted average.

The PET-test can thus be seen in the second column and represents the meta-average, corrected for potential publication selection bias even if there was none reported via the FAT-test. This time the null hypothesis of $\beta_M = 0$ is being rejected at the highest commonly used significance level, the estimated Kaldor-Verdoorn effect is around 0.42, vastly lower than the average reported estimate would suggest.

But do these results still hold when controlled for potentially omitted variables? A further advantage of MRA is the use of dummy variables (called 'moderator variables from now on) to account the impact of omitted variables and their impact on publication selection bias. Especially in economics, where researchers are often forced to work with pre-compiled data sets and important variables might not be able to get taken into account omitted variable bias might indeed be one of the biggest drivers of misspecification. Similarly, other study-specific properties might be of interest to the researcher if a good summary of the respective literature is desired. Other often used variables include the proxies used to represent the respective effect, nature and origin of the dataset and estimation technique used in the paper as well as year and journal of potential publication, sources of funding, etc. Moderator variables might indeed be one of the key advantages of MRA as it allows to get a more thorough understanding of the underlying mechanism of scientific publication.

It might even retrieve more information from a set of studies than the studies themselves were intended to give. MRA can check for potential differences in years of publication, where certain theoretical explanations might have gained a lot of momentum and thus researchers were compelled to publish results more favourable to that idea - or at least, consider this idea more prominently in their econometric model than other less prominent ones. Another way to take this into account would be to introduce a moderator variable for every study citing a specific publication and so on. Curiosity really is the only frontier here.

Table 2 shows the augmented FAT-PET test and checks for estimates focussing on OECD countries ($doecd$), the manufacturing sector ($dmanuf$) and whether an estimate included the capital stock as explanatory variable ($dcap$). All of the added moderator variables are significant, while their respective impact on the reported estimates varies strongly.

Dependent Variable: TSTATISTIC
Included observations: 98 after adjustments

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	0.340107	0.366462	0.928082	0.3558
P	0.737366	0.031266	23.58372	0.0000
DOECD	−0.678915	0.030886	−21.98131	0.0000
DMANUF	0.513804	0.031290	16.42092	0.0000
DCAP	0.241540	0.036337	6.647219	0.0000
R-squared	0.917106	Mean dependent var		5.651130
Adjusted R-squared	0.913541	S.D. dependent var		5.745916
S.E. of regression	1.689528	Akaike info criterion		3.936448
Sum squared resid	265.4688	Schwarz criterion		4.068334
Log likelihood	−187.8859	Hannan-Quinn criter.		3.989793
F-statistic	257.2286	Durbin-Watson stat		0.998710
Prob(F-statistic)	0.000000			

Table 2: Augmented FAT-PET Test

While the addition of the stock of capital in the estimation regression tends to augment the overall Kaldor-Verdoorn effect by 0.24 percentage points (0.51 percentage points in the case of estimates using manufacturing sector data), countries studying the Kaldor-Verdoorn effect on OECD countries seem to find much lower estimates than elsewhere. Nevertheless the findings of the simple FAT-test stay - no sign of publication selection bias is apparent. Furthermore, after the inclusion of our moderator variables the overall estimated meta-average β_M indicates that the literature on the Kaldor-Verdoorn effect seems to have been quite precise. While the average estimate of all reported studies points is 0.78, the result of the PET-test suggests that the Kaldor-Verdoorn effect, corrected for potential publication selection bias is around 0.74 - very close to the consensus of the reviewed literature.

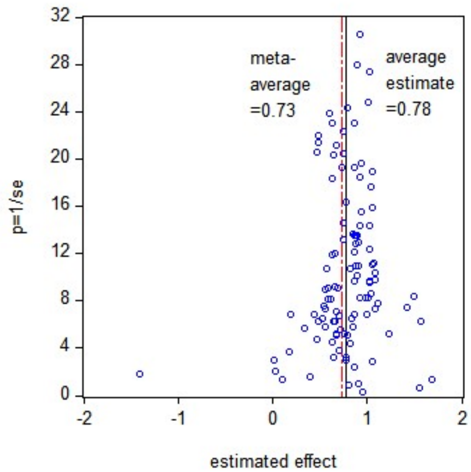

Figure 2: Average and Meta-Average

5 Conclusion

This study tried to summarise the plethora of empirical estimates surrounding the Kaldor-Verdoorn effect, the relation between growth in aggregate output/demand and productivity growth. Conventional literature reviews are very limited in both the amount of pages dedicated to such an enterprise as well as the level of detail they can delve into. Using a method known as Meta-regression analysis (MRA) the estimates of 13 studies along with several study-specific properties were accumulated and analysed by the FAT-PET test using weighted least square (WLS) regression.

Two main findings can be drawn from this study. First, the studies analysed did not indicate the presence of publication selection bias. This is unusual as nearly all MRA studies find existing publication selection bias in the studied field, although the strengh of the bias might be quite different every time. It has to be pointed out however that theoretically there might still be publication selection bias which could not be measured by the FAT-test.

Secondly, the meta-average the average estimate controlled for publication selection bias with a coefficient of 0.73 is very close to the descriptive average before the MRA of 0.78. This indicates that publication selection bias might indeed not be existent in the reviewed literature. This might have important implications for many different fields of research, such as economic development, trade and growth theory. It is also a surprising finding to many economists who would assume the Kaldor-Verdoorn effect to be weaker.

References

Kenneth J. Arrow. The Economic Implications of Learning by Doing. *The Review of Economic Studies*, 29(3):155, June 1962. ISSN 00346527. doi: 10.2307/2295952.

Betsy Jane Becker and Meng-Jia Wu. The Synthesis of Regression Slopes in Meta-Analysis. *Statistical Science*, 22(3):414–429, 2007. ISSN 0883-4237.

Chris Doucouliagos. Publication Bias in the Economic Freedom and Economic Growth Literature. *Journal of Economic Surveys*, 19(3):367–387, July 2005. ISSN 1467-6419. doi: 10.1111/j.0950-0804.2005.00252.x.

Nicholas Kaldor. *Causes of the Slow Rate of Economic Growth of the United Kingdom: An Inaugural Lecture*. Cambridge University Press, 1966. Google-Books-ID: Mi87AAAAMAAJ.

John McCombie, Maurizio Pugno, and Bruno Soro, editors. *Productivity Growth and Economic Performance*. Palgrave Macmillan UK, London, 2002. ISBN 978-1-349-42876-2 978-0-230-50423-3. doi: 10.1057/9780230504233.

Martin Paldam. Meta-Analysis in a Nutshell: Techniques and General Findings. *Economics: The Open-Access, Open-Assessment E-Journal*, 2015. ISSN 1864-6042. doi: 10.5018/economics-ejournal.ja.2015-11.

P. M. Romer. Capital Accumulation In The Theory Of Long Run Growth. RCER Working Paper 123, University of Rochester - Center for Economic Research (RCER), 1988.

Andrew K. Rose and T. D. Stanley. A Meta-Analysis of the Effect of Common Currencies on International Trade*. *Journal of Economic Surveys*, 19(3):347–365, July 2005. ISSN 1467-6419. doi: 10.1111/j.0950-0804.2005.00251.x.

Doris Weichselbaumer and Rudolf Winter-Ebmer. A Meta-Analysis of the International Gender Wage Gap. *Journal of Economic Surveys*, 19(3):479–511, July 2005. ISSN 1467-6419. doi: 10.1111/j.0950-0804.2005.00256.x.