Université de Montréal

Essays in Innovation, Inequality and Risk

par

Marlène KOFFI

Département de sciences économiques
Faculté des arts et des sciences

Thèse présentée en vue de l'obtention du grade de
Philosophiæ Doctor (Ph.D.)
en Discipline

June 2020

Université de Montréal

Faculté des arts et des sciences

Cette thèse intitulée

Essays in Innovation, Inequality and Risk

présentée par

Marlène KOFFI

a été évaluée par un jury composé des personnes suivantes :

Joshua Lewis
(président-rapporteur)

Panousi Vasia
(directeur de recherche)

Baris Kaymak
(membre du jury)

Marti Mestieri
(examinateur externe)

(représentant du doyen de la FESP)

A mon Père, A ma Mère, A mon Frère ...

Remerciement

Je tiens tout d'abord à exprimer ma profonde gratitude envers ma directrice de recherche Vasia Panousi, qui a accepté de m'encadrer et me guider dans cette thèse de doctorat. Je lui suis infiniment reconnaissante pour avoir eu confiance en moi, pour avoir cru en moi, pour son inlassable disponibilité, sa gentillesse, ses conseils, son soutien et sa rigueur. Je suis très heureuse de l'avoir plutôt surprise à son arrivée en lui demandant d'être mon encadreur. Merci encore pour ces collaborations et ces échanges sur ces sujets de recherches passionnants.

Je remercie aussi Marti Mestieri et Dimitris Papanikolaou pour leur soutien, leur confiance et nos différentes discussions au travers desquelles j'ai pu bénéficier de leurs conseils et de leur expérience.

J'adresse aussi mes remerciements à Joshua Lewis pour sa disponibilité et sa générosité, à Christopher Rauh, Immo Schott, Baris Kaymak et Marine Carrasco pour leurs conseils. Je suis aussi reconnaissante envers tous les professeurs du département de sciences économiques de l'Université de Montréal. Je remercie aussi le Centre Inter-universitaire de Recherche en Economie Quantitative (CIREQ) et la Faculté des études supérieures et postdoctorales pour leur soutien à la recherche estudiantine.

Mes sincères remerciements aux doctorants en sciences économiques de l'Université de Montréal, particulièrement à Magnim Farouh pour son soutien infaillible, à Ismael Mourifié, à tous mes camarades de promotion N'Golo Koné, Fatim Diabagaté, Lucienne Talba, Lionel Siwe, Alex Tieumeuna, Samuel Gingras et Idriss Tsafack. Merci d'avoir meublé ses années par votre présence!

Je ne peux terminer sans exprimer toute ma gratitude aux membres de ma famille sans qui je ne n'aurais pu tenir ses années. Ils m'ont toujours témoignée leur amour, leur soutien et ont toujours cru en moi et pour moi.

Résumé

Cette thèse s'articule autour de trois chapitres en économie de l'innovation et de la science. Pour ce faire, elle développe des modèles empiriques et théoriques pour analyser l'innovation technologique et scientifique et produire des recommandations politiques. Le premier chapitre utilise l'apprentissage automatique et les sciences de données pour construire un indicateur de l'innovation technologique. À l'aide d'une base de données unique sur les brevets au Canada, nous construisons un indice de qualité des brevets pour répondre à deux questions principales : l'absence d'une base de données systématique sur les brevets et leur valeur au Canada ainsi que l'évaluation du secteur pharmaceutique, l'un des principaux secteurs leaders de l'innovation au Canada . Les résultats révèlent que notre indice de qualité est lié à la performance économique des entreprises, à leur productivité et à la productivité agrégée. Le deuxième chapitre examine les innovations dans la recherche universitaire. Plus précisément, se focalisant sur les sciences économiques, ce chapitre vise à relier l'innovation et les inégalités en analysant la reconnaissance des idées des femmes. Des données bibliométriques issues de la recherche en économie sont utilisées pour étudier les biais de genre dans les citations. Sur la base des techniques d'apprentissage profond, on peut (1) établir les similitudes entre les articles (2) établir un lien entre les articles en identifiant les articles qui citent, les articles cités et les articles qui devraient être cités. Cette étude révèle qu'en moyenne, les articles qui ne sont pas cités sont 20% plus susceptibles d'être écrits par des femmes que par des hommes. Ce biais d'omission est plus répandu lorsqu'il n'y a que des hommes dans l'article citant. Dans l'ensemble, pour avoir le même niveau de citation que les articles rédigés par des hommes, les articles rédigés par des femmes doivent être supérieurs de 20 percentiles dans la distribution du degré d'innovation de l'article. Enfin, le dernier chapitre analyse l'innovation dans une perspective plus macroéconomique, en se concentrant sur les entrepreneurs. En effet, les entrepreneurs sont au cœur du développement économique et de l'innovation. Cependant, l'activité entrepreneurial reste très risquée. Quelles sont donc les opportunités de diversification des risques d'investissement pour les entrepreneurs ? Pour répondre à cette question, nous étudions le rôle de l'intégration financière. Avec un modèle théorique en temps continu et avec des agents hétérogènes, nous montrons que l'ouverture financière produit des gains de bien-être substantiels pour les entrepreneurs et peut donc les aider à diversifier le risque d'investissement. Nos résultats sont également étayés par une analyse empirique.

Mots-clés: Innovation, Fluctuations économiques et Croissance, Entrepreneuriat, Inégalité, Inégalité entre les sexes, Recherche universitaire.

Abstract

This thesis is organized into three chapters in the economics of innovation and science. In doing so, it develops empirical and theoretical models to analyze technological and scientific innovation and produce policy recommendations. The first chapter uses data science and big data techniques to build an indicator of technological innovation. Using a unique database on patents in Canada, we build a patent quality index to answer two main questions: the absence of a systematic database on patents and their value in Canada and the evaluation of the pharmaceutical sector, one of the leading innovating sectors in Canada. The results reveal that our quality index is linked to the economic performance of firms, their productivity, and aggregate productivity. The second chapter looks at innovations in academic research. Specifically, focusing on economics, this chapter aims to connect innovation and inequality by analyzing the recognition of women's ideas in the field. Bibliometric data from research in economics are used to investigate gender biases in citation patterns. Based on deep learning and machine learning techniques, one can (1) establish the similarities between papers (2) build a link between articles by identifying the papers citing, cited and that should be cited. This study finds that, on average, omitted papers are 20% more likely to be female-authored than male-authored. This omission bias is more prevalent when there are only males in the citing paper. Overall, to have the same level of citation as papers written by males, papers written by females need to be 20 percentiles upper in the distribution of the degree of innovativeness of the paper. Finally, the last chapter analyzes innovation from a more macroeconomic perspective, focusing on entrepreneurs. Indeed, entrepreneurs are at the core of economic development and innovation. However, entrepreneurship remains very risky. What are the opportunities for investment risk diversification for entrepreneurs? To answer this question, we investigate the role of financial integration. With a theoretical model featuring a continuous-time dimension with heterogeneous agents, we show that financial openness produces substantial welfare gains for entrepreneurs and therefore can help its agents to diversify the investment risk. Our results are also supported by empirical analysis.

Keywords: Innovation, Economic Fluctuations and Growth, Entrepreneurship, Inequality, Gender Inequality, Academic Research.

Contents

List of tables

List of figures

Introduction

Since Shumpeter (1911/1934) via Romer (1986), it is more and more consensual to perceive scientific and technological innovations as a means of ensuring sustained growth and improving standards of livings. However, this role of innovation as a pioneer of economic growth faces two major difficulties. On the one hand, even if there is a lot of theoretical economic models describing the process of innovation, it remains very difficult to find an empirical correspondence to these models and therefore to really measure innovation. On the other hand, as the idea of creative destruction already mentioned, innovation is not always followed by a distribution of gains for economic agents, and this is all the more true when it comes to new ideas (Kogan et al. (2020)). Thus, to better capture the benefits of innovation, it is first necessary to be able to acquire an effective measure of innovation, then to understand the unequal dimension that innovation can take and finally to evaluate the possibilities of diversification of potential risk that could be linked to innovation. This thesis aims to contribute to these three dimensions by analyzing innovation from an empirical and microeconomic point of view, but also to open up perspectives by adding a more theoretical and macroeconomic dimension.

Specifically, Chapter 1 uses a new panel dataset constructed from information provided by the Canadian Intellectual Property Office to study the relationship between patents, innovation and growth in the Canadian pharmaceutical industry. First, because Canadian legislation does not require citing previous significant patents, we perform textual analysis on patent documents to create an indicator of patent quality or significance. Our indicator assigns higher quality to patents or innovations that are novel, i.e. different from the prior knowledge stock, and influential, i.e. related to subsequent patents. Second, we match the firms in our patent dataset to their balance-sheet information from Compustat North America. We are then able to validate our patent-quality measure by relating it to various measures of firm value and performance. The results indicate that the anticipation of the granting of a breakthrough patent increases firm profitability, on average, for up to five years before the grant. This increase in profitability is reflected in increased markups, as opposed to increased employment or investment. Third, we construct firm- and aggregate-level TFP measures. We show that our patent-quality indicator is positively related, statistically and economically, with firm productivity. The conclusion from this analysis is that significant patents or innovations increase firm productivity as captured by measured TFP.

Moving forward, in chapter 2 investigates a more abstract concept of innovation in ideas. In fact, it analyzes the recognition of women's innovative ideas. Bibliometric data from research in economics are used to investigate gender biases in citation patterns. Based on deep learning and machine learning techniques, one can (1) establish the similarities between papers (2) build a link between articles by identifying the papers citing, cited and that should be cited. This study finds that, on average, a paper omits almost half of related prior papers. There are, however, substantial heterogeneities among the authors. In fact, omitted papers are 15% to 30% more likely to be female-authored than male-authored. First, the most likely to be omitted are papers written by women (solo, mostly female team) working at mid-tier institutions, publishing in non-top journals. In a group of related papers, the probability of omission of those papers increases by 6 percentage points compared to men in similar affiliation when the citing authors are only males. Overall, for similar papers, having at least one female author reduces the probability of omitting other women's papers by up to 10 percentage points, whereas having only male authors increases the probability of being omitted by almost 4 percentage points. Second, the omission bias is twice as high in theoretical fields that involve mathematical economics than it is in applied fields such as education and health economics. Third, men benefit twice as much as women from publishing in a top journal, in terms of likelihood of being omitted. Lastly, being omitted with respect to past publications affects future productivity and reduces the probability of getting published in a top journal. Finally, peer effects and more editorial board diversity tend to counteract and reduce the omission bias.

Finally, Chapter 3 investigates the possibility of diversification of the risk related to investment. It focuses on the diversification of entrepreneurial risk for financially integrated economies. In doing so, it uses a continuous time, general-equilibrium model with heterogeneous agents facing a time-varying idiosyncratic investment risk (uncertainty shock). First, by contrast to model with no time-varying risk, this novel framework gives the implications of idiosyncratic risk for business cycle fluctuations and talks about stabilization policy. Second, in a similar model with only an aggregate risk, the cost of capital flows outweighs the gain from risk-sharing. Countries do not gain a lot from financial integration even considering extreme values of risk aversion. At the opposite, in the presence of a time-varying idiosyncratic risk, the results get reversed and sizeable welfare gains emerged. Three key mechanisms help in getting those results: a price effect, a reallocation effect, and a wealth effect. Agents in a country hit by a bad shock are less willing to invest and reallocate their portfolio in favour of the less risky asset. This avoids a substantial drop in the aggregate price of capital compared to the autarchy situation. Therefore, their balance sheets are less proportionally hit by the shock and they can recover using their savings in the least risky asset. The overall economy becomes less volatile. The welfare gains from financial globalization are higher and could go above 10%, depending on the systematic risk of the country. The chapter ends by highlighting some empirical support of the main theoretical channel described. Therefore, the results also call for more cautiousness from policy makers in attempting to limit capital movement without taking into account heterogeneity at the individual production side.

Chapter 1

Patents, Innovation and Growth in Canadian Pharmaceuticals

[1]

1.1. Introduction

Context. Scientific and technological innovations are crucial for long-term sustainable economic growth. Patents are a frequently-used measure of such innovations. Recently there has been a resurgence of academic research interest on the importance of patents for growth, however most of these studies have focused almost entirely on US patents. By contrast, there has been no systematic analysis of patents granted by Canadian governments, either across provinces or across time. This is despite a number of concerns that have been raised about the state of innovation in Canada. For example, the consensus among researchers and policy-makers is that Canada is lagging other members of the Organization for Economic Cooperation and Development (OECD) in the rate of growth of patent registrations. There are also concerns about the share of Canadian patents that are held and exploited abroad. Furthermore, the Canadian patent process exhibits much larger delays, compared to that in the US. Specifically, the average waiting time between patent application and patent granting is eight years in Canada, compared to two years in the US. For example, in 1921, Canadian medical scientists Frederick Banting and Charles Best discovered insulin, for which they received the Nobel Prize in medicine in 1923. However, due to process delays in Canada, they filed for a patent with the US instead, where they were approved in 1923. In the same year, they sold the patent rights to the University of Toronto for $1 each. In 1996, James Gosling, the Canadian inventor of the programming language Java, filed for a Canadian patent but only received a request for its examination by the Canadian Intellectual Property Office (CIPO) in 2000. In 2003, he was granted a patent by the US. By 2005, he had not

[1]This chapter is coauthored with Vasia Panousi (Université de Montréal).

received news about the final outcome of the Canadian examination, so he abandoned the process. Nonetheless, Canada appears to maintain an innovation advantage in the areas of pharmaceuticals and medical devices.

Contribution. In this paper, we investigate the relationship between patents, innovation and growth in the Canadian pharmaceutical industry. We start by constructing a novel longitudinal database on Canadian pharmaceutical patents since 1869, using data obtained after communication with the Canadian Intellectual Property Office (CIPO). The database contains information on all pharmaceutical patents granted by Canadian governments since 1969, such as the name and address of the applicant company, the names and nationalities of the inventors, detailed textual descriptions of each patent, patent claims, dates of filing and granting of each patent, and industry classification codes. Using the raw data, we also construct variables that track the ownership of patents across time, firms, and inventors. We present descriptive statistics and stylized facts about our data. In that sense, this paper provides information about Canadian patents in a similar fashion to that of the NBER database available for patents granted by the USPTO, and is the first systematic longitudinal study of patents across Canada. .

Second, we perform textual analysis using big-data and machine-learning methodologies in order to construct new measures of patent similarity and patent quality. This is especially important because in Canada, contrary to the US, there are no patent citation counts. This is because the Canadian legislature does not require inventors to cite prior related art, whereas US legislation does. As a result, patents filed in Canada very rarely contain references to previous patents upon which they build. Furthermore, the CIPO does not appear to be relying on any other systematic measures of patent quality across time, provinces or industries. Given that 90 percent of patents are improvements on prior innovations, this is a substantial problem. In fact, it could explain the differences in the time lags between filing and granting in Canada, compared to the US: in Canada, it takes eight years vs two years in the US. Our patent-quality indicator, constructed from comparison of textual patent documents, assigns higher quality to patents that are novel and influential, i.e. different from the prior stock of knowledge but related to subsequent patents.

Third, we validate our measure of patent quality via case-study analysis and via examination of the relationship with similar-patent citations. To begin with, the patents at the top 1 percent of our quality measure distribution appear to be patents related to the first steps in the process of significant medical discoveries. For example, patents that initiate substantial improvements in the cure of HIV/AIDS or in the cure of various types of cancer. By contrast, patents at the bottom 1 percent of the quality distribution appear to be either extremely similar to previous patents or contribute to the minor alleviation of lesser problems, such as indigestion or stomach aches. Furthermore, we design a web-crawling algorithm that identifies similar patents granted in the US, for which Google patents provides the number of citations. We then examine the relationship between our Canadian patent-quality measure and the US patent-citation measure for similar patents. We find that this relationship is positive and significant, statistically and economically. These results indicate that our measure does in fact identify significant patents and innovations and that it could be used by Canadian officials to potentially expedite the patent examination and granting process.

Fourth, we construct a new bridged database, connecting the pharmaceutical firms from our patent database to their balance-sheet and stock-price information from Compustat and CRSP North America. In a novel data-collection endeavor, compared to the few other existing studies of pharmaceutical industries, we take into account mergers and acquisitions (M&A) across companies. We use a combination of hand-collection and of automated web-crawling techniques to gather data on M&A from a variety of sources, including newspapers, company websites, and Bloomberg libraries. We also adjust for changes in company names, organization, and ownership across countries. We then validate our patent-quality measure via case studies, by examining its relationship with measures of firms' economic performance, such as firm value and profitability, and by examining the relationship with traditionally-used measures of innovation, such as expenditures in research and development (R&D). We are thus able to determine that our patent-quality measure constitutes a quantifiable, comprehensive, consistent, and objective innovation-quality indicator across time and space in Canada.

Finally, we use Compustat data to construct firm-specific as well as industry-wide TFP measures. We find that there is a positive and statistically significant relationship between aggregate TFP and breakthrough patents, where breakthrough patents are defined as those at the top 10 percent of our patent-quality measure. We also document that firm TFP increases on average for the five years before the grant of a breakthrough patent and even for a couple of years after the grant, depending on controls. Overall, there is a statistically and economically significant positive link between our index of innovation and measured productivity. Furthermore, the results remain significant after including year and firm fixed effects. This implies that we are capturing differences in innovation across firms, as opposed to aggregate trends. Although the analysis here relies on correlations, one potential conclusion is that significant patents enhance firm productivity as captured by measured TFP.

1.2. Related literature

In 2012, in an attempt to explain that not all patents reflect "true" or "fundamental" innovations, former University of Toronto president David Naylor stated that "Canada's innovation landscape is cluttered with brokers, buffer bodies, boutiques and regional boondoggles. Cleaning up this landscape would save millions if not billions of dollars" (Schwanen 17). Concerns have also been raised about the share of Canadian patents that are held and exploited abroad, eg. Johnson 2002, CIPO 2016a and 2016b.

A number of studies have used brief snapshots of Canadian patent data from CIPO, for example Lexchin 1993, Calabrese et al. 2000, Hall and Bagchi-Sen 2001, Doloreux 2004, Amara-Landry 2005, Doloreux-Parto 2005, Albert-Laberge 2007, Council of Canadian Academics 2013, Moyse 2015, Fortin-Hadfield 2016, Holness 2016, Boadway-Tremblay 2017, CIPO 2017, Greenspon-Rodigues 2017, Impact Center Canada 2017. However, no systematic analyses have been conducted across firms, industries, provinces, and time. Furthermore, existing Canada studies have never utilized the rich information contained in the detailed documents with patent descriptions submitted to support each patent application.

A number of studies have examined various aspects of patents and innovation in the US. The academic work on US patents has been assisted a lot by the database created by a number of NBER researchers, who also made the data publicly available. Our work aims to create a similar dataset for patents in Canada. Among other things, US studies have demonstrated that the number of patents granted may not be a great measure of innovation. This is because fluctuations in the number of patents granted are often the result of changes in patent regulation or in the quantity of resources available to the US patent office (e.g. Griliches, 1990; Hall and Ziedonis, 2001). As a result, a larger number of patents does not necessarily imply greater technological innovation (Griliches, 1998). Alexopoulos (2011) proposes an alternative innovation measure that is based on books published in the field of technology. Though the measure in Alexopoulos (2011) overcomes many of the shortcomings of patent counts, it is only available at the aggregate level and for only the later part of the twentieth century.

Patent citations may also not be a perfect innovation proxy, due to strategical reasons for the pursuit of patenting, see for example the work of Hall-Ziedonis 2001 and Abrams et al. 2013. Even more, citation of prior patents is not a requirement in Canadian legislation and the simple suggestion of citing prior influential work only entered Canadian legislation in 1973. Our proposed innovation-quality indicator is based on the degree of linguistic similarity across patent documents and it therefore overcomes the problem of the inconsistencies in the patent-citation measure and also the vagueness of the legal language in the Canada Patent Act, which requires an innovation to be "novel, non-obvious and useful". In this dimension, the paper most similar to ours is by Kelly et al. (2017), who perform textual analysis and create indicators of patent quality for US patents, using the information from the NBER database. Kogan et al. (2017) use data from CRSP to construct measures of the economic value of US patents granted to US publicly-traded firms. However, their analysis does not account for mergers and acquisitions. Li et al. (2018) is the only other systematic analysis of pharmaceutical patents. They use data from the Thomson database to examine novelties in US pharmaceutical innovations using details about the molecular structure of new medications. However, our analysis also contains information on the prices of new patented pharmaceuticals, which is not available in the US. Finally, we find that in Canadian pharmaceutical companies the granting of significant patents increases firm profitability but not firm employment. This is in contrast to the results of Kline et al (2018), who show that, in the US, an initial allowance of an ex-ante valuable patent lead firms to increase employment. This result indicates that there may be differences in firm innovation behavior between the US and Canada or that pharmaceutical companies behave differently from the average company in terms of sharing the benefits of innovation and increased profitability.

1.3. Legal framework for Canadian patents

The first federal Patent Act of 1869 established patent grants for a term of 15 years. The second federal Patent Act was passed in 1872 and allowed foreigners to register patents. In the 1880s and 1890s this Act was amended to extend patent terms from 15 to 18 years. Next, between 1900 and 1919, the Patent Office and the post of the Commissionaire of Patents were established by statute. In 1923, the third federal Patent Act added provisions for inventions by public servants. The fourth Patent Act was passed in 1935, with procedural provisions for obtaining patents related to national defense and atomic energy. In 1991, the Patent Office and the post of the Commissionaire of Patents were incorporated into the new Canadian Intellectual Property Office (CIPO). In 1993, the requirement that an invention be "not obvious" was included to the previous requirements of "novel and useful". In 1996, the agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) was passed and, in 2001, patent duration was extended to 20 years as a result of a TRIPS-related ruling by the World Trade Organization (WTO).

Quebec has also been promoting innovation efforts for many decades now. In fact, the first patent in Canada was granted in 1791 by the legislature of Quebec. Furthermore, in July 2015, the Québec government launched its unique across Canada "First Patent" program, which encourages businesses to patent their inventions and offers eligible businesses a subsidy on expenses related to obtaining a first patent. Additionally, as of January 1, 2017, Quebec established the "Patent Box" program, which reduces the tax rate on patent-related business income so as to support businesses that wish to carry their innovations through to the commercialization stage. Some exceptions to these concerns are the industries related to artificial intelligence, pharmaceuticals and medical devices. This paper starts the investigation in the outlier growth of the Canadian pharmaceutical industry by using textual analysis based on new data for Canadian patents.

In Canada, the first applicant to file a patent application is entitled to obtain the patent. Any public disclosure of an invention before filing may make it impossible to obtain a patent. Public disclosure of all materials related to a patent application is required at 18 months after filing with the Canadian Intellectual Property Office (CIPO). The average waiting period between filing and granting is eight years in Canada, compared to two years for the US. However, legal protection of patent claims starts at the 18-month mark and it therefore applies retroactively before the date of granting. In the US, legal patent rights are guaranteed for after the date of patent granting. Another major difference between the two countries is that Canadian legislation does not require citation of previous patents on which a new patent is based, whereas the US legislation does mandate for such citations. Of course, patent citations have well-known problems when used as proxies of patent quality. However, as 90 percent of new patents are improvements upon previous patents and as Canadian inventors are not obligated to cite previous state of the art, the complete lack of any patent-quality measure might be one of the reasons for the big delays in the patenting process in Canada.

1.4. Patent database description

We first construct a novel big database using data provided by the Canadian Intellectual Property Office (CIPO) on all the patents granted by Canadian authorities since 1869. These data contain information on almost 3 million patents filed in Canada over two centuries. The raw data were obtained after communication with the CIPO, which provided us with CDs containing patent information in the form of xml links, excel files, and pdf documents. We then used a combination of algorithms to combine the information into a unique longitudinal database on patents and the firms they belong to. The main aspects of the raw data are described in the appendix.

In 2016, 1,683,625 patents have been granted in Canada from which 12.5% are still in activity and the remaining had expired or been lapsed. From 200 patents in 1869, the number of patents issued has reached more than 25,000 in 2016. Patents issued in Canada has slightly increased from 1869 to 1929, and decreased in the period 1929-1945, period of the great depression and the world war II. From 1945 to 1971, we assist to a boom in patenting activity, followed by a gradual drop during the period 1971-1997 where the patent granted falls from 29 593 to 6,667 . But since 1997, the number of patents issued have increased again. The total number of patents issued followed the dynamic of the business cycles in Canada, with high patenting activity associated with economic booms and low patenting activity associated with recessions.

Focusing on more recent years, we see an increasing trend in the number of patents issued. This number has tripled from 1997 to 2017 with an average annual growth rate of 7.5%. It shows the capacity of Canada to transform ideas into innovations. Nevertheless, there is a probability of 10% that a patent issued is abandoned by the owner in the five years following the date of issuance. This number is a first attempt to see how the applicants value the patents, at least the fact of patenting in Canada.

1.5. Pharmaceutical patents

For the rest of the paper, we focus on the pharmaceutical industry in Canada. The reasons for this are related to the specific importance of pharmaceuticals for Canada but also, at this point in time, to computational constraints. Canada is the only OECD country with a public health care plan that does not include the cost of prescription drugs. Canadians pay the second highest pharmaceutical costs out of all OECD countries, the US being first. According to the Canadian Institute for Health Information these costs were estimated to rise 4.2 percent in 2017. In a recent House of Commons committee meeting on health, the Assistant Deputy Health Minister admitted that high prescription drug costs will rise under pending free trade agreements. While patent protections have increased in trade agreements, research and development in the pharmaceutical industry decreased.

We select patents with the international patent classification (IPC) code A61P. The International Patent Classification (IPC) is a hierarchical patent classification system used in over 100 countries to classify the content of patents in a uniform manner. It was created under the Strasbourg Agreement (1971), one of a number of treaties administered by the World Intellectual Property Organization (WIPO). The classification is updated on a regular basis by a Committee of Experts, consisting of representatives of the Contracting States of that Agreement with observers from other organisations, such as the European Patent Office. The first letter represents "section", where A is human necessities. The two-digit number represents "class". The final letter makes up the subclass. 61 is medical or veterinary science; hygiene. P is specific therapeutic activity of chemical compounds or medicinal preparations. In this subclass, the term "drugs" includes chemical compounds or compositions with therapeutic activity. In this subclass, therapeutic activity is classified in all appropriate places. We focus on the time period 1991-2017 for the construction of our patent- or innovation-quality measure and on the period 2000-2017 for the match of the patent and Compustat databases. This final dataset contains information on 15,919 patents, belonging to 5,000 different companies. Of those companies, about 100 are independent or "mother" companies, while the rest are subsidiaries, reflecting a high degree of oligopolistic power in the pharmaceutical industry in Canada. For most of the analysis we use the date of granting of a patent.

1.6. Patent similarity index

Term frequency (TF) gives the frequency of each word in each document in the corpus. Specifically, tt is the ratio of number of times the word appears in a document, compared to the total number of words in that document. Clearly, TF increases as the number of occurrences of the word within the document increases. In the end, each document has its own TF.

Inverse data or document frequency (IDF) is used to calculate the weight of rare words across all documents in the corpus. Thus, the words that occur rarely in the corpus have a high IDF score, because they are more informative for the similarity calculations. For example, words like "virus" and "cancer" are more useful than words like "patent" or "inventor" for the comparison of pharmaceutical patents, and therefore should enter more prominently into the similarity calculation.

The combination of these two metrics yields the TF-IDF score of each word in each document in the corpus.

A low TF-IDF score may indicate that the word appears infrequently in the document (low TF) or that it is a very common word that appears in many documents (low IDF). A high TF-IDF score indicates that a word appears relatively frequently in one document but it does not appear in most other documents, and is therefore crucial for the content of this particular document.

However, the traditional TF-IDF methodology does not take into account the chronological ordering of the patents. Hence, it cannot capture the novelty of a patent, given the history of innovations leading up to that patent. Instead construct a retrospective version of IDF or "backward-IDF" of term w for patent p, denoted by BIDFwp, as the log frequency of documents containing w in any patent granted prior to patent p.[2] This retrospective IDF changes as the frequency of use of each word changes over time, thereby providing a temporally appropriate weighting of each word. In this sense, the backwards-looking IDF reflects the history of knowledge and innovations up to the arrival of each new patent.

Clearly, $\rho \in [0,1]$. Patents that use the same words with the same frequency have similarity 1, whereas patents with no common terms have similarity 0.[3] We set similarities below 5% to zero, so as to reduce the computational burden of the analysis, thereby eliminating about 90% of pairwise comparisons.

This section describes the construction of our patent similarity measure. We use textual analysis on pharmaceutical patent documents via the statistical language Python to translate textual data into numerical data. Specifically, the textual analysis creates links between each new invention and the set of existing and subsequent patents. We thus construct a measure of textual similarity to quantify the commonalities in the topical content of each pair of pharmaceutical patents. This procedure identifies significant or high-quality patents as those with content distinct from prior patents (novelty) and similar to future patents (impact).

We pre-process the textual information by dropping "stopwords", capitalization, punctuation, and unnecessary symbols. Stopwords are commonly-used words, such as "a", "the", "and" etc., which are not informative when comparing different patent documents. We also lemmatize and/or stemmatize words, so that we can group together words that are closely related. For example, the words "electric" and "electricity" are closely related, so they are both assigned to the common root "electri". We also identify collocations and n-grams, using relevant scientific dictionaries. In the end, we are able to construct a corpus matrix where rows are patents and columns are the "informative" words in each patent.

We then employ an adjusted Term Frequency-Inverse Document Frequency (TF-IDF) methodology to construct a similarity measure across each pair of patents (each two-row combination of the corpus matrix), based on the words (columns) in each patent. For each pairwise patent comparison, we create a dictionary that contains the union of the words appearing in both patents. The TF is the number of times a word appears in a patent

[2]Note that patent numbers are assigned in the order in which they are granted.

[3]The similarity measure is related to the Pearson correlation, except for the fact that the TF-BIDF is not centered before the dot product is applied.

document, adjusted for document length. The IDF is the number of patent documents in which the word shows up. One limitation of the standard IDF is that it does not account for language changes over time. We therefore use an adjusted IDF (AIDF), which captures the number of documents in which a word shows up before time t, in the set of documents before t. The TF-AIDF method is then:

$$TFAIDF(word, document, year) = TF(word, document) * AIDF(word, year) \tag{1.6.1}$$

We next construct the cosine similarity measure, which captures the degree of of similarity of two patents. The pairwise similarity is reflected in the cosine distance metric, which is the angle between the vector representations of two patents. The cosine lies in the interval between 0 and 1. If the cosine close to 1: patent p and p' are very similar. If the cosine close to 0: patent p and p' are very different. For U and V the vector representations of patent p and p', we have:

$$cos(p,p') = \frac{UV}{\|U\|\|V\|} \tag{1.6.2}$$

We provide an example of U and V in the next section.

1.6.1. Example

Let's compare two pharmaceutical patents, for simplicity using only their abstracts. The raw data are as follows.

First patent. Patent number: 2481369. Title: INHIBITORS OF SERINE PROTEASES, PARTICULARLY HEPATITIS C VIRUS NS3 - NS4 PROTEASE. Assignee: VERTEX PHARMACEUTICALS INCORPORATED. Date of issuance: 2012-07-10. Abstract: The present invention relates to compounds of formula IA (see formula IA) as defined herein, that inhibit serine protease activity, particularly the activity of hepatitis C virus NS3-NS4A protease. As such, they act by interfering with the life cycle of the hepatitis C virus and are also useful as antiviral agents. The invention further relates to compositions comprising these compounds either for ex vivo use or for administration to a patient suffering from HCV infection. The invention also relates to methods of treating an HCV infection in a patient by administering a composition comprising a compound of this invention. The invention further relates to processes for preparing these compounds.

Second patent. Patent number: 2812261. Title: MACROCYCLIC PROLINE DERIVED HCV SERINE PROTEASE INHIBITORS. Assignee: ENANTA PHARMACEUTICALS, INC. Date of issuance: 2017-02-21. Abstract: The present invention discloses compounds of Formula (I) or pharmaceutically acceptable salts, esters, or prodrugs thereof: Formula (I) which inhibit serine protease activity, particularly the activity of hepatitis C virus (HCV) NS3-NS4A protease. Consequently, the compounds of the present invention interfere with the life cycle of the hepatitis C virus and are also useful as antiviral agents. The present invention further relates to pharmaceutical compositions comprising the aforementioned compounds for administration to a subject suffering from HCV infection. The invention also relates to methods of treating an HCV infection in a subject by administering a pharmaceutical composition comprising the compounds of the present invention.

The cleaned data for each abstract are as follows.

Patent number: 2481369. Abstract: [compoun, formul, formul, inhibit, serin, proteas, activit, activit, hepatitis, virus, proteas, interfer, liv, cycl, hepatitis, virus, antivir, agen, composit, compris, compoun, viv, administrat, patien, suffer, infect, method, treat, infect, patien, administ, composit, compris, compoun, proces, prepar, compoun]

Patent number: 2812261. Abstract: [disclos, compoun, formula, pharmaceutic, salt, ester, prodrug, formula, inhibit, serin, proteas, activit, activit, hepatitis, virus, proteas, compoun, interfer, lif, cycl, hepatitis, virus,antivir, agent, pharmaceutic, composit, compris, compoun, administrat, subject, suffer, infect, method, treat, infect, subject, administ, pharmaceutic, composit, comprise, compoun]

The respective vectors for each abstract are:

$$U = TfIdf(activity, pat1, year) * activity + TfIdf(administer, pat1, year) * administer + \\ TfIdf(antiviral, pat1, year) * antiviral + ... \tag{1.6.3}$$

and

$$V = TfIdf(activity, pat2, year) * activity + TfIdf(administer, pat2, year) * administer + \\ TfIdf(antiviral, pat2, year) * antiviral + ... \tag{1.6.4}$$

while the cosine similarity measure is:

$$Cos(U,V) = \frac{UV}{\|U\|\|V\|} = 0.8 \tag{1.6.5}$$

The cosine is 0.8, meaning that these two patents are similar.

1.6.2. Descriptive statistics

Figure (1) plots the distribution of the similarity score, focusing on patent pairs that are 0–20 years apart. The similarity distribution is highly skewed. Patents tend to be highly dissimilar, with only a small fraction of pairs very closely related. The median similarity score across patent pairs is 7.8%, whereas the average similarity score is 10.2%. In the right tail, the 90th and 95th percentiles of similarity scores are 17.6% and 22.9%, respectively. In other words, the matrix is sparse.

Fig. 1. Cumulative density function of patent-similarity or cosine measure

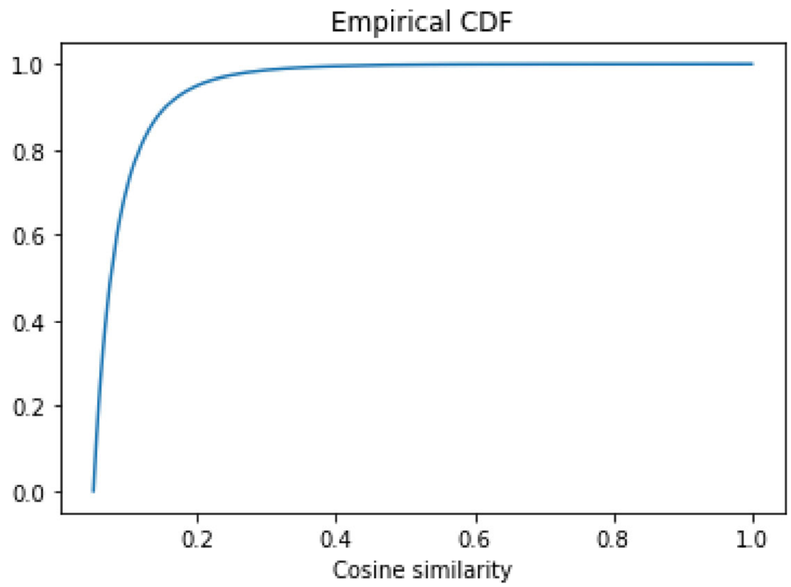

1.7. Patent significance index

In this section, we construct a measure of quality for each patent, by aggregating its comparisons with other patents into a single index. A patent has high quality when it combines novelty, compared to the past, and impact, compared to the future. In other words, the most important or significant patents are conceptually different from the predecessors and they also influence future scientific advances. The latter is reflected in high similarity to subsequent innovations.

First, we measure a patent's novelty as the (inverse of) its similarity with the existing patent stock at the time it was filed. We refer to this as "backward similarity". Second, we measure a patent's impact by its "forward similarity". The forward similarity measure estimates the correlation between each patent and the technological innovations over the next τ years.

The indicator of patent significance then combines novelty and impact, attributing higher scientific value to patents that are both novel relative to the status quo and influential for future research. The index normalizes by backward similarity, because a patent may have high forward similarity, and therefore high numerator, either because it reflects a breakthrough or because it is a follower in a technology area with many other followers, in which case it will have a high backward similarity as well. Overall, the significance indicator follows the same logic as that behind indicators relying on future citations. Specifically, the numerator is the sum over similarity with future patents—which is directly analogous to the sum of future citations.

In this section, we use the cosine similarity measure from above to construct a measure of patent or innovation quality. Here, we build a summary measure of patent quality that incorporates both the patent's impact (forward similarity or FS) and novelty (backward similarity or BS). We construct a measure of the scientific importance of a patent, as the ratio of the patent's future impact (FS) to its novelty (BS). We refer to the measure as "relative forward similarity" and interpret it as an overall measure of patent quality. In particular, our summary measure attaches higher scientific value to patents that are more novel relative to their predecessors but are related to subsequent research. Forward similarity measures the strength of association between the patent and future technological innovation, and normalizing by backward similarity emphasizes the novelty of the patent. A patent may have high forward similarity because it is a "follower" in a technology area with many other followers, in which case it is likely to also have a high backward similarity as well. On the other hand, its high forward similarity may indicate a new and impactful breakthrough, in which case it is likely to have low backward similarity, and thus an especially high relative forward similarity. Further, another reason why a patent might have high forward similarity is that it uses general language that is not distinct to any particular technology but is stylistically common. In this case, normalizing by backward similarity counteracts the effect of general language on measured impact.

We compare each patent granted in each year t to all patents in the previous five years and in the subsequent five years. The same can be done for all different comparison horizons, T.

Hence, the BS measure is the sum of pairwise cosine similarities of patent p, issued in t, with patents issued in $t - T$

$$BS^0_{-T} = \Sigma_{p'} cos(p,p') \tag{1.7.1}$$

The FS measure is the sum of pairwise cosine similarities of patent p, issued in t, with patents issued in $t + T$

$$FS_0^T = \Sigma_{p'} cos(p,p') \tag{1.7.2}$$

Finally, the patent quality measure, defined as q, is:

$$q^T = FS_0^T / BS_{-T}^0 \tag{1.7.3}$$

1.7.1. Descriptive statistics

Table decomposes the variation in patent quality q into variation that arises from differences in the calendar year the patents were filed (which could be the case, for example, if there systematic differences in the quality of innovation across years), differences between technology classes (which might reflect, for example, differences in general purpose versus specific purpose technologies), and differences across patent assignees (which might arise, for example, if firms are heterogenous in innovation quality). Since many patents have no assignees, we perform the analysis separately with and without assignee fixed effects. For comparison we perform the same exercise for the (logarithm of one plus) the number of forward citations the patent receives. In the interest of space, we focus on forward similarity (and forward citations) in the five years following a patent filing. Technology class fixed effects account for a relatively small share of the overall variation (less than 10%). This is true for both text-based quality and citations. In contrast to technology class, assignee fixed effects account for approximately 20% of the overall variation for both quality and citations. This is an important result that suggests that innovativeness varies predictably across assignees. Finally, patent year cohort effects account for a significant share of variation, particular for patent quality. Though it is possible that these time effects capture variation in the rate of technological innovation, they also likely reflect the presence of other nuisance factors, for instance shifts in language or variation in USPTO standard for granting a patent.

1.7.2. Validation

In this section, we examine the relationship between the patent significance index and more traditional measures of patent quality, such as patent citations. Usually, a pharmaceutical company files for a patent in Canada as well as in the US. However, US legislation requires citations. We therefore match Canadian patents to their equivalents in the US by designing an algorithm that uses Google patents to collect the number of citations for similar patents that were granted in the US. We further examine the relationship with patent

Table 1. Validation Quality-Citation

Log(citations)		
Log(quality)	0.22***	0.31***
Rsquare	0.42	0.63
Observations	19150	19150
Grant Year FE	Y	
Assignee FE	Y	
Grant Year* Assignee FE		Y

This table shows the regressions coefficients of the relation between the quality of a patent and the number of citations of this patent. Standard errors are robust and clustered at the assignee level. ($* = p < 0.10$, $** = p < 0.05$,$* * * = p < 0.01$)

citations. The results are shown in Table 1. There is a positive and statistically significant correlation between high patent quality (high q) and number of citations of that patent.

1.8. Patent significance and firm performance

This section performs validation tests of our patent significance measure, q, by relating q to measures of firm performance. These measures are constricted via the match of our patent database with the database of Compustat North America. For the purposes of this analysis, we define breakthrough patents as those at the top 10 percent of our patent quality measure q. Economic measure of significance, as opposed to the above which is more on the "true" or technological aspect. The two need not be perfectly correlated, especially if there are commercialization problems, as may be for Canada.

In this section, we discuss the relation between patent quality and market valuations. Market values are by definition private values; they measure the present value of pecuniary benefits to the holder of the patent. By contrast, our quality measure is designed to ascertain the scientific importance of the patent. The relationship market value and scientific importance can be ambiguous. For instance, a patent may represent only a minor scientific advance while being very effective in restricting competition, thus generating large private rents. The relation between the private and the scientific value of innovation—as measured by patent citations—has been the subject of considerable debate in the literature.[4]

[4]For instance, Hall et al. (2005) and Nicholas (2008) document that firms owning highly cited patents have higher stock market valuations. Harhoff et al. (1999) and Moser et al. (2011) provide estimates of a positive relation using smaller samples that contain estimates of economic value. By contrast, Abrams et al. (2013)

In what follows, we revisit the empirical literature that studies this relationship using our text-based measure of patent quality. We do so at two levels of granularity. Section 1 analyzes patent level data, where the estimated market value of each patent is based on stock market reactions in a narrow window around the issuance date, following the methodology of Kogan et al. (2017). In section 2 we perform the analysis at the firm level, relating differences in firm valuation ratios (Tobin's Q) to differences in the quality of firms' patent portfolios, following Hall et al. (2005).

1.8.1. Patent significance and firm value

We examine if our innovation significance measure explains differences in firm value. Our analysis closely follows that of Hall et al. (2005), who estimate the relationship between a firm's Tobin's Q and its "knowledge stock". The knowledge stock is defined as the depreciating balance of the firm's $R\&D$ investment, its patents, and its patent-citation count, according to:

$$SX_{f,t} = (1-\delta) \cdot SX_{f,t-1} + X_{f,t} \tag{1.8.1}$$

where, for each variable X for firm f in year t, $X_{f,t}$ represents flows and $SX_{f,t}$ represents accumulated stocks. We assume that the depreciation rate is $\delta = 15\%$. The variable X may be new R&D investment or successful patent applications or patent citations. We introduce a fourth knowledge stock variable based on our patent quality measure. We define firm-level patent quality for firm f in year t as:

$$q^{\tau}_{f,t} = \sum_{j \in J_{f,t}} q^{\tau}_{j} \tag{1.8.2}$$

where, $J_{f,t}$ is the set of patents filed by firm f in year t. We then create a "quality-weighted" patent stock that accumulates (1.8.2) according to (1.8.1).

The firm-level regression specification is:

$$log(Q_{it}) = b_0 + b_1 \frac{SRD_{it}}{A_{it}} + b_2 \frac{SP_{it}}{SRD_{it}} + b_3 \frac{SQ_{it}}{SP_{it}} + b_4 D(RD_{it} = 0) + v_i + u_t + u_{it} \tag{1.8.3}$$

where $SRD_{f,t}$, $SPA_T f$,t, $SCITES_{f,t}$, and $q_{f,t}$ are the stocks of R&D expenditure, number of patents, patent citations, and the patent quality measures constructed as in (14).

The $R\&D$ stock is scaled by total assets, $A_{f,t}$, the patent stock by the R&D stock, and the citation stock by the patent stock. We scale our patent quality stock by the stock of patents by count, giving it an interpretation as the average quality of patents held by the firms. We

use a proprietary dataset that includes estimates of patent values based on licensing fees and show that the relation between private values and patent citations is non-monotonic.

estimate the market value regressions using quality and citation stocks over horizons τ of 1, 5, or 10 years after the application date. For our baseline results, we restrict the sample to patenting firms (that is, firms that have filed at least one patent). Also, at is the fixed effect for year t and accounts for any time specific effect that moves around the value of all the firms in a given year. We also include a dummy variable for missing *R&D* observations. We cluster standard errors by firm.

Table 2. Tobin's Q and Quality

Log Q	(1)	(2)	(3)	(4)	(5)
R & D capital stock / total Assets	0.74***	0.65***	0.66***	0.56***	0.53***
	(3.80)	(4.38)	(3.74)	(3.86)	(4.97)
Patent stock/ R & D stock	0.37**	0.55**	0.32	0.40	0.35
	(1.97)	(2.15)	(1.63)	(1.31)	(1.08)
Quality index stock/patent stock		0.13**		0.15**	0.13**
		(2.32)		(2.58)	(2.34)
Dummy(R & D=0)	-1.11***	-0.98***	-1.09***	NA	NA
	(-3.67)	(10.91)	(-3.22)		
N	1332	728	1332	728	726
R2	0.19	0.19	0.19	0.17	0.15
Year FE	Y	Y	Y	Y	Y
Firm FE			Y	Y	Y

Breakthrough patents and firm value. We regress the logarithm of a firm's Tobin's Q on the following firm stocks: the stock of *R&D* expenditures; the stock (number) of patents; the stock of patent quality. All stocks are constructed as shown in equation. The depreciation rate is assumed to be $\delta = 0.15$. Standard errors are clustered at the firm level.

Our main coefficient of interest is b_3 which estimates the relationship between quality-weighted patent stock and firm value. Table 2 presents the results. Examining column (2), we see a strong and statistically significant relation between Tobin's Q and the patent quality stock. A one-standard deviation increase in the (per-patent) quality stock is associated with a 0.15 log point increase in Tobin's Q—evaluated at the median—which is economically significant given that the unconditional standard deviation in log Tobin's Q is equal to 0.63. For comparison, a one-standard deviation increase in the citation-weighted stock in column (3) is associated with a 0.13 log point increase.

The results indicate that the stock of patent quality has predictive power for firm value, as captured by Tobin's Q, on top of that of the stock of research and development spending and the total number of patents owned by the firm. The relationship remains statistically and economically significant when controlling for a variety of controls, as shown in the different columns of the table. Taken together, our findings in Section 1 and 2 show that our quality indicators are systematically related to market values, even controlling for patent citations. Given that these estimates are based on data from the later part of the sample, when citation data are broadly available, these results reinforce the view that our text-based measure captures information about patent quality that is not fully incorporated in patent citations

1.8.2. Patent significance and firm performance

We define a "breakthrough" patent as one that falls in the top 5% of the quality distribution (among all patents in all years). Our baseline results use quality with a 5-year forward window. We also compare against an alternative definition of breakthrough patents based on the 5% of patents with the most forward citations over the same horizon (and likewise adjusted for year fixed effects).

An advantage of our innovation measure is that it allows us to analyze the relation between innovation and economic outcomes at a fairly granular level. We next examine patterns at the assignee level.

Next, we focus our analysis to firms we can match in Compustat—and therefore have much more detailed information—and examine the response of firm profitability to the event of having a breakthrough innovation. Given that the distribution of these breakthroughs is highly skewed—over 90% of firm-year observations have no breakthroughs, while a small fraction (1%) of the observation have more than 15 breakthroughs, we define our main variable of interest as a dummy variable that takes the value of one if the firm had a breakthrough in a given year or zero otherwise. Given the increased level of granularity, the appropriate dating of these breakthroughs becomes more important. As our baseline case, we date patents as of the year the patent application is filed–as opposed to when the patent is issued. We do so because firms may utilize the innovation that is associated with patent even before the application is approved by the CIPO. This is important especially for Canada where the duration between the application and the grant of the patent is approximately 7 to 8 years.

The dependent variable is the growth in average profits from t to t+h. We focus on the growth in average profits over a period, rather than on the year-to-year changes in profitability to smooth out transitory variations in profitability. We consider two definitions for profitability. First, we focus on gross profitability, defined as sales minus costs of good sold. This specification informs us on the extent to which innovation is associated with higher firm growth. In addition, we also examine gross profits scaled by the number of employees; this definition informs us on whether innovation enhances labor productivity. We winsorize all variables at the 1% level. Since the exact timing of when these breakthrough innovations may affect profits is somewhat ambiguous, we examine horizons of up to ten years after the patent applications, as well as up to five years prior.

Our ideal thought experiment compares two otherwise identical firms, one of which generated a breakthrough innovation and another that did not. As a result, the vector of controls $Z_{f,t}$ includes firm variables that are related to future profitability, but also the variables which predict the likelihood of successful innovation by the firm, as we document in the section above. Thus, we control for the logarithm of firm size (defined as total book assets); the log of the current level of profitability by the firm; a dummy for whether the firm filed for a patent in year t; the log of (one plus) its number of patent applications; firm age based on the first appearance in Compustat; the stock of patents as of year $t-1$ (in logs); and the share of patents that are breakthrough innovations as of year $t-1$. Standard errors are clustered by firm and year.

We further validate our quality index using firm performance indicators. The left-hand-side variable is the average change in firm profitability in the five years before a company gets granted a breakthrough patent or the average change in firm profitability in the five years after a company gets granted a breakthrough patent. The right-hand-side variable include a set of controls X: the log of the total book assets (to capture firm size); the log of the current level of profitability; a dummy for whether the firm has been issued a patent in year t; the log of (one plus) its number of granted; the firm age based on its first appearance in Compustat; the log of the stock of patents as of year t-1; and the share of breakthrough patents as of year t-1; an indicator which takes the value of 1 in years when a firm gets a breakthrough-patent grant and 0 otherwise; time- and firm-fixed effects.

$$log(\frac{1}{h}\Sigma_{t=1}^{h}pr_{it}) - log(pr_t) = c + a_h D(Breakthrough_{it}) + gX_{it} + v_i + u_t + u_{it+h} \qquad (1.8.4)$$

Figure (2) shows the relationship between breakthrough patents and growth of profits by employee. This figure plots the relationship between firm profits by employee growth and a dummy variable that takes the value of one if the firm has a breakthrough patent. Point $t=0$ indicates the granting year of the patent. Years $t-\tau$ are before granting. Years $t+\tau$ are after granting. Controls include a dummy variable for whether a firm has filed any patents during the period, the (log) number of patents and industry-year fixed effects. Patent quality or significance is measured as the ratio of the 5-year forward similarity to the 5-year backward similarity. Breakthrough patents are those in the top 10% of the patent-quality measure (net of year fixed effects). Employment and profit data data are from Compustat. Profits are sales minus cost of goods sold. Standard errors are clustered by firm and year. The

Fig. 2. Breakthrough patents and growth of profits by employee

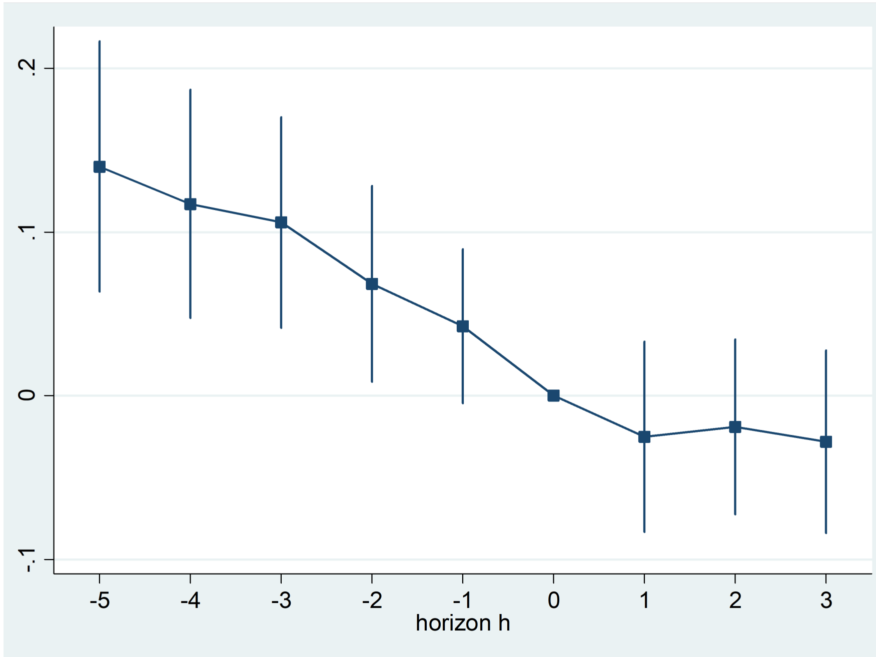

Plots the relationship between firm profits by employee growth and a dummy variable that takes the value of one if the firm has a breakthrough patent. Point $t = 0$ indicates the granting year of the patent. Years $t - \tau$ are before granting. Years $t + \tau$ are after granting. Controls include a dummy variable for whether a firm has filed any patents during the period, the (log) number of patents and industry-year fixed effects. Patent quality or significance is measured as the ratio of the 5-year forward similarity to the 5-year backward similarity. Breakthrough patents are those in the top 10% of the patent-quality measure (net of year fixed effects). Employment and profit data data are from Compustat. Profits are sales minus cost of goods sold. Standard errors are clustered by firm and year.

graph shows that the anticipation of the granting of a breakthrough patent increases firm profitability (per worker), on average, for a period of five years before the patent grant. This evidence suggests that firms start adjusting their behavior before the grant of a significant patent.

Figures (3) and (4) show the relationship between breakthrough patents and firm growth of employment and of capital expenditures. These relationships are not statistically significant, either before or after the granting of the breakthrough patent. Combining this with the profitability results from above, it would appear that breakthrough patents increase firm markups via increase in sales. The increase in the profit margin is not translated into increases in employment or investment in capital expenditures.

Fig. 3. Breakthrough patents and firm employment growth

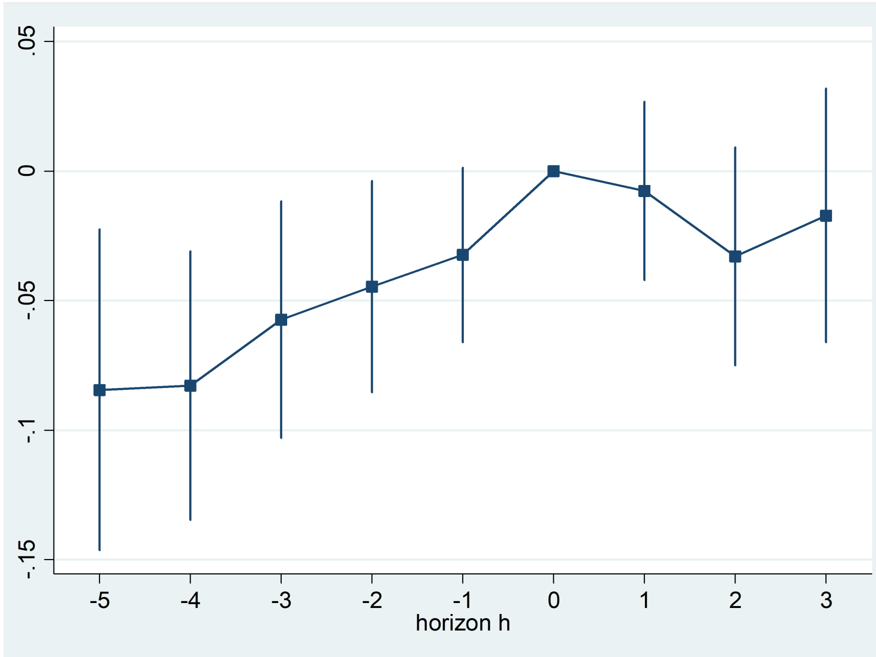

Plots the relationship between firm employment growth and a dummy variable that takes the value of one if the firm has a breakthrough patent. Point $t = 0$ indicates the granting year of the patent. Years $t - \tau$ are before granting. Years $t + \tau$ are after granting. Controls include a dummy variable for whether a firm has filed any patents during the period, the (log) number of patents and industry-year fixed effects. Patent quality or significance is measured as the ratio of the 5-year forward similarity to the 5-year backward similarity. Breakthrough patents are those in the top 10% of the patent-quality measure (net of year fixed effects). Employment data are from Compustat. Standard errors are clustered by firm and year.

1.9. Patent quality and TFP

We construct firm-specific measures of TFP, using data from Compustat and two alternative methods. First, TFP as Solow residual. Second, TFP adjusted for selection effects, using the methodology in Olley-Pakes. Results are similar across both methods. The figures refer to the first TFP calculation. We also aggregate the firm-specific TFP measures, using as weights the share of the company sales in total industry sales, to generate an aggregate TFP index.

Fig. 4. Breakthrough patents and firm capital-expenditure investment

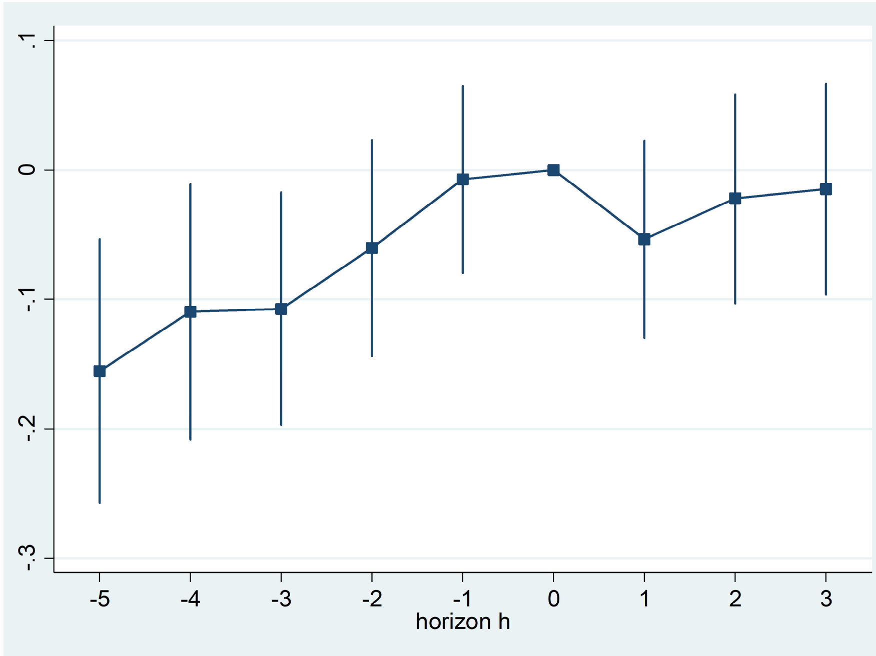

Plots the relationship between firm capital-expenditure growth and a dummy variable that takes the value of one if the firm has a breakthrough patent. Point $t = 0$ indicates the granting year of the patent. Years $t - \tau$ are before granting. Years $t + \tau$ are after granting. Controls include a dummy variable for whether a firm has filed any patents during the period, the (log) number of patents and industry-year fixed effects. Patent quality or significance is measured as the ratio of the 5-year forward similarity to the 5-year backward similarity. Breakthrough patents are those in the top 10% of the patent-quality measure (net of year fixed effects). Capital expenditure data data are from Compustat. Standard errors are clustered by firm and year.

Fig. 5. TFP distribution

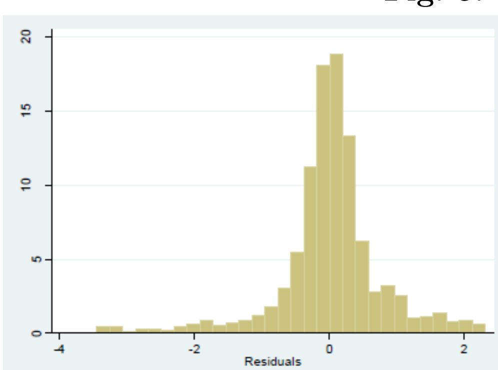

(a)

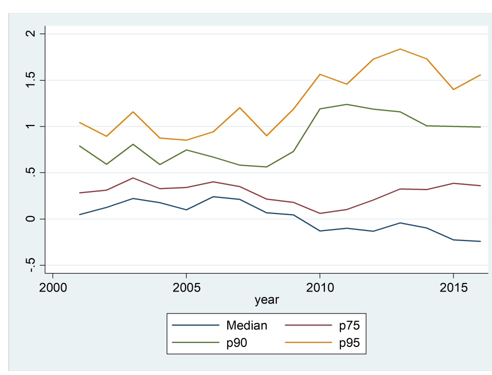

(b)

Panel (a) shows the pooled distribution of Total factor productivity. Panel (b) shows the distribution of Total factor productivity overtime.

Fig. 6. Breakthrough patents and TFP

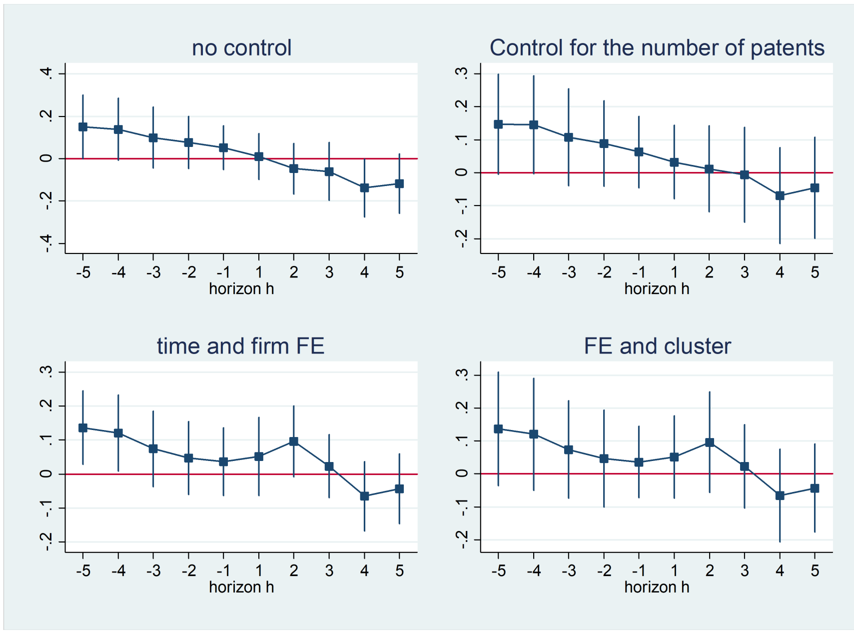

Breakthrough patents and firm TFP. Panel A (B) plots the relationship between firm-level TFP based on and a dummy variable that takes the value of one if the firm has a breakthrough patent. Point t=0 indicates the granting year of the patent. Years t-5 are before granting. Years t+5 are after granting. Controls include a dummy variable for whether a firm has filed any patents during the period, the (log) number of patents and industry-year fixed effects. Patent quality or significance is measured as the ratio of the 5-year forward similarity to the 5-year backward similarity. Breakthrough patents are those in the top 10% of the patent-quality measure (net of year fixed effects. Profits are Compustat sales minus costs of goods sold. Profits per worker is profits divided by the number of employees from Compustat. Standard errors are clustered by firm and year.

Panel (a) of Figure 5 shows the pooled distribution of TFP. Figure shows the evolution of the aggregate TFP distribution over time. Panel (b) of Figure 5 shows that there is a positive and statistically significant relationship between the aggregate TFP measure and the number of breakthrough patents. Figure shows that firm TFP increases on average for the five years before the grant of a breakthrough patent and even for a couple of years after the grant, depending on controls.

Overall, there is a statistically and economically significant positive link between our index of innovation and measured productivity. Furthermore, the results remain significant after including year and firm fixed effects. This implies that we are capturing differences in innovation across firms, as opposed to aggregate trends. Although the analysis here relies on correlations, one potential conclusion is that significant patents enhance firm productivity as captured by measured TFP.

1.10. Conclusion

In this article, we have built an index to measure the scientific value of a patent. This index is based on textual information contained in patent submission documents. Then, It makes it possible to exploit all the richness of the database and to build an indicator that disregards any desire of applicants to cite or not another patent. We believe that this indicator could be very useful, in particular for identifying key patents and thus reducing the administrative time between submission and obtaining the patent. In such a case, Canada would avoid the relatively large number of patent lapses and could better retain innovations on its territory. Besides, our quality indicator is linked to business performance. This, therefore, makes it possible to generate a market value for these intangible assets. Finally, the aggregated quality measure through breakthrough patents makes it possible to better capture technological innovation and therefore to overcome the shortcomings of existing methods.

Chapter 2

Innovative Ideas and Gender Inequality

2.1. Introduction

Context: Women face a lower entry rate and a higher exit rate than men in industries or fields that require mathematical skills and analytical abilities.[1] As a result, women are underrepresented in those fields, especially in top-ranked positions (Ceci et al. 2014; Ginther and Kahn 2004). Preferential choices such as family, risk aversion, and competitiveness, along with discriminatory factors, have been suggested as potential explanations for this gap. Yet one mechanism received less or no attention: the recognition of women's works. Knowing that being recognized and valued for their work can be a motivation for starting or continuing in a field, it is necessary to consider the question of the credit given to women's work.[2] Especially since ideas are at the core of the research and innovation process, a bias in the credit granted can also create social suboptimality via a "misallocation" of human capital.

Contribution: This paper analyzes the state of intellectual property in academic research in economics, with an emphasis on the recognition of women's works. In this sense, intellectual property at the academic research level will work through the recognition of individual work and the acknowledgement of relevant prior literature. Academic research provides an ideal framework for analyzing how women's ideas are perceived, used and referred to in society. Therefore, the paper explores whether the innovative ideas of female

[1]In 1999-2000, 13% of women received a bachelor degree in education versus 4% for men; 2% of women received a bachelor degree in engineering versus 12% for male (2001 Baccalaureate and Beyond Longitudinal Study, Zafar (2013)). Antecol and Cobb–Clark (2013) reach the same conclusion using survey data from the National Longitudinal Study of Adolescent Health over the period 1994-2008 on a survey database. Preston (1994) documents the higher exit rate of women in math-intensive fields. Hunt (2016) uses survey data from the National Survey of College Graduates to examine the difference in exit rates of women in science and engineering compared with other fields.

[2]A parallel can be drawn between respect for intellectual property. Indeed, the various processes for the protection of intellectual works (Trademarks, Copyrights, and Patents) aim to encourage authors to engage in innovative activities by guaranteeing them recognition, even exclusivity, over their production.

authors are listed as they should be in the references of the articles that follow. The paper focuses on economics within academia for two main reasons. First, the representation gap is among the largest in economics (Bayer and Rouse 2016). Second, many voices have recently been raised against gender discrimination in economics research, which seems to be more prevalent than in other life sciences or engineering (Ginther and Kahn 2004, Wu 2018, Sarsons 2019). Thus, this paper sheds new light on the lack of recognition of women's works. Further, it contributes to the existing works by exhibiting the heterogenous pattern in the omission bias, pointing out the more vulnerable female authors. Finally, it proposes possible ways to address the gender omission problem.

Methods: To achieve the research objectives, this article uses bibliometric data on articles published in major economic journals. These data come from sources like Web of Science, Ideas Repec, and Econlit among others. In the second step, the textual analysis based on the tools of big data and machine-learning adds key information to the analysis. On the one hand, these techniques have made it possible to extract the references of articles. On the other hand, they enable to establish a link between the article citing, the cited articles and the articles that should be cited. Two key indexes are constructed from this analysis. The first is an omission index, which is the novelty of the current paper. It measures the propensity with which an article forming part of the existing literature of some papers is omitted from the references of those papers. It captures the fact that an article that has several similarities with another one in the future, is not mentioned in the references of the latter.[3] The second is the innovation index, which measures the quality of an article. Indeed, this index provides an alternative way to assess the quality of an article. As opposed to citations, this measure is less likely to be biased.[4] Similar to Kelly et al. (2018) and Koffi and Panousi (2019), an article is considered very innovative (hence of high quality) if it is new and influences future research. The index of omission coupled with the innovation index allows to contrast what should be and what we can observe with the citations. Further, a combination of a probabilistic algorithm and manual collection is used to identify the gender of the authors of an article. Finally, the observable characteristics of each article are combined to build an author-level database for evaluating the effects of omission for authors in terms of future publications.

Findings: Turning to the findings, this paper first validates the metric used to compare two papers. Indeed, the more two papers are similar in the sense of the metric, the greater the probability that one is cited by the other. It is also a way of seeing that the omission index

[3]This omission may be voluntary or involuntary.
[4]See Lampe (2012) and D'Ippoliti (2018) among others.

does capture a no-citation relationship between two rather similar articles. Conversely, the more innovative an article is (higher quality article), the more it will be cited. The average innovation index of a journal is also positively related to the impact factor of that journal.

Second, this paper documents the state of intellectual property in academic research in economics. I find that on average, conditional on the distance metric constructed, a paper omits almost half of related prior papers. However, this finding is robust to alternative distance measures between papers. Third, omitted papers are 15% to 30% more likely to be female authored than male authored. Mixed team papers (with both male and female authors) tend to fall in between both genders. The papers most likely to be omitted are written by women (solo, mostly female team) working at mid-tier institutions, publishing in non-top journals. In a group of related papers, the probability of omission of those papers increases by 6 percentage points compared to men in similar affiliation when the citing authors are only males. Overall, for similar papers, having at least one female author reduces the probability of omitting other women's papers by up to 10 percentage points, whereas having only male authors increases the probability of being omitted by almost 4 percentage points. Moreover, the omission bias is twice as high in theoretical fields that involve mathematical economics than in applied fields such as education and health economics. In addition, in top journals, even papers written by women published in top journals are not exempted from the omission bias. This indicates that controlling for the quality of a paper using the journal of publication, we are more likely to pick up on a lower bound. Indeed, if there is a bias in the standards imposed on men and women, then articles published by women in top journals are of better quality than those published by men, and yet the bias still exists.

Finally, being omitted with respect to past publications reduces the probability of getting published in a top 5 journal in the future by up to 4%. To ensure that one is not capturing an effect from the type of journal, the paper focuses on a journal recognized by a wide range of rankings as a top economics journal. All those regressions include observable characteristics of the papers such the affiliation of the most prolific authors, the primary field, the year of publication, and the previous publication record of the authors.

Finally, a potential explanation for why women could be omitted from references is given. Taste-based discrimination could be one reason, but the argument of statistical discrimination seems to be the most relevant in this case. In fact, the omission bias is more prevalent for women at mid-tier institutions. Authors may thus prefer male authors because of their historically more numerous citations than women and the prestigious role they may have in the profession.[5] The analysis ends with a difference in difference estimation to see if the change in the editor's gender could lower the omission of women from references. The change in the editor's gender seems to reduce the omission bias against women. However, the coefficients are not statistically significant at the time of the change, but become statistically significant if one permits a delay in the effectiveness of the policy. Lastly, having at least one author in the citing paper and the cited paper in the same institutions ("peer effects") tend to reduce the omission bias against women significantly.

2.2. Related Literature

Overall, this article demonstrates that the lack of recognition of women's work is also noticeable through the non-reference of articles written by women. In addition, through the subjects discussed and the techniques used, this study builds on several areas of the economic literature.

First, the question of whether women get enough credit and therefore recognition for their research is at the core of this paper. In this sense, this paper is complementary to Sarsons (2019). Indeed, Sarsons (2019) tests the uncertainty about the individual contributions of co-authors favors men in terms of tenure rates compared to women. Here, I explicitly use article references to assess to whom credit is most often attributed and if this is done to the detriment of women. Moreover, the findings of Sarsons (2019) suggest that women are worse off when they collaborate with men. Similarly to Hengel and Moon (2019), I show that women also fare worse when they do not collaborate with men. In fact, mixed gender teams received treatment midway between that received by single-gender teams.

This paper is linked to the general literature on gender discrimination in academic research. More specifically, three key points emerge from the recent literature.

The first one is the presence of stereotypes. Wu (2018) highlighted that female authors are most often associated with physical characteristics while male authors are most associated with intellectual characteristics. The second element is the difference in standards and evaluations between men and women. For example, Hengel (2019) shows that women experienced longer delays in the review process and are asked to make much more revisions before getting published. In the same line, Card et al. (2019) show that to publish in the same journal as males, females must have higher quality articles. As in Wu (2018), this paper uses textual analysis techniques to extract relevant information. This paper additionally constructs two indices revealing hidden patterns that traditional numerical data do not

[5] Controlling for the existing number of citations reduces the bias. However, because the analysis focuses on omissions, so citations, checking for the existing citations is actually a "bad control".

highlight. Further, it adds to this literature by arguing that beyond higher standards and stereotypes, women still face a lack of recognition of their work even when they publish high quality papers compared with their male colleagues. Moreover, in a general discrimination analysis, this paper also addresses a question raised by Hammermesh (2018), namely that merit may not always go to the rightful person.[6]

The last point is the existence of gender bias in citation patterns. Citations as well as the journal of publication (Hilmer, Ransom and Hilmer (2015), Heckman and Moktan (2019)) are commonly used measures to evaluate the quality of a paper. However, Fong and Wilhite (2012, 2017) show how citations may not necessarily reflect the merit of the cited article or are manipulated to increase the journal impact factor. Citations could therefore reflect a strategic decision (Lampe (2012)) or characterize a network (D'Ippoliti (2017)).[7]. At this general level, this paper departs from and complements the existing literature by first building a citation database over time. Second, this paper contrasts realized citations and expected citations. Third, this paper uses an alternative measure of the scientific quality of a paper.[8] Focusing on gender, Ferber (1986, 1988) and Dion et al. (2018) show that women's papers are mostly cited by women's papers. The current findings are in line with those of Ferber (1986, 1988) and Dion et al. (2018). Indeed, in an omission perspective, women's papers are more likely to be omitted by men's papers.

In addition, one key question is why we care about citations or missing citations. Ellison (2013), Hamermesh and Pfann (2012), Jensen et al. (2009) argue that citations are important in determining labor market outcomes. They signal reputation and are important for hiring, salaries, tenure, and grants. In line with those findings, this paper further shows that being omitted influences an author's future publication possibilities. Missed authors tend to have lower chance of publishing in top economic journals.

The present paper also tackles the general role of the editorial board in economics. Abrevaya and Hamermesh (2012) find no evidence of relative favoritism toward one's own gender in the reviewing process. Card et al. (2019) maintain that editors tend to be more favorable to female authors, although they impose higher standards on women. Bransch and Kvasnick (2017) suggest that the share of women who publish in top journals does not increase with

[6]For more literature on gender and academia, see Moss-Racusin et al. (2012), Chari and Goldsmith-Pinkham (2017), Teele and Thelen (2017), Ductor et al. (2018), Auriol et al. (2019), Lundberg and Stearns (2019), Hospido and Sanz (2019)

[7]Additional evidence of the networking effect can be found in this study by Colussi (2017), which demonstrated that publications in a journal are influenced by the social connections, faculty colleagues and Ph.D. students.

[8]This measure is based on a similar measure built in the patent literature. See Koffi and Panousi (2019), Kelly et al. (2019)

the fact of having female editors. Compared with those papers, the current analysis investigates the role of the editors' gender on the omission of women from references. The results show that having women as editors or more editorial board diversity in general may reduce omissions against women.

The remainder of the paper is organized as follows. First, the data are described and evidence of gender bias in omissions is provided. Second, the main empirical strategy is described. Third, the effect of omissions on future productivity is assessed. The paper ends with a discussion of potential solutions to improve the recognition of women's works, followed by a conclusion.

2.3. Data description

The raw data are collected from two main websites, the Web of Science (WoS) database and IdeasRepec (IR). Together, these sites constitute the largest depository of academic research in economics. A web crawling algorithm is used to collect information from Ideas Repec (IR). This information is then organized into a novel database.

First, a corpus is created from all papers published in the top 16 journals in economics over the period 1991-2019. Details about journal ranking can be found in Laband and Piette (1994), Kalaitzidakis et al. (2003, 2011), Kodrzycki and Yu (2006), Engemann and Wall (2009), Bornmann et al. (2018), Thomson and Reuters Clarivate Analytics, and IR. The full list of journals is provided in the appendix. As is well known, published papers are submitted to a range of controls by reviewers so that they contain all relevant information concerning the prior literature. In all, the sample includes the five general-interest journals traditionally considered as the "top 5" (t5), i.e. American Economic Review, Econometrica, Journal of Political Economics, Quarterly Journal of Economics and Review of Economics Studies, as well as 11 renowned special-interest or field journals. The corpus excludes proceedings papers, comments, articles of less than three pages, book reviews, bibliographical items, articles without references and without abstracts, editorial material, letters, and corrections. The WoS database is merged with the IR database using the title of the article, the journal of publication, and the authors' last names. Because about 40% of authors' appellations on WoS consist of initials and last names, the authors' full names were validated by the *Cited*

Reference API.[9] Overall, the merged database contains 24,033 papers and their associated information.

Figure 1 shows that the number of publications in the sample has increased six-fold between 1991 and 2018. Figure 2 demonstrates that the number of pages per article has also increased over the same period. Specifically, the 95^{th} percentile of the page distribution was 40 in 1991 and it reached 55 in 2018. The number of authors has increased over time as well, from 1.7 authors per paper on average in the 1990s to about 2.2 authors per paper in the 2010s. The average number of authors over the entire sample period is approximately 2.[10]

Fig. 1. Number of articles over time

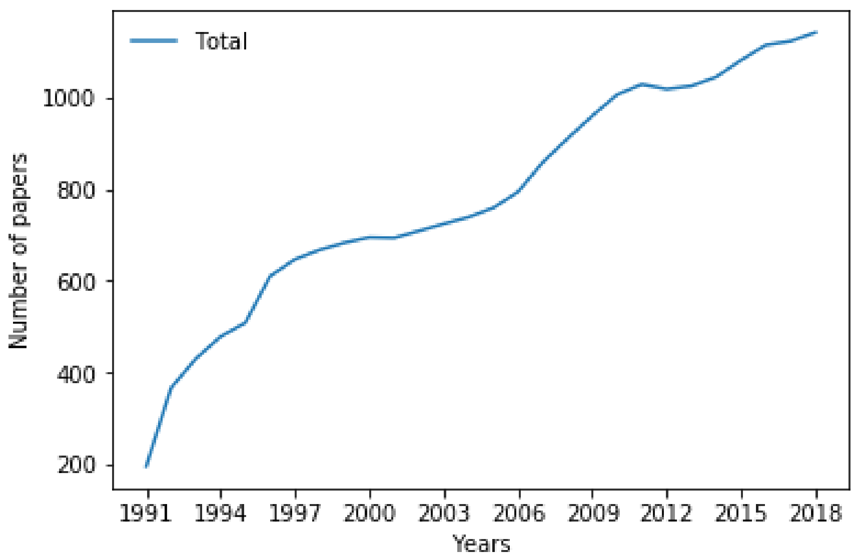

The figure shows the evolution of the number of articles over time from 1991 to 2018. The sample included publications in the 16 economic journals retained in this analysis: American Economic Review, Econometrica, Journal of Econometrics, Journal of economic growth, Journal of economic literature, Journal of economic perspectives, Journal of Economic Theory, Journal of Finance, Journal of Financial Economics, Journal of International Economics, Journal of Labour Economics, Journal of Monetary Economics, Journal of Political Economics, Quaterly Journal of Economics, Review of Economics Studies, Review of Financial Economics. The selected papers exclude proceedings papers, comments, articles of less than pages, books reviews, bibliographical items, articles without references and without abstracts, editorial material, letters and corrections.

[9]The “complete" names obtained by Rcrossref are also sometimes incomplete (initial of the first name + last name) or correspond to the names of the first authors in alphabetical order. For example, for a paper written by Abhijit Banerjee and Esther Duflo , Rcrossref could give “Banerjee A. and Duflo, E." or “Banerjee A." as authors’ names.
[10]Those findings are in line with Card et al. (2013).

Fig. 2. Distribution of the number of pages per article

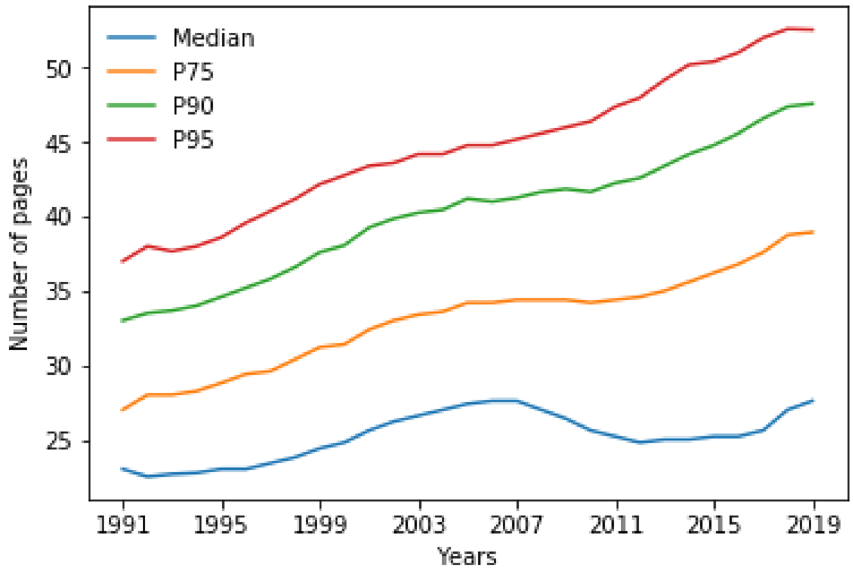

The figure shows the distribution of the number of pages per article over time since 1991 to 2018. The blue line is the median; the orange line is the 75th percentile; the green line is the 90th percentile and the red line is the 95th percentile.

Second, a set of automated web-crawling algorithms are designed for collecting the following information from WoS and IR for each paper in the corpus: The title, the abstract, the keywords,the JEL codes, the references, the journal of publication, the date of the paper's publication, the names of the authors, the institutions where the authors are employed, and the number of citations. The references and the number of citations are two important variables for the analysis, so a bit more detail is given below about their collection.

The terminology "citations" refers to the number of times a paper is cited by subsequent papers. Citations can also be collected from Google Scholar (GS). However, for the same paper, the number of citations in WoS and in GS is often different. This is because WoS considers citations that occur post-publication and only from other published papers, whereas GS counts citations already from the working stage of a paper and from any other papers, documents, or articles. In what follows, consistent with the fact that the year of publication is used as the reference year, the preferred specifications will use WoS as the main source of citations.

The terminology "references" refers to the bibliographical references of the related literature provided by each published paper. One difficulty here is that the references provided in WoS are often incomplete. However, WoS provides the digital object identifier (DOI) of each published paper. Hence, an algorithm is built to link each unique DOI to all other information about this DOI available in IR, including a complete list of references. To control for the fact that some articles could be quoted while unpublished, an additional web crawling algorithm was designed, allowing for recovery of the title of the unpublished paper at the moment of citation. On net, the resulting database contains 914,371 references with an average of 38 references per article. This is total, not restricting to top 16. Some of the references are to papers published in journals lower than top 16. Those are not considered. In the end, about 30% of the references (up to 50% in recent years) are to top 16 papers, and these are the ones used for comparisons.[11]

Table 1 shows the average number of citations at different horizons, using the connections between references and articles. The average of number of citations is 9.2, but the distribution is highly skewed. The standard deviation is 18 and the 95th percentile is 36.[12]

Table 1. Citations of Published Papers

	Full sample	Male	Female	Mixed	Unknown
	(1)	(2)	(3)	(4)	(5)
	Overall database				
Average	9.2	9.94	6.6	7.91	2.88
Standard deviation	18	20.51	10.8	14.2	5.54
Median	3	4	3	3	1
75th percentile	10	11	8	9	3
90th percentile	23	25	18	20	8
95th percentile	36	39	27	32	13
	WoS database				
Average	79	84	61	72	38
Standard deviation	196	209	102	162	158
Median	27	29	26	26	12
75th percentile	77	82	73	74	35
90th percentile	190	201	161	177	76
95th percentile	307	325	230	278	116

The table shows the mean value of the number of citations. Most recent citation update was in January 2019.

Third, the authors' institutional affiliations are classified into three categories, based on internationally acknowledged rankings of economics departments and organizations: top tier (H-group, rankings 1-10), middle tier (M-group, rankings 11-19), and lower tier (L-group, rankings 20-30). For multiple coauthors, the paper's affiliation is taken to be the affiliation

[11]In the online appendix, the database is extended to around 100 journals.

[12]Table 8 in the appendix shows that after five years, more than 25% of the papers are still at zero citations.

of the coauthor at the highest-ranked institution. The details of the ranking procedure are provided in the appendix. Papers written by authors in H-group institutions account for 32% of the sample, those written by authors in M-group institutions account for 46% of the sample, and authors in L-group institutions account for 21% of the sample.

Fourth, the gender of the authors is determined via a combination of automated algorithms and hand-collection efforts. First, a couple of probabilistic Bayesian algorithms for gender attribution are employed, namely the *Genderize.io* and *gender-guesser* in-built applications of the programming language Python. The algorithms are based on an in-built large databases of names collected from the US census, from international dictionaries, and from social media. So the algorithm yields the probability that a certain first or last name is associated with the male or female gender. This library of names was then augmented in three ways. First, via the merging of a database of inventors' names from the World Intellectual Property Organization (WIPO), for an additional eight million names, as well as the merging of the IR list of names of the top 10% of female economists.[13] Second, via specialized algorithms for web searches related to particular types of names. For example, in cases where the authors' native country can be determined, an algorithm was designed for finding country-specific name-gender probabilities. For instance, in the general sample, the name "nicola" is identified as female with probability 70%, whereas in Italy it is identified as male with probability 99%. Third, via hand collection and first-party verification. For example, for the 200 most cited authors identified as female and for the 200 most cited authors identified as male, the genders were verified by visual inspection of authors' personal websites and from other publicly available sources. When each of the relevant algorithms identifies the gender associated with a particular name with probability higher than 80%, then the author's gender is considered as determined. For names that were assigned a probability lower than 80% by at least one of the relevant algorithms, a manual search was conducted to recover the author's gender, based on personal websites, web sources that refer to the authors using their pronouns, and articles including authors' photographs. On net, 85% of the total 22,053 authors' names were identified as male or female. The rest were related to unisex first names, such as "taylor", or to first names that consistently appear initialized, and for which no other information was retrieved about the author.

Fifth, the "gender composition" of each team of coauthors is identified. In particular, four categories of gender composition are defined. A paper is identified as male-authored, if all the coauthors are men. A paper is identified as female-authored, if all the coauthors are women. A paper is identified as mostly-male-authored, if most of the coauthors are

[13]IR, January 2019, https://ideas.repec.org/top/top.women.html

Fig. 3. Female authors

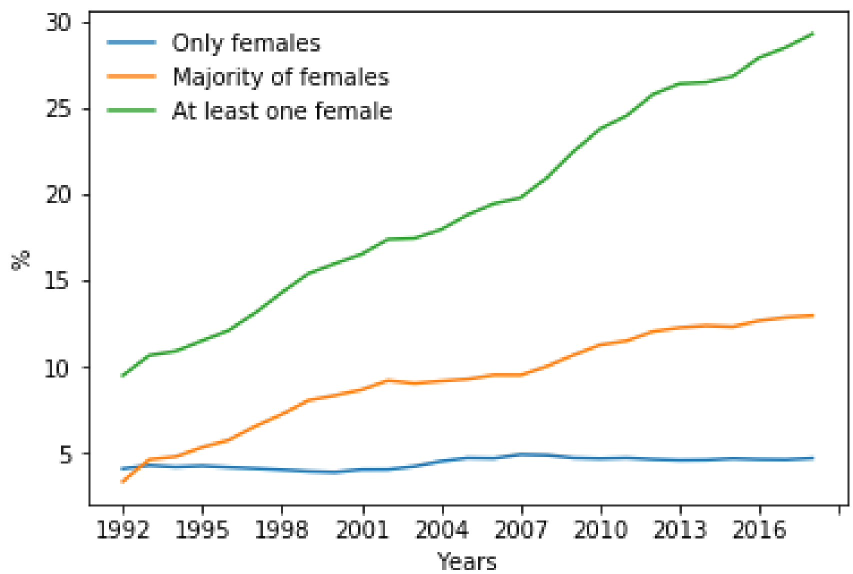

The figure shows the evolution (five-year moving averages) of the share of articles with at least one female authors for a given year (green line), the share of articles with at least one female authors with a majority of female authors (orange line) and only female authors (blue line). The share of a certain category is the total number of articles in this category (at least one female author, majority of female author, only female authors) over the total number of articles.

men. A paper is identified as mostly-female-authored, if most of the coauthors are women. Almost 75% of the articles in the database are male-authored, about 5% are female-authored, and the remaining are authored by mixed-gender teams. Articles with at least one female author (sole- or team-authored) make up only about 20% of the database. Figure 3 presents the evolution of the share of papers with at least one female author. In the early 1990s, papers with at least one female author constituted only 10% of published papers, whereas in recent years this number is closer to 30%. Similarly, the share of mostly-female-authored papers has increased over time. However, the share of female-authored articles has remained constant since 2010. The gender composition of the authors differs systematically across fields. For example, in labor economics and in the economics of education, about 8% of the papers are female-authored and about 23% have at least one female coauthor, compared to 3% and 15%, respectively, in the fields of theory, finance, and macroeconomics.[14] Papers written by mixed-gender teams tend to include more references than papers written by solo-gender teams, whether male or female, with the difference statistically significant at the 1% level. The accumulation of citation is different depending on the gender of the papers. In

[14]These findings are similar to those in Card et al.(2019).

fact, male-authored papers accumulate citations more quickly, compared to female-authored papers papers. Papers written by mixed-gender teams accumulate citations at rates similar to those for male-authored papers, at least between the first five years after publication. Further, on average across fields, about 75% of citations are for prior papers (related or not) that are "male" (the blue bars), where male indicates an all-male team (solo or all-male authors). Most of the remaining references are to prior published papers written by mixed-gender teams (the grey bars). Only a tiny fraction of the references are for papers written by female teams (solo or all female, the orange bars). This fraction is marginally higher in the fields of labor, education, and IO.

Overall, the different figures in this section present some summary statistics of the corpus and the main variables. Figure 1 shows that the number of articles published in the top 16 economics journals increases over time, from about 200 in 1991 to over 1000 in 2018. Most of this evolution is due to the addition of journals to the database over time.

Figure 2 shows that the top-skewness of the distribution of the number of pages per published article increases over time. In particular, the median number of pages remains constant at about 25, whereas the 75th (90th) percentile increases from 27 (34) to 37 (46) pages over the sample period. The number of authors per paper remains roughly unchanged, on average at two authors per paper.

Figure 3 shows that the fraction of published economics papers with only female coauthors (solo or all-female team) has remained stable over time at 5%. By contrast, the fraction of papers with at least one female author (green line) has increased from 15% in 1991 to almost 30% in 2018. In addition, the fraction of mostly-female authored papers has increased from 5% to about 12%.

2.4. Similarity and omission indexes

In this section, two main indexes are constructed, using the techniques of textual analysis, natural language processing, and unsupervised machine learning. The first index, termed the “similarity" index, determines how similar two papers are in terms of context. This index is basically a linguistic distance metric that allows to evaluate how closely related two papers are, based on their subject matter. The similarity index thus indicates which prior papers should be cited by a current paper, based on language and thematic similarities. If gender were irrelevant, then, all else equal, this similarity should be the only factor shaping the references or citations of past papers in a current paper. The second index, termed the "omission" index, captures the number of papers that are omitted from the references of a current paper, when they should be cited, according to the similarity index.

2.4.1. Natural language processing

Each paper is linked to the chronological sets of pre-existing and subsequent papers using commonalities in the topical content of each pair of papers. In turn, the topical content is culled from titles, abstracts, keywords and, in more advanced analysis, from the text of the paper, especially from the introduction and the related literature. These so-called textual data are cleaned and then taxonomized into sets of words. For example, cleaning involves dropping words that appear very frequently across papers, and therefore cannot be used to determine the degree of similarity across papers, such as "the" or "and".[15] Taxonomy encompasses, among other things, the attribution of parts of speech to each word that is not dropped during the cleaning stage.

The set of words includes individual words as well as word expressions and collocations or n-grams. Collocations are basically combinations of multiple words and they are frequently occurring in many fields of economics. For example, "public debt" is a bigram and "capital income taxation" is a trigram. These collocations are identified via unsupervised machine-learning, which means that they are allowed to automatically evolve over time, following the evolution of the language used by economists. For example, the collocation "idiosyncratic income risk" became more frequent after the seminal 1994 paper by Aiyagari, while the collocation "unconventional monetary policy" started being used after the onset of the 2007 financial crisis in the US.

2.4.2. Term frequency-Inverse document frequency

The Term frequency - Inverse document frequency (TFIDF) is a metric often used in machine learning to identify the relative frequency of a word in a corpus or collection of documents. The term frequency (TF) component gives the frequency of each word in the set of words in each document. Specifically, the TF is the ratio of the number of times the word appears in a document over the total number of words in that document. Clearly, the TF increases as the number of occurrences of the word within the document increases. For each word w in each paper p, the TF is therefore computed as:

$$TF(w,p) = \frac{Card(w \in p)}{Card(p)} \tag{2.4.1}$$

[15]The Appendix provides details on the data-cleaning process.

where $Card(w \in p)$ is the number of times the word w appears in paper p, and $Card(p)$ is the cardinal of p or the number of words in paper p.

The inverse document frequency (IDF) is defined as the logarithm of the inverse ratio of the number of documents in which a word appears over the total number of documents in the corpus. Let C be the corpus or the set of all documents in the database and $Card(C)$ the cardinal or the number of papers in the corpus C. The IDF is then computed as:

$$IDF(w) = -\log\left(\frac{\sum_P \mathbb{1}_{w \in P}}{Card(C)}\right) \tag{2.4.2}$$

Thus, the words that appear in every document will have IDF = 0, whereas the words that occur less frequently in the corpus will have a high IDF, because they are more informative for assessing similarities across documents. For example, for the similarity comparison of two papers, words like “taxation" or “bayesian" will be more useful than words like “paper" or “model", hence they should enter more prominently into the similarity calculation.

The TFIDF of a word is then the product of the TF of the word times the IDF of the word, or $TFIDF = TF \cdot IDF$. A low TFIDF may indicate that the word appears infrequently in the document (low TF) or that it is a very common word that appears in many documents (low IDF). A high TFIDF indicates that a word appears relatively frequently in one document but it does not appear in most other documents of the corpus, hence it is crucial for the content of this particular document.

However, the traditional IDF does not take into account the evolution of the natural language or the introduction of new vocabulary or terminology over time. Following Kelly et al.(2018), the IDF is therefore adjusted to give a higher weight to newly-introduced concepts. Let $\tilde{C}(t)$ be the corpus of papers before a given date t. Then, the adjusted IDF (ATFIDF) is given by:

$$AIDF(w,t) = -\log\left(\frac{\sum_{p \text{ prior to } t} \mathbb{1}_{w \in p}}{Card(\tilde{C}(t))}\right) \tag{2.4.3}$$

Thus, the AIDF is the logarithm of the inverse ratio of the number of papers published before paper p in which a word appears over the total number of papers published before paper p. Basically, it is a retrospective version of the IDF. Because the AIDF varies with time and across words, it attributes importance or weight to each word depending on the degree of utilisation of the word over time. As a result, it reflects the state of the art or the frontier of innovation up until the arrival of each new research paper. Clearly, the $ATFIDF = TF \cdot AIDF$.

2.4.3. Similarity index

The similarity index, which measures the textual or conceptual similarities across two papers, is basically a cosine similarity distance measure. The cosine similarity is a measure of similarity between two non-zero vectors of an inner product space. It measures the cosine of the angle between the vectors, where the cosine of a 0-degree angle is 1, and the cosine of a 90-degree angle is 0. Each of the two papers to be compared is represented by a vector based on the ATFIDF of each word. Let U and V be the respective vector representations of papers p and p':

$$U = [tfaidf(w_1,p,t), tfaidf(w_2,p,t), ..., tfaidf(w_n,p,t)]^T$$
$$V = [tfaidf(w_1,p',t), tfaidf(w_2,p',t), ..., tfaidf(w_n,p',t)]^T$$

Then, the cosine similarity, $\lambda_{p,p'}$, is the angle between these two vectors:

$$\lambda_{p,p'} = cos(p,p') = \frac{U.V}{\|U\|\|V\|} \tag{2.4.4}$$

Clearly, $cos(p,p') \in [0,1]$. Papers that are very similar tend to use the same words with the same frequency, so their vector representations have a trigonometric angle closed to 0 and, therefore, the cosine similarity measure will take a value close to 1. At the opposite end, papers that have no common concepts will yield a cosine of around 0.[16]

2.4.4. Omission index

Next, the relevant prior literature of paper p, denoted by $\mathcal{P}_p$, is defined as the n-papers, denoted by p_i, with the highest cosine:

$$\mathcal{P}_p = \{p_1, p_2, ..., p_n\} \text{ such that for } i \in [0,n], \lambda_{p,p_i} > \lambda_{p,p'}, \quad \forall p' \in C \backslash \{p_1, p_2, ..., p_n\} \tag{2.4.5}$$

The preferred specification uses $n = 5$, thereby examining which out of the top-five most related prior papers are omitted from the references of a current paper. However, the qualitative results are robust to higher values of n. Next, the omission index for the comparison between similar papers p and p', of which p' was published first, is a binary variable, denoted by $omit_{p,p'}$, which takes the value of 1 if paper p cites paper p' in its references, and 0 otherwise:

[16] In the analysis that follows, similarities below 0.05 are set at zero, and zeros are dropped, so as to reduce the computational burden and to remove extremely limited similarities.

$$omit_{p,p'} = \begin{cases} 1 \text{ if p does not cite p' conditional on p' in } \mathcal{P}_p \\ \\ 0 \text{ if p cites p' conditional on p' in } \mathcal{P}_p \end{cases}$$

This index therefore determines if the relevant prior literature is included in the references of a current paper or not, and to what extent.

2.4.5. Descriptives

Figure 4 plots the empirical CDF of the pairwise similarity index in panel (b) and the relationship between the similarity index and the probability of being cited in panel (c). As demonstrated in panel (c), papers with high relative similarity are more likely to be linked by a citation. In other words, the probability that one paper cites another is increasing in the similarity index.

Moreover, about 55% of papers do not cite all of their top 5 most similar, and therefore relevant, prior papers. In other words, more than half of all published papers omit from their references at least one of the most related prior contributions to the literature. This finding is stable across alternative specifications of the similarity index and of the omission index, as shown in detail in the appendix. Panel (b) retains the number of related papers to $n = 5$, but imposes a tougher restriction on the cosine similarity. As can be seen, the results are robust when the requirement of more similar prior literature is imposed. Panel (c) shows that the results are robust to the increase in the number of papers in the related prior literature, from $n = 5$ to $n = 10$. Finally, panel (d) shows that the results are unchanged when both the number of related papers and the degree of similarity across related papers are required to be higher. Overall, almost 60% of papers cite one of their most related prior papers, while 30% of papers cite two related prior publications. It is worth noticing that due to the huge economic literature, the authors must choose between the articles which they wish to cite and those which they will not cite. However, the problem arises when this choice systematically excludes a category of people based on criteria, such as gender. This is what we will investigate in the next section.

2.5. Omission and gender: Descriptive analysis

In sum, the average published paper omits almost half of related prior published papers. This section presents a descriptive investigation of the potential gendered pattern of these omissions of prior literature.

Fig. 4. Pairwise similarity

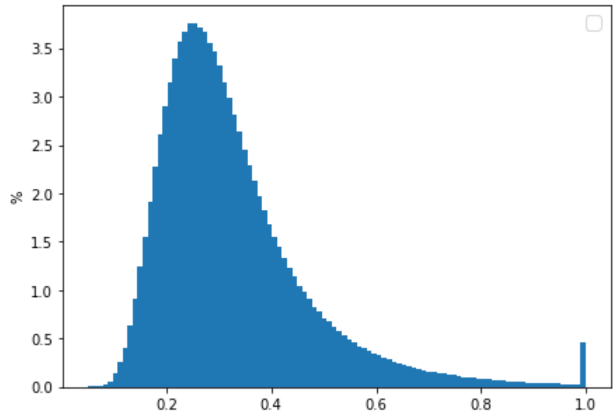

(a)

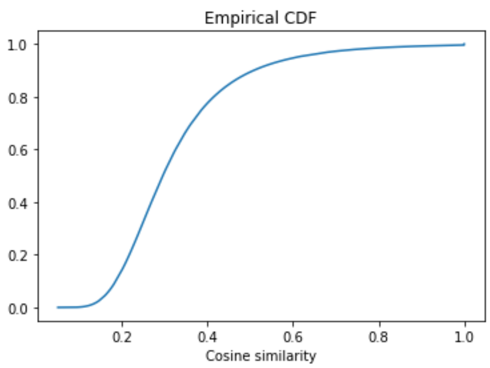

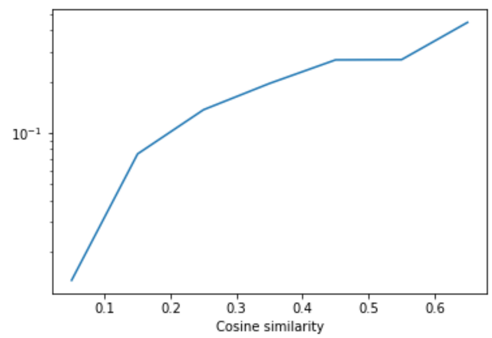

(b) **(c)**

Panel (a) plots the distribution of the relative cosine. For paper P and P_{max} such as: $P_{max} = argmax_{P' \in C} cos(P,P')$, for any given article P' in the database, the relative cosine of paper P and paper P' is defined as:$\tilde{\lambda}_{p,p'} = \frac{\lambda_{p,p'}}{\lambda_{p,p_{max}}}$. Panel (b) shows the empirical cumulative distribution function (CDF) of the relative cosine. Panel (c) shows the conditional probability that paper p cites a paper p' as a function of the similarity score between those two papers. The computation excludes similarity score lower than 0.05.

2.5.1. Gender and intersectionalities

Figure 5 plots the evolution of the gender composition of the authors of the most similar papers.

Panel (a) shows that the fraction of papers whose most related literature, as determined by the similarity index, includes at least one female author has increased from 40% in 1991 to 70% in 2018. However, the fraction of papers that completely omit any of those related papers has remained flat at 30% and the fraction of papers that cite some of those related papers has increased only from 10% to 30%. Panel (b) shows that the fraction of papers whose

Fig. 5. prior literature with female authors

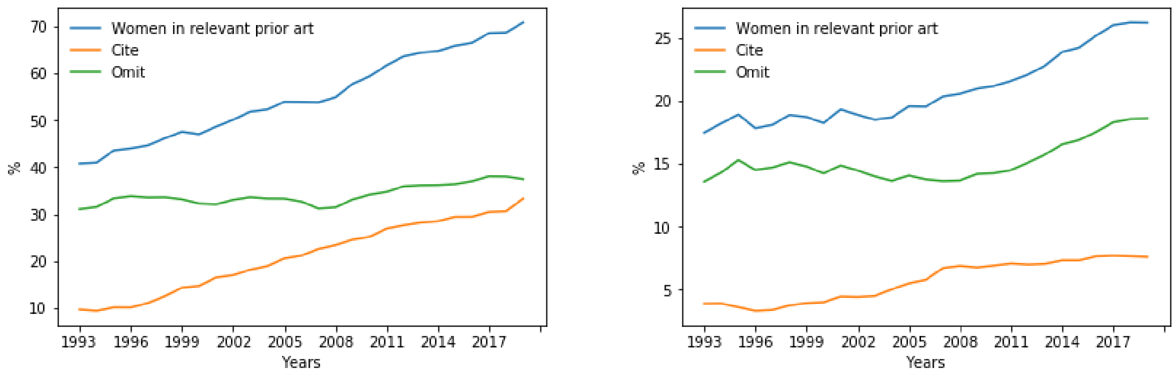

(a) At least one female author

(b) Majority of female authors

Panel (a) shows the fraction of papers with at least one female in their relevant prior literature (blue line), the fraction of papers that cites at least one paper with a female author (orange line), the fraction of papers that does not cite at least one paper with a female author (green line). Panel (b) shows the fraction of papers that have at least one paper with a majority of females in their relevant prior literature (blue line), the fraction of papers that cites at least one paper with a majority of females in their relevant prior literature (orange line), the fraction of papers that does not cite papers with a majority of females in their relevant prior literature (green line). The curves are five-year moving average and normalize by the 1993 values.

most related literature includes majority-female authored papers increased from about 15% to about 25% over the sample period. However, the fraction of papers omitting these related papers has increased by 2 percentage points to 17% over the sample period. The fraction of papers citing some of these related papers has also increased by 2 percentage points and stands at 7% in 2018. Essentially, despite the increasing and substantial representation of women authors in the related literature since 1991, the degree of their complete omission from the references remains high and unchanged over the past three decades, whereas the degree of their partial omission has declined a bit, but by no means as fast as the increase in representation.

Figure 6 presents the gender distribution of the references, as a function of the gender structure of the current paper's author team. The vertical axis is the gender of the citing team. The horizontal axis is the probability of a related paper being cited, when the authors are male (the blue bars, solo or all male team), female (the orange bars, solo or all female team), or mixed-gender (the grey bars). Comparing the orange part of the bottom and middle bars, we can see that women are 3 times more likely to be cited by women authors than by men authors.

Furthermore, in almost all fields, the probability of being omitted, when relevant, is higher for women than for men. The only exception is IO, where men are omitted with

Fig. 6. Fraction of references to a given gender by gender of the citing paper

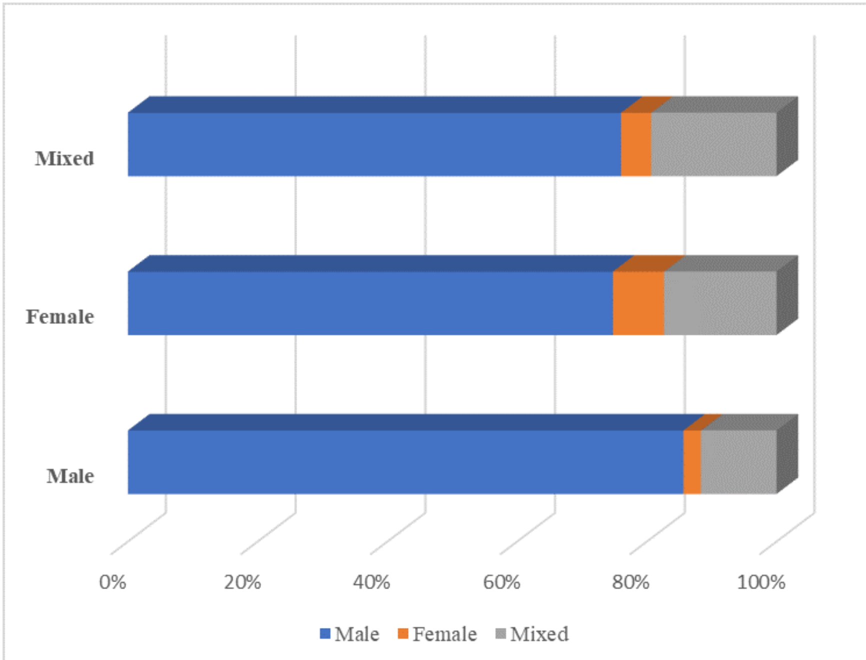

This figure presents the share of references that are attributed to a given gender depending on the "gender" of the citing paper. *Male* designed paper written by only men; *female* designed paper written by only women; *Mixed* designed paper written by a team of females and males.

higher probability than women. The omission of papers written by women is on average 30% bigger in theoretical fields compared to applied fields. In fact, the gap between articles written by men and articles written by women is almost 50% bigger in theory than in applied economics. In applied economics, relevant papers written by women are two times more likely to be omitted than cited, and the same is true for related papers by men. While the odds for men do not change in theoretical fields, the odds for women worsen, with related papers written by women now three times more likely to be omitted than cited.

Moreover, the existing number of citations of an article can be a determining factor for the choice of this article. The number of citations of a paper up to the date of publication of the article that should cite it is used as a proxy of the existing citation number by the time an article chose to quote another one. Citing papers tend to quote papers with the highest number of citations among their relevant prior literature. On average, the papers in the relevant prior literature that are cited have 40% more citations than those left out. Moreover, in general papers cited (relevant or not) have much more citations than those left out.

Fig. 7. Likelihood of citation by gender of the citing paper

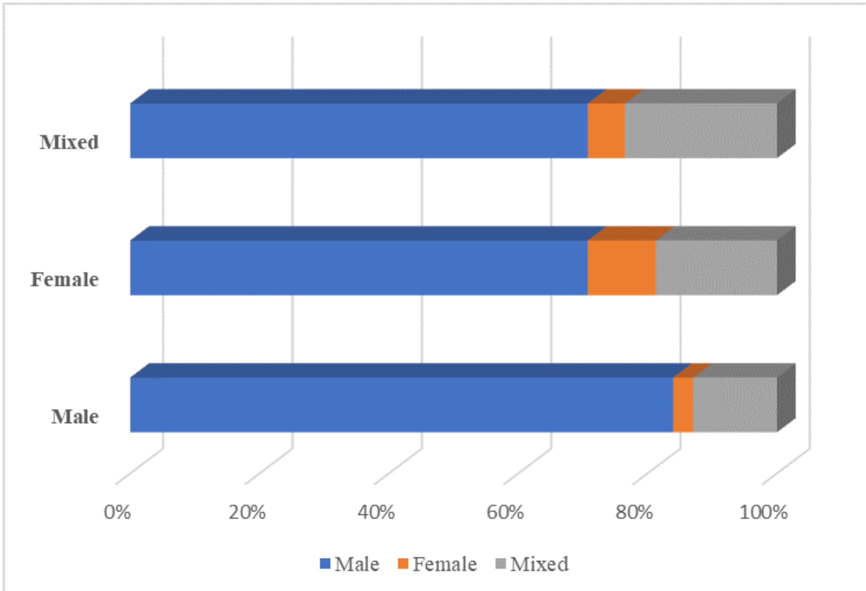

This figure decomposes the likelihood of being omitted compared to the one of being cited with respect to the gender of the cited/omitted and the gender of the citing paper. *Male* designed paper written by only men; *female* designed paper written by only women; *Mixed* designed paper written by a team of females and males.

Figure 7 shows the probability that a paper written by authors of a gender type X cites another paper written by authors of a gender type Y. Male authors papers have the highest probability to be cited by any gender group. This makes sense as 70% of the papers in the database are from male authors. Male authors tend to cite more papers written by other male authors (80.2%), papers written by mixed gender teams (13%) and paper written by women (3%). The probability that a male author cites another male author is 15% higher than the probability that a male author is cited by a woman author or a mixed gender team. Surprisingly, papers written by authors from a gender type X has the highest probability to be cited by papers written by authors from the same gender type. Therefore, with a 10% probability, papers written by women cite other papers written by women. This is three times bigger than the probability that a paper written by male authors cites a paper written by female authors and two times bigger than the probability that a paper written by mixed author team cites a paper written by female author. Unambiguously, articles that refer the most to papers written by women are articles written by women. Men authors refer refer the least to papers written by women.

2.5.2. Gender distribution of omissions

An article may not refer to papers written by women or any other papers in a group of closely related literature but could have a completely different gender structure in terms of the references it is citing. This section aims to conduct a counterfactual analysis by looking at the gender structure of the references of all the articles in the database. For example, suppose paper p is citing N other prior papers in the data. From this, we can obtain the observed distribution of references across authors' genders. Next, take the N most related papers to paper p, as determined by the similarity index. From this, we can obtain the "target" distribution of references across authors' genders, i.e. the references that should have occurred, had similarity been the only determinant and had gender not been a factor. Hence, actual citations are "accurate" when the observed distribution equals the target distribution of references. Figures 5 and 9 compare the gender distribution of the actual references to the target distribution of references.

Figure 5, panel (a), compares the actual distribution versus the target distribution. A positive difference means that the actual distribution of a certain gender type is higher than the target distribution of this gender type. A negative difference means that the actual distribution of a certain gender type is lower than the target distribution of this gender type. The distribution of the difference for females has a much greater weight in 0 compared to that of males. It also tends to be more oriented to the left, while the distribution of the difference for males tends to be more oriented to the right. For papers written by male authors, the average difference (actual-target) is about 3%, implying that papers written by male authors are more cited compared to what is suggested by the target distribution. For papers written by women, the average difference is at -1.5%, meaning that papers written by women are less cited compared to the target distribution. In general, for a negative difference between the actual and the target distribution, papers written by women are more likely to be totally omitted than cited to a lesser extent, compared to the target distribution. By contrast, papers written by men are more likely to be cited to a lesser extent than totally omitted from the references, compared to the target distribution. Panel (b) complements this argument by plotting the distribution of the difference of both differences. The histogram has more weight on the left. It means that if males are *over-cited*, they are more *over-cited* than females. Similarly, when they are *under-cited*, they are less *under-cited* than females.

Figure 9 plots the difference between the actual and the target distribution over time for different genders. Roughly, 75% of the papers exhibit a negative difference for females, i.e. they should cite more papers written by women than they actually do. This number

Fig. 8. Experiment 2: Actual distribution-Target distribution (1/2)

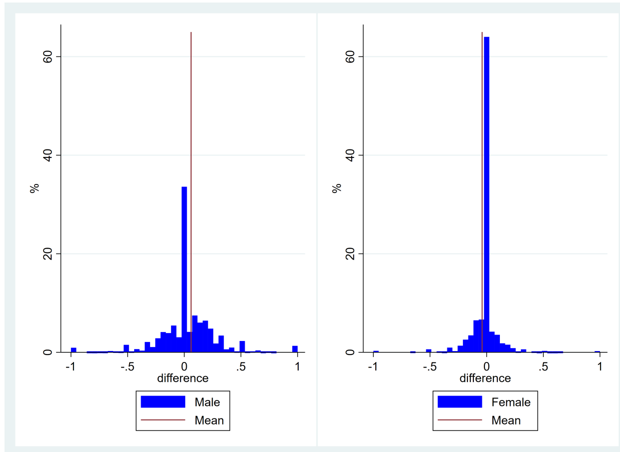

(a) Males versus Females

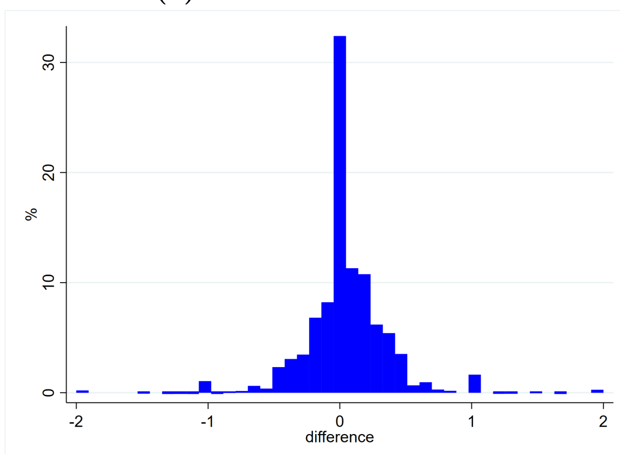

(b) Difference Males-Difference Females

This figure plots the difference between the actual distribution and the target distribution of gender type g. For each paper p, the actual distribution of gender is the share of papers in its references belonging to each category of gender (only males, only females, mixed gender). The target distribution of gender is the share of the closest papers (in the sense of the relative cosine) in its prior literature belonging to each category of gender (only males, only females, mixed gender). For each paper, the difference between the actual distribution of gender g and the target distribution is taken. Panel (a) plots the distribution of this difference for males versus females. A positive difference means that the actual distribution of a certain gender type is higher than the target distribution of this gender type. A negative difference means that the actual distribution of a certain gender type is lower than the target distribution of this gender type. Finally, a null difference means that the actual distribution of a certain gender type is the same as the target distribution of this gender type. Panel (b) takes the difference of the difference for males and for females. A positive double difference means that males are more *over-cited* or less *under-cited* than females.

is only half as large for papers written by males. Overall, there is a slight increase in the share of articles that have a positive difference (actual > target) for female citations, from

Fig. 9. Experiment 2: Actual distribution-Target distribution (2/2)

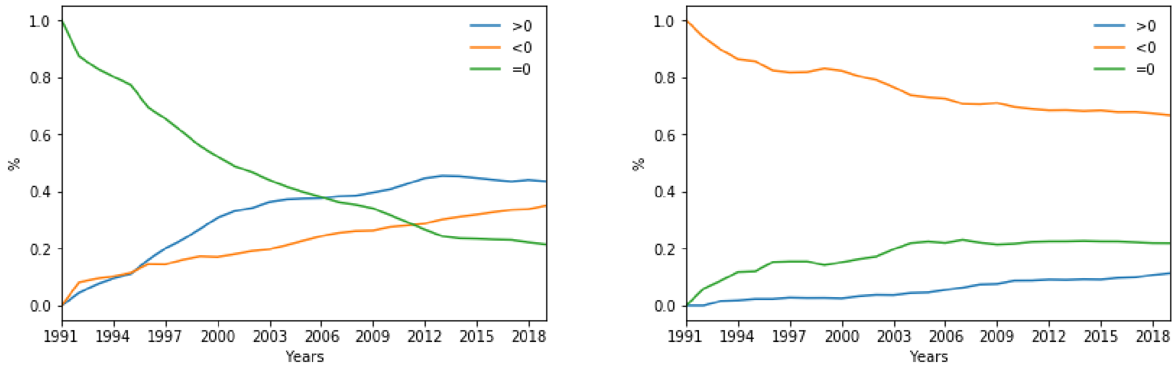

(a) Males **(b)** Females

This figure plots the difference between the actual distribution and the target distribution of gender type g. For each paper p, the actual distribution of gender is the share of papers in its references belonging to each category of gender (only males, only females, mixed gender). The target distribution of gender is the share of the closest papers (in the sense of the relative cosine) in its prior literature belonging to each category of gender (only males, only females, mixed gender). For each paper, the difference between the actual distribution of gender g and the target distribution is taken. Panel (a) and (b) plot the evolution of the share of papers with a positive, negative and null difference with respect to gender type male, female and mixed gender teams.

almost 0 in the early 1990s to 10% in 2018. However, there is a noticeable increase in the share of articles that have a positive difference for male citations, from less than 10% in the early 1990s to almost 40% in 2018. The lack of citations of papers written by women is roughly constant over time. In other words, as shown in panel (a) of Figure 9, since 1991, there has been a constant decline in the accuracy of citations for papers written by men (the green line). Instead, about 40% of publications (blue line) over-cite male papers (solo or all male authors), meaning that the observed number of citations of male papers is higher than the target, which is based on similarity alone, and about 30% of publications (orange line) under-cite male papers. By contrast, as shown in panel (b), for papers written by women (solo or all female teams), the accuracy remains roughly constant over the 1991-2019 period, with about 20% of published papers (green line) having accurate references. The fraction of publications over-citing female papers is basically flat at zero (blue line), whereas the fraction of papers under-citing women has declined a bit over the period, but still stands at the very high level of 80% (orange line).

2.6. Omission and gender: Empirical analysis

The previous section presented an empirical analysis suggesting that there is a potential role of gender for determining omissions of prior related art from the references of a current paper. This section proceeds with a rigorous empirical analysis that controls for a number of factors that may influence the observed citation patterns in economic publications.

2.6.1. Benchmark probability model

Assume that paper i was published in year t and that paper j published in year t'. Assume that paper j belongs in the relevant prior literature of paper i, according to the similarity index. Let $gend_j$ be the variable that defines the gender of paper j's authors. This is the variable of interest. To make the papers as similar as possible, the estimation equation includes a wide range of control variables. Let Z_j^1 and Z_i^2 be sets of controls for papers j and i, such as the journal of publication, the authors' affiliation, the number of authors, and the number of references. Controlling for the journal is a way of conditioning on the quality of the paper. Similarly, the affiliations of the authors make it possible to exclude the fact that a potential bias could be because women have less visibility if they are more likely to be affiliated to lower-ranked institutions. The number of authors can also be a confounding factor. Indeed, if women are more likely to write articles alone than men, this could be misleadingly captured by the gender variable. The number of references of an article makes it possible to exclude the fact that the bias is systematic with articles with few references which will, therefore, choose just a handful of articles to be cited. Let $Z_{i,j}^3$ be a set of controls about observed commonalities across papers i and j, such as the primary field of study. Let $Z_{t,t'}^4$ be a set of controls about the year of publication of each paper. The number of years between the cited and the citing papers can also affect the likelihood of a paper being cited. The determinant of a paper omission are investigated given that this paper is similar to the citing one. Then, the probability of paper i omitting paper j, when it should have cited it according to the similarity index, termed $omit_{ij}$, is given by the following logit model:

$$omit_{it,jt'} = \beta_0 gend_j + \beta_1 Z_j^1 + \beta_2 Z_i^2 + \beta_3 Z_{i,j}^3 + \beta_4 Z_{t,t'}^4 + \epsilon_{it,jt'} \qquad (2.6.1)$$

where standard errors are clustered at the paper level.

The results of this estimation are presented in Table 2, for a number of different specifications and controls. The dependent or outcome variable is the probability of omission. It captures the probability that paper i cites paper j in the data, given that j is in the relevant prior literature of i, according to the similarity index.

Table 2. Relationship between omission and gender

	Outcome variable: Omission			
	(1)	(2)	(3)	(4)
f_j	0.193***	0.195***	0.148***	0.131***
	(0.043)	(0.045)	(0.045)	(0.045)
top5 j		-0.623***	-0.552***	-0.650***
		(0.018)	(0.019)	(0.020)
field		-0.923***	-0.923***	-0.861***
		(0.019)	(0.019)	(0.020)
Lag		0.048***	0.042***	0.056***
		(0.002)	(0.002)	(0.002)
G_{prior}		-0.230***	-0.189***	-0.174***
		(0.047)	(0.047)	(0.048)
Nref		-0.054***	-0.048***	-0.043***
		(0.001)	(0.001)	(0.002)
Authors				-0.015
				(0.012)
Institution of j FE			Y	Y
Institution of i FE				Y
Journal of i FE				Y
Year of publication of i FE				Y
Field FE				Y

This table shows the relationship between the omission and the gender of the omitted paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. f_j represents papers written by only women. The reference variable is *male*, which represents papers written by only men (the two other gender structure are added but not shown in the table to ease the reading. See appendix for more details). *Top 5 j* is binary and indicates if paper j is published in a top 5 journal or not. *field* is binary and indicates if paper i and paper j have the same primary field. *Lag* is the difference between the publication year of paper i and the publication year of paper j. G_prior is the share of paper written by at least one female author in the relevant prior literature. *Nref* is the number of references recovered from the database. *Authors* is the number of authors writing the paper. The equations are estimated using a logit model. The odd ratio for a variable is the exponential of its given coefficient. Standard errors are clustered by papers and reported in parentheses. The total number of observation is 117694. ($* = p < 0.10$, $** = p < 0.05$,$* * * = p < 0.01$

The variable "female" in the first row refers to an all-female author team (solo or multiple authors). The associated coefficient is the odds ratio that prior relevant paper j is omitted from the citations of paper i, when the author team of paper j is al female, compared to all male. Regardless of specification, the coefficient is tightly estimated and is about 20% and it is always highly statistically significant. In other words, the odds of being omitted from the references are 20% higher for papers written by all-female teams, compared to those by all male-teams.

The variable $top5j$ controls for whether a prior paper is published in a top five economics journal. The associated coefficient, which is tightly estimated at around -0.55 across specifications, shows that the odds of being omitted from the references are almost 55% lower, if the paper is published at a top 5 journal, rather than a non-top5 journal.

The variable $field$ denotes same primary field of the citing and cited papers. Its coefficient is tightly estimated at -0.9 and it is statistically significant across all specifications. It shows that the odds of being omitted from the references are 90% lower, if the paper belongs to the same primary field as the citing paper. Hence, belonging to the same field is a necessary but not sufficient conditions for being cited. However, if the papers belong to different fields, it is very unlikely that the prior paper will be cited.

The variable lag captures the interval between publication years for citing and cited papers. To coefficient is tightly estimated at 0.03 across specifications and it is always statistically significant. It shows that one additional year between publication dates increases the odds of being omitted by 3% on average.

The variable G_{prior} captures the gender structure of the relevant prior literature. Its coefficient is estimated in the interval [-0.2, -0.6] across specifications and it is always statistically significant. It indicates that a one unit increase in the share of papers with at least one female author in the relevant prior literature reduces the odds of being omitted by 40%, on average. Put differently, if relevant prior literature includes one more paper with at least one female author, this decreases the odds of being omitted by 8%.

The variable $Nref$ is the total number of references of a paper. Its coefficient is tightly estimated at -0.045 across specifications and it is statistically significant. It indicates that one additional bibliographical references reduces the odds of being omitted by about 4%.

Overall, the results for the main variables of interest are highly significant and quantitatively similar across specifications. This finding is robust to the inclusion of fixed effects for the field, for the institutional affiliation of the authors of i and j, and for the journal and year of publication of paper i.[17]

2.6.2. Two-sided gender

This section examines in more detail the role of the gender structure for citing and cited papers on the probability of omission. The dependent or outcome variable is the probability

[17] When there are multiple coauthors, the paper's affiliation is taken to be the affiliation of the coauthor at the highest-ranked institution. See appendix

of omission. It captures the probability that paper i cites paper j in the data, given that j is in the relevant prior literature of i, according to the similarity index. The controls now include a number of cross-gender variables:

$$omit_{it,jt'} = \tilde{\beta}_0 gend_j + \tilde{\beta}_1 Z_j^1 + \tilde{\beta}_2 Z_i^2 + \tilde{\beta}_3 Z_{i,j}^3 + \tilde{\beta}_4 Z_{t,t'}^4 + \tilde{\beta}_5 \cdot gend_i + \tilde{\beta}_6 \cdot gend_i \cdot gend_j + \epsilon_{it,jt'} \quad (2.6.2)$$

In an ideal setting, the interaction effect reflects the difference-in-differences in the relative omission bias for a paper j written by an all-female team versus a paper j written by an only-male team, when paper i is written by all-women relative to when paper i is written by all-men.

The results are presented in Table 3. The variable f_j indicates only female authors in the cited paper (solo or all-female team). The variable f_i indicates only female authors in the citing paper (solo or all-female team). The variable $A1f_j$ indicates at least one female author in the cited paper. The variable $A1f_i$ indicates at least one female author in the citing paper. The variables $F_j \cdot F_i$, $F_j \cdot A1F_i$, etc. are cross-variables for citing and cited papers. The specification includes a dummy for when the gender structure is the same in i and j, namely it takes the value 1 when both papers have all female, all male, or at least one female teams. The coefficient on this dummy is -0.017 and statistically significant at the 1% level. All other Z-controls are also included.

Take column (1), in which the citing paper i has an all female team, f_i, and paper j has an all female team, f_j.

First, the coefficient on f_j corresponds to $\tilde{\beta}_0 = 0.036$ and it is statistically significant. It means that having only male authors in citing paper i increases the probability to get omitted for an all-female relevant paper j by 3.6 percentage points (pp), compared to an all-male paper j. In other words, conditional on an all-male citing team, the probability of omission is 4pp higher for female papers than for male papers.

Second, the coefficient on f_i corresponds to $\tilde{\beta}_5 = -0.006$, which however is not statistically significant. It means that, for an all-male relevant paper j, the probability to be omitted is the same, regardless of whether citing team i is all-female or all-male.

Third, the coefficient on $f_i \cdot f_j$ corresponds to $\tilde{\beta}_6 = -0.106$ and it is statistically significant. It means that having only female authors in citing paper i decreases the probability to get omitted for a relevant all-female paper j by 10.6pp, compared to an all-male paper i. In other words, a complete change in the gender structure of the authors of i from male to female is associated with a substantial, statistically and economically, increase of almost 11pp in the probability of citation for relevant female-authored papers.

Table 3. Omission and two-sided gender

	Outcome variable: Omission			
	(1)	(2)	(3)	(4)
f_i	-0.006			-0.004
	(0.006)			(0.006)
f_j	0.036***		0.036***	
	(0.005)		(0.005)	
$f_i \cdot f_j$	-0.106***			
	(0.022)			
$A1f_i$		0.005*	0.005	
		(0.003)	(0.003)	
$A1f_j$		0.021***		0.021***
		(0.003)		(0.004)
$A1f_i \cdot A1f_j$		-0.063***		
		(0.007)		
$f_j \cdot A1f_i$			-0.073***	
			(0.012)	
$A1f_j \cdot f_i$				-0.058***
				(0.013)
N	77832	109055	94760	88725
R-sqr	0.08	0.08	0.08	0.08

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the gender of the citing paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. f_x represents paper x written by only women. $A1f_x$ represents paper x with at least one female author. All the specifications include controls for paper j published in a top 5 journal; paper i and paper j having the same primary field; difference between the publication year of paper i and the publication year of paper j; the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. The equations are estimated using a logit model. The table displayed the marginal probabilities. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

Column (2) presents the results for the case where paper i has at least one female coauthors, $A1f_i$, paper j has at least one female coauthors, $A1f_j$, and the gender-interaction term is $A1f_i \cdot A1f_i$. The results and the corresponding interpretations are qualitative similar to those in column (1). Column (3) presents the results of the gender structure $(f_i, A1f_j)$, with gender-interaction term $f_i \cdot A1f_j$. Column 4 presents the results of the gender structure

$(A1f_i, f_j)$, with gender-interaction term $A1f_i \cdot f_j$. Again, the results are qualitatively similar to those in column (1). On average, across all columns, $\tilde{\beta}_0 > 0$ and statistically significant, $\tilde{\beta}_6 < 0$ and statistically significant, and $\tilde{\beta}_5 = 0$.

2.6.3. Robustness

Tables 4 and 5 present the robustness of the results to the inclusion of a various potentially relevant controls. The dependent or outcome variable is the probability of omission, when relevant. It captures the probability that paper i cites paper j in the data, given that j is in the relevant prior literature of i, according to the similarity index. This table shows that the same overall qualitative picture for $\tilde{\beta}_0, \tilde{\beta}_5, \tilde{\beta}_6$ emerges when a number of robustness tests are performed. Column (1) of table 4 repeats the results from column (1) of Table 3, namely the results for gender structure $(f_i, f(j))$ with gender-interaction term $f_i \cdot f_j$.

Next, column (2) adds the cosine similarity to the set of controls. The coefficient on the cosine similarity is negative and statistically significant at -0.75. This means that an increase in the textual similarity across paper i and prior relevant paper j reduces the probability of omission of j by about 75pp. Incidentally, this confirms the use of the cosine similarity as an index of contextual proximity across papers. Furthermore, $\tilde{\beta}_0 > 0$ and significant, $\tilde{\beta}_6 < 0$ and significant, and $\tilde{\beta}_5 = 0$.

In fact, an increase in the similarity between i and j is associated with a reduction in the probability that i omits j from 75% to 95% on average. Specifically, one additional unit in terms of the similarity between i and j reduces significantly the probability that i omits j by 75% if j is a paper written by men and by only 60% if j is written by women. This effect is even larger (by 10pp) for papers with high level of similarity, compared to those with lower level of similarity.

Table 5 shows that switching from only males to only females in i reduces the probability to be omitted for paper j published in a top 5 (respectively non top 5) written by only women by 10 percentage points (respectively 10 percentage points). The omission bias is present in top 5 and in non top 5 journal.

Furthermore, switching from only males to only females in i reduces the probability to be omitted for paper j from a top tier affiliation written by only women by 10 percentage points. Overall, switching from only males to only females in paper j from a top affiliation, increases the probability to get omitted by 4 percentage points when i is written by only males. Switching from only males to only females in paper j from a top affiliation, reduces the probability to get omitted by 6 percentage points when i is written by only females. Similarly, switching from only males to only females in i reduces the probability to be omitted for

Table 4. Omission and gender: Robustness (1/2)

	Outcome variable: Omission		
	(1)	(2)	(3)
f_i	-0.006	-0.004	
	(0.006)	(0.006)	
f_j	0.036***	0.035***	-0.018
	(0.005)	(0.005)	(0.013)
$f_i \cdot f_j$	-0.106***	-0.096***	
	(0.022)	(0.022)	
cos_{ij}		-0.757***	-0.762***
		(0.01)	(0.01)
$cos_{ij} \cdot f_j$			0.134**
			(0.06)
N		77832	113280
R-sqr		0.107	0.11

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the effect of the similarity between two papers. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. f_x represents paper x written by only women. cos_{ij} is the value of the cosine between i and j. All the specifications include controls for paper j published in a top 5 journal; paper i and paper j having the same primary field; difference between the publication year of paper i and the publication year of paper j; the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. The equations are estimated using a logit model. The table displayed the marginal probabilities. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

paper j from a mid tier affiliation written by only women by 17 percentage points. Overall, switching from only males to only females in paper j from a mid-tier affiliation, increases the probability to get omitted by 6 percentage points when i is written by only males. Switching from only males to only females in paper j from a mid-tier affiliation, reduces the probability to get omitted by 11 percentage points when i is written by only females. Finally, switching from only males to only females in i reduces the probability to be omitted for paper j from a low tier affiliation written by only women by 6 percentage points (but the effect is not significant).

The results are also robust to alternative time periods or alternative time length between the citing paper and the omitted or cited papers.

Table 5. Omission and gender: Robustness (2/2)

	Outcome variable: Omission				
	Journal j		Institution j		
	(4)	(5)	(6)	(7)	(8)
	Top5	Non Top5	Top tier	Mid tier	Low tier
f_i	-0.009	0.001	-0.030**	0.002	0.007
	(0.01)	(0.006)	(0.013)	(0.01)	(0.011)
f_j	0.050***	0.032***	0.038**	0.061***	0.019
	(0.010)	(0.006)	(0.016)	(0.009)	(0.012)
$f_i \cdot f_j$	-0.093**	-0.098***	-0.102*	-0.171***	-0.061
	(0.039)	(0.025)	(0.056)	(0.043)	(0.05)
N	27735	50097	18959	25461	11497
R-sqr	0.13	0.10	0.12	0.09	0.09

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the gender of the citing paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. f_x represents paper x written by only women. All the specifications include controls for paper j published in a top 5 journal; paper i and paper j having the same primary field; difference between the publication year of paper i and the publication year of paper j; the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. The equations are estimated using a logit model. The table displayed the marginal probabilities. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

Table 6 examines the robustness of the qualitative result $\tilde{\beta}_0 > 0$ and significant, $\tilde{\beta}_6 < 0$ and significant, and $\tilde{\beta}_5 = 0$, across different fields in economics. As can be seen, the pattern is especially strong in columns (1)-mathematical economics and econometrics, (2)-microeconomics, (3)-macroeconomics, (4)-international economics, and (5)-finance. Those fields show higher probability of omission of relevant papers that include at least one female, when the citing team consists mostly of men, compared to mostly women. The effect of under-citation of women, compared to the target, is bigger in terms of magnitude in microeconomics, followed by macroeconomics and mathematical economics. Finance and international economics complete the list. By contrast, the effect of under-citation of female teams is not as strong in the fields of labor and education (column (6)), industrial organization (column (7)), and in column (8), which includes public, urban, and health economics.

Table 7 examines the robustness of main qualitative result when authors share the same institutions ("peer effects"). A "peer effect" between papers i and j is said to exist when at

Table 6. Omission and gender: Field of study

	Outcome variable: Omission							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mathe-matical	Micro	Macro	International Economics	Finance	Labour - Education	IO	Other fields
$A1f_i$	0.005	0.012	-0.010	0.013	0.013*	-0.001	-0.069**	0.005
	(0.008)	(0.008)	(0.011)	(0.009)	(0.007)	(0.008)	(0.034)	(0.009)
$A1f_j$	0.027***	0.038***	0.024**	0.019*	0.013*	-0.006	-0.021	0.020**
	(0.009)	(0.008)	(0.011)	(0.010)	(0.007)	(0.011)	(0.035)	(0.009)
$A1f_i \cdot A1f_j$	-0.091***	-0.069***	-0.048**	-0.064***	-0.055***	-0.043**	0.051	-0.034*
	(0.018)	(0.019)	(0.023)	(0.018)	(0.014)	(0.017)	(0.080)	(0.020)
N	18462	18354	9009	10187	28528	11562	1088	11031
R-sqr	0.131	0.105	0.115	0.131	0.098	0.132	0.184	0.121

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the gender of the citing paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. $A1f_x$ represents paper x with at least one female author. All the specifications include controls for paper j published in a top 5 journal; the relative cosine; paper i and paper j having the same primary field; difference between the publication year of paper i and the publication year of paper j; the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. The field section is defined based on the Journal of Economic Literature (JEL) codes. The category other fields includes public economics, agricultural economics, general economics, urban economics, law and economics, business administration, economic history, economics systems. The equations are estimated using a logit model. The table displayed the marginal probabilities. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

least one of the coauthors on each side have the same affiliation. For example, if team i and team j both include a coauthor from the Harvard economics department. This reduces the probability of related paper j being omitted from the references of paper i by 10pp, if paper i is all-male authored, and by 7pp if i is all-female authored. This suggests that peer effects are more beneficial for men than for women, although both genders benefit from being in a very close network.

For the sake of brevity, this section presents just a subset of the robustness analyzes that have been performed. In addition to what is presented here, the following robustness checks have been evaluated: Control for proxies of experience and seniority (maximum publications, maximum publication in top 5, date of first publication in the sample), existing recognition

Table 7. Omission and gender: Peer effects

	Outcome variable: Omission	
	(1)	(2)
f_j	-0.003	-0.006
	(0.005)	(0.005)
same affiliation	-0.099***	-0.101***
	(0.005)	(0.005)
f_j·(same affiliation)		0.031**
		(0.015)
N	94760	94760
R-sqr	0.165	0.165

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the effect of being in the same affiliation. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. f_j represents paper j written by only women. The control variables include the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

(existing citations and past publications), effect over time, alternative year difference between papers, consider up to 10 most similar papers, use of words by gender, Soft cosine and LDA with entropy similarity measures, extended sample to around 100 economic journals, the order of the name and gender, androgynous names versus non-androgynous names. The results are robust to all those alternatives specifications. Moreover, to investigate the difference by field, a comparison was made between Economics, Mathematics, and Sociology. Those results can be found in the online appendix.

2.6.4. Aggregate omission index

Because sometimes an article can be cited and other times it can be omitted, in aggregate, a gender bias may no longer appear. To analyze whether even when aggregating, the bias exists, the dependent variable in this subsection is the overall omission score for a given paper, denoted by O_{pt}. This omission score is calculated using three alternative measures. First, the (logarithm of) the raw count of the number of times a paper is relevant but has been omitted. Second, the intensity of omissions. Here, intensity is defined as the difference between the number of times a paper is relevant but has been omitted versus the times when the paper is relevant and has been cited. Third, the compensation for omissions. Here, compensation is defined as the difference between the number of times a paper is relevant but has been omitted minus the total number of times the paper has been cited. In all three cases, the true relevance of a paper is determined by the similarity index. The main specification is:

$$O_{pt} = \beta_1 \cdot gend_{pt} + \beta_2 \cdot \tilde{Z}_p^1 + \tilde{Z}_t^2 + \tilde{\epsilon}_{pt} \tag{2.6.3}$$

where $\tilde{Z}_p^1$ is a set of paper-level controls, such as the number of coauthors, the authors' affiliation, the field of the paper, the journal of publication, and a dummy that controls for NBER membership; and $\tilde{Z}_t^2$ captures publication-year fixed effects. Field- and journal- fixed effects are included. The standard errors are clustered by publication year and journal. The baseline regression is over a five-year horizon, but results are similar for different horizons.

The main coefficient of interest is β_1, and it measures the extent to which papers written by female teams are under-cited. Table 8 presents the results. The dependent or outcome variable is the overall omission score, as measured by the raw count of past omissions (column 1), the intensity of omissions (column 2), and the compensation for past omissions (column 3). In all specifications, the coefficient on fem (female author team, solo or group) is positive and statistically significant. This means that papers written by women are associated with a higher omission index, even after adjusting for the overall number of citations.

Measured by the raw count of omissions, papers written by women have a 0.06 higher omission index, compared to papers written by men. Measured by intensity, papers written by women have a 0.2-unit higher omission index, compared to papers written by men. This corresponds to an 18% increase in the median value of the omission index. Measured by compensation, papers written by women have a 0.6-unit higher omission index, compared to papers written by men. This corresponds to 40% of the average absolute value of the compensation score.

Note that being published in a top 5 journal, being affiliated with a top institution, and having one additional coauthor, are all associated with a significant reduction in the omission score.

2.7. Citations and publication quality

The existing literature has suggested that citations are a noisy signal of quality, for example in the case of patents. The analysis above also indicates that citations may not accurately reflect the quality of a published paper in economics, as they tend to systematically

Table 8. Aggregate omission

	Outcome variable: Omission		
	Raw Count	Intensity	Compensation
	(1)	(2)	(3)
female	0.054**	0.180*	0.638**
	(0.03)	(0.11)	(0.26)
mixed	0.026	0.104**	0.446***
	(0.02)	(0.05)	(0.17)
unknown	0.046	0.248*	1.682***
	(0.03)	(0.13)	(0.25)
N	11162	11162	11162
R-sqr	0.072	0.061	0.172

This table shows the relationship between the omission index and the gender a given paper. The dependent variable, omission index, indicates: the raw count of omission; the intensity of omission (sum omitted over sum cited and in prior literature) and the compensation (sum omitted minus sum citation). *female* represents papers written by only women. *at least one female* represents paper with at least one female author. The specifications include controls variables for the journal of publication, the institution, the number of authors, the year of publication, the field. Standard errors are clustered by journal of publication and years reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$* * * = p < 0.01$)

omit the contributions of female economists and groups of female economists. Therefore, this section, following Kelly et al.(2018) and Koffi and Panousi(2019), constructs an alternative index for measuring the quality of a publication in economics, using a textual and linguistic comparison across different papers. By constructing a measure of the quality which ignores the authors' willingness to refer to articles (therefore without bias in this sense), we can assess the differential relationship that exists between men and women by comparing our measure of quality with no inherent gender bias and citations where gender bias has been highlighted in the previous section.

2.7.1. Quality index

Specifically, the quality index, denoted by q, has two dimensions, which together capture the degree of innovativeness of a paper. First, more innovative papers are more distinct from prior related papers, in that they offer a novel idea or method to the pre-existing stock of knowledge. Second, more innovative papers are more likely to influence the framework or the methodology of future papers. In other words, the concept of innovation used here reflects the novelty as well as the influence of a publication. High-quality papers are both novel (distinct from prior papers) and influential (similar to future papers). Overall, the most important or significant papers introduce new concepts in the literature that make them different from their predecessor but very useful for future advances in economics.

The “novelty" of a paper is captured by a backward-similarity index, which is the sum of pairwise relative cosine similarities of paper p, published in t, with papers p' published in $t-T$:

$$BS^{0}_{-T}(p) = \Sigma_{p'}\tilde{\lambda}_{p,p'}$$

Papers with low backward similarity are dissimilar from the past literature, hence they are innovative with respect to the existing state of art.

The “influence" of a paper is captured by a forward-similarity index. The forward similarity is the sum of pairwise cosine similarities of paper p, published in t, with papers p' published in $t+T$:

$$FS^{T}_{0}(p) = \Sigma_{p'}\tilde{\lambda}_{p,p'}$$

Papers with high forward similarity have a higher impact on future publications. For example, they may open up a rich future line of literature, or they may propose an empirical methodology that many future papers will use.

Therefore, the quality index q will be a combination of the novelty and of the impact of a paper, as measured, respectively, by the backward and the forward similarity

$$q^{T}(p) = \frac{FS^{T}_{0}(p)}{BS^{0}_{-T}(p)}$$

In the benchmark specifications, $T = 5$, but the results are robust to alternative windows.

The q-index is a measure of the true underlying scientific quality or innovativeness or importance of a paper. If a paper has a high forward similarity (high numerator) and also a high backward similarity (high denominator), this could mean that the paper is a follower among other followers in a research area. Hence, it will have a low q-index, compared to a paper with a high forward similarity and a relatively low backward similarity. In that respect, it operates similar to the citations measure. Citations capture the degree to which a paper has influenced future papers. However, different from citations, the q-index also penalizes the lack of novelty (inverse of backward similarity). For example, a paper may have high forward similarity because it uses general language in the description. For those kinds of papers, the normalization by the backward similarity will counteract the high forward similarity and, yielding a relatively low level of quality, ceteris paribus. In the baseline, the horizon chosen to compute the q-index is 5 years.[18]

2.7.2. Quality and citations

To analyse the link between the quality index and the number of citations received, the following equation is estimated:

$$C_{pt} = \gamma_1 \cdot Q_{pt} + \gamma_2 \cdot \hat{Z}_p^1 + \hat{Z}_t^2 + \hat{\epsilon}_{pt} \tag{2.7.1}$$

where

C_{pt} is the logarithm of the number of citations of paper p published in year t; Q_{pt} is the quality index of paper p published in year t; $\hat{Z}_p^1$ is a set of paper-level controls, such as the number of coauthors, the coauthors' affiliation,the field of the paper,the journal of publication, and NBER membership; $\hat{Z}_t^2$ captures publication-year fixed effects. Field- and journal- fixed effects are included. The standard errors are clustered by publication year and journal.

Table 9 presents the results of this regression for horizons of 1, 3 and 5 years. Regardless of specification, there is a positive and statistically significant correlation between the quality index and the number of citations received by an article. For the 5-year window, a one-standard-deviation increase in the quality index, q, is associated with an increase in citations of 0.4, on average, which corresponds to an increase of the median number of citations of about 30%.

For the two other variables, the number of authors and the fact of having one NBER member among the author, the relation with the number of citations is also positive and significant. Similarly, to Hamermesh (2019), one additional author raises the number of citation but does not raise it in a one to one relation. In fact, one additional author increases the number of citations (log of one plus) by 0.038. Interestingly, conditional on having the same level of quality, having a NBER member author increases the citation count by 0.227.

In addition to predicting the level of the citations over the same horizon, the quality index is also related to future citations. There is a positive and significant relationship between the quality index on the 0-5 year-window as well as the forward citations after the 5th year post publication. Therefore, the text-based quality index is positively correlated to the number of citations, which in turn has been a widely used measure of a paper's importance.

[18]Figure 6 (in the appendix), panel (a), plots the pooled distribution of the q-index. This quality distribution is less skewed, compared to the citation distribution (see tables 5 and 1).

Table 9. Quality index and citations: 5-year horizon

log(1 + cite)	(1)	(2)	(3)	(4)	(5)	(6)
log(1 + q)	0.517***	0.513***	0.697***	0.609***	0.585***	0.587***
	(0.068)	(0.061)	(0.047)	(0.038)	(0.037)	(0.036)
authors		0.159***	0.135***	0.061***	0.038***	0.038***
		(0.012)	(0.010)	(0.008)	(0.008)	(0.008)
nber		0.485***	0.459***	0.383***	0.243***	0.227***
		(0.026)	(0.027)	(0.018)	(0.020)	(0.020)
N	17173	17173	17173	17173	17173	17173
R-sqr	0.02	0.10	0.10	0.28	0.32	0.33
Publication Year FE			Y	Y	Y	Y
Journal FE				Y	Y	Y
Institution FE					Y	Y
Field FE						Y

This table shows the relationship between the quality index, q, and the number of forward citations, *cite*. The controls include dummies for journal, field, affiliation, year of publication, number of authors, and NBER membership. See equation 2.7.1 in text for the construction of q. Standard errors are clustered by journal of publication and year and are reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

Moreover, the quality measure has the advantage of identifying important prior literature without relying on the willingness of the authors to cite or not prior relevant papers.

2.7.3. Quality, citations and gender

This section examines the relationship between the quality index, the number of citations, and the gender of the authors. To that end, equation 2.7.1 is augmented with the addition of a set of dummies that control for the gender structure of the authors.[19]

$$C_{pt} = a_1 \cdot Q_{pt} + a_2 \cdot gend_{pt} + a_3 \cdot gend_{pt} \cdot Q_{pt} + a_4 \cdot Z_p + \theta_t + \epsilon_{pt} \tag{2.7.2}$$

The variable *gend* takes the value 1 if there is at least one female author, and 0 if the paper is written by all-male teams; or a dummy variable that take 1 if the paper is written by only

[19]A selection bias may be present, as women and men may not have publish in top 5 journals with the same probability.

women and 0 if the paper is written all-male teams (solo and co-authored).[20] All the other variables are the same as in equation 2.7.1.

Fig. 11. Quality, Citation and Gender by journal

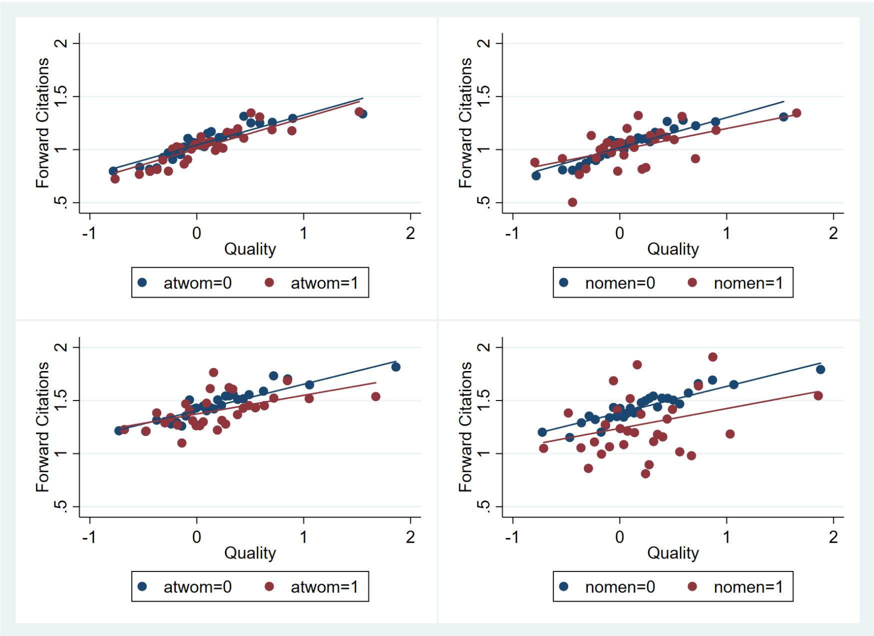

The figure plots the link between the number of forward citations and the q-index for papers written by males and papers written by at least one female (left) and for males and females (right). The variable *atwom* indicates a dummy variable that takes one if the paper is written by at least one female and 0 otherwise. The variable *nomen* indicates a dummy variable that takes one if the paper is written by only females and 0 otherwise. The upper graphs are for the non top 5 journals and the lower graphs for the top 5 journals. The binned scatter plot controls for journals, field, institutions, year of publications, number of authors, NBER member. The corresponding estimates are in table 7.

Figure 11 shows the results of the estimation. The first row column is non-top-5 and the second row is top 5. Left row is at least one female, right row is only female. Panel at the left on the first row refers to papers that are published in non-top-5 journals, written by all-male teams versus by at-least-one-female team. The male papers have a quality index that is pretty close to the average of the q-distribution, whereas the female papers have a quality level similar to the 70th percentile of the q-distribution. In other words, to receive the same number of citations, papers with even just one female coauthor have to be substantially better, in terms of innovation and influence, compared to male-authored papers. Panel at the left on the second row shows that the same increase in quality, from the p50 to the p70 of

[20]The quality index is here standardized. For each value, the mean is subtracted and the result is divided by the standard deviation. This is helpful for the interpretation of the regression coefficient in estimation with a cross variable and a continuous variable. The results are not dependant of the standardisation.

the q-distribution, is required for top 5 publications, when the paper has at least one female coauthor, compared to an only-male team of coauthors.[21] Panel at the right on the second row shows that, for the same number of citations, the jump in quality needed for all-female teams of authors is even higher, specifically from the p50 to the p90 of the q-distribution. In other words, top 5 journal publications receive the same number of citations when they are written by men and have median quality, or when they are written by all-female teams with quality at the top 10% of the distribution. The comparison across these three panels shows that, for the same number of citations, papers written by teams that include women require an increase in quality from the p50 to the p70 when the journal level increases from non-top-5 to top 5; and an additional jump in quality to the p90, if all authors of top 5 publication are to be women. Or, for the same number citations, as long as there is one male in the team, the increase in quality from the median is the same in all tiers of journals, when a woman is added to the coauthors. However, when all coauthors are women, an additional increase in quality in required at the top 5 in order to maintain the same number of citations.

Figure 12 is a key result showing the effect of the bias generated by the omissions. In a counterfactual analysis, the total number of omissions is added to the number of citations. This is interpreted as the number of citations an article would have received if all the papers with which it shares the most similarities had cited it. The gap corresponding to the 20 percentile in terms of innovativeness index disappears completely with this compensation. In other words, if it were not because of the omission bias, the standards to be cited would be the same for women as for men.

2.7.4. Robustness

In this section, equation 2.7.2 is estimated separately for top 5 journals versus non-top-5 journals, and for top-tier institutions versus non-top-tier institutions. The results are presented in the appendix.

At an average quality level, a paper written by men will received around 3.8% more citations than a paper written by a mixed gender team if the papers are published in a non top 5 journals. The difference between mixed gender team and men team is insignificant for papers published in a top 5 journal. By contrast, the difference between papers written by men and papers written by women (in terms of citations given a certain level of quality) is insignificant when the paper is published in a non top 5 journal and significant when

[21]Note, however, that the test here does not have enough power.

Fig. 12. Counterfactual: citation compensating with omission cases (Top 5)

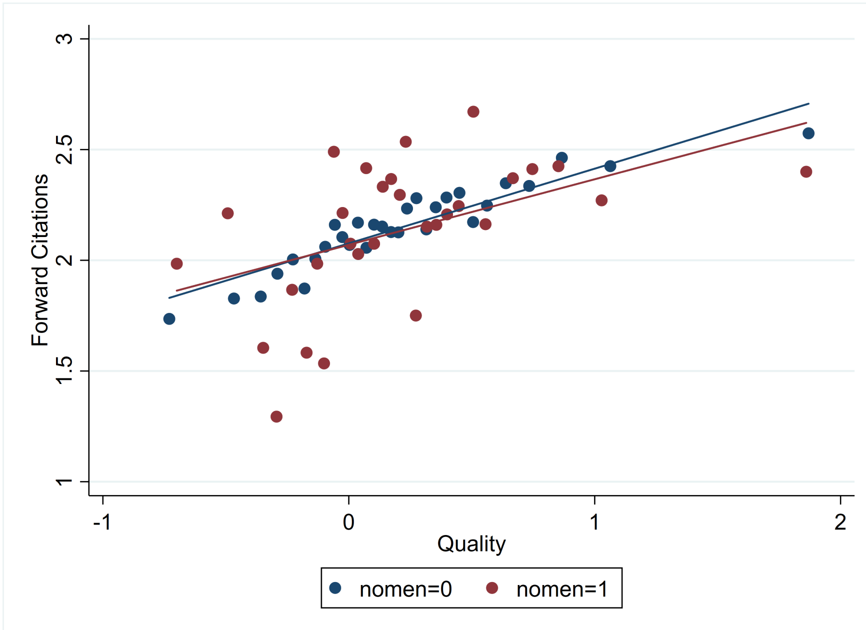

The figure plots the link between the number of forward citations compensating by the number of omissions and the q-index for papers written by males and papers written by females for top 5 journals. The variable *atwom* indicates a dummy variable that takes one if the paper is written by at least one female and 0 otherwise. The variable *nomen* indicates a dummy variable that takes one if the paper is written by only females and 0 otherwise. The binned scatter plot controls for journals, field, institutions, year of publications, number of authors, NBER member. The corresponding estimates are in table 7.

the paper is published in a top 5 journal.[22] Indeed, papers written by women in a top 5 journal received about 15% less citations for a paper at the average of the quality distribution compared to a paper written by men in a top 5 journal. Women tend to benefit less in terms of citations from publishing in top 5 journal compared to their males' colleagues.

Also, papers written by women will tend to receive less citations compared to papers written by men. However, the effect is only significant for the mid-tier institutions. Indeed, papers having at least one female author received 5% less citation compared to men for a mean quality paper.

[22]The line of best fit on Figure 11 shows that for papers written by women versus papers written by men in a non top 5 journal, the lines are crossing each other. This may hide a cross over interaction. For high quality papers, papers written by women may be less cited whereas this is the opposite for low quality papers. This could explain why the coefficient is significant when we consider mixed gender team and insignificant when we focus only on women.

2.8. Gender omission bias and future productivity

This section examines the relationship between the history of past omissions and the future productivity of a given author. The analysis uses three measures of future productivity. First, the probability to be published in a top 5 journal within the next 3 years. Second, the total number of forward citations of all papers by the same author over the next 5 years. Third, the quality of future publications by the author over the next 5 years.

The regression specification for the case where the measure of future productivity is the probability of publication in the top 5 is:

$$Top5_{p,t,a} = \theta_1 \cdot H\_Omission_{p,t,a} + \theta_2 \cdot gender_a + \theta_3 \cdot \Gamma_a^1 + \theta_4 \cdot \Gamma_p^2 + \Gamma_t^2 + \varepsilon_{p,t,a} \tag{2.8.1}$$

Here, $Top5_{p,t,a}$ is a binary variable that indicates if paper p by coauthor a is published in a top 5 journal in year t. Next, $H\_Omission_{p,t,a}$ is the omission history of author a prior to paper p. This captures all the previous times any paper written by author a belonged in the relevant literature, according to the similarity index, but was omitted from the references.[23] The control variables include author-specific variables, such as gender, and characteristics of past publications, such as quality, citations, and number of top 5 publications. Paper characteristics such as field, year of publication, and time lag between two publications are also included.

Table 10 presents the results. As can be seen by the highlighted coefficients in the first row of the table, an increase in the history of omissions reduces the probability of being published in the top 5 within the next 3 years. The coefficient is tightly estimated in the interval $[-0.01, -0.04]$, depending on the various controls included in columns (1) through (11), and it is always statistically significant at the 1% level. Overall, moving from a relatively low level of past omissions to a relatively high level is associated with a $10 - 15\%$ decline in the odds of getting published in the top 5 within the next 3 years. Furthermore, on average, men in mid-tier institution with relatively high history of omissions, have 2 times higher odds of being published in the top 5 within the next 3 years, compared to similar women. Essentially, the probability of publishing in a top 5 journal is reduced by 15% for the men versus by approximately 30% for women. The difference between men and women in top- and low-tier institution is also significant at the 10% level.

[23]Alternative measures for the history of omissions have also been used. For example, the average number of past omissions, the omission index for the paper published by author a just before paper p, and the intensity of omission of past publications. The qualitative results remain unchanged across the different measures. In addition, results remain qualitatively similar if the dependant variable is the journal impact factor.

Table 10. History of omissions and future publications

	Outcome Variable: Probability of publication in top 5 within next 3 years										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
H_Omission	**-0.025***	**-0.036***	**-0.037***	**-0.039***	**-0.034***	**-0.034***	**-0.033***	**-0.011***	**-0.011***	**-0.016***	**-0.010***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.006)	(0.004)
H_Citation		0.020***	0.019***	0.021***	0.010***	0.003	0.009**	0.012***	0.012***	-0.016***	0.012***
		(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)	(0.004)
gender			-0.028**	-0.026*	0.002	0.001	0.009	-0.007	-0.006	-0.022	0.009
			(0.014)	(0.014)	(0.010)	(0.012)	(0.012)	(0.013)	(0.013)	(0.018)	(0.018)
H_Quality				0.004***	0.004***	0.003***	0.001*	0.000	0.000	0.002*	0.000
				(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
H_top5					0.024***	0.020***	0.021***	0.035***	0.034***		0.035***
					(0.004)	(0.004)	(0.004)	(0.004)	(0.004)		(0.004)
nber								0.020***	0.020***	0.002	0.020***
								(0.004)	(0.004)	(0.005)	(0.004)
star									0.038**		
									(0.017)		
Gend*H_Omission											-0.013
											(0.010)
Current paper gender								Y	Y	Y	Y
Past publications										Y	
Time lag						Y	Y	Y	Y	Y	Y
Year FE								Y	Y	Y	Y
Affiliation FE								Y	Y	Y	Y
Field FE					Y	Y	Y	Y	Y	Y	Y
N	32227	32227	32137	31290	31290	22957	22957	22521	22521	22521	22521
R-sqr	0.004	0.006	0.007	0.009	0.177	0.176	0.185	0.232	0.232	0.075	0.232

This table presents the relationship between the history of past omissions and the probability to publish in a top 5 journal within the next 3 years. The history of omissions, $H_$Omissions, is measured as the cumulative number of past omissions. The control variables include author-specific variables, such as gender, characteristics of past publications (quality, citations, number of top 5), paper characteristics such as field, year of publication, time length (lag) between two publications. Linear probability model. Standard errors clustered at author level and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

Table11 presents the results for different future productivity measures. In the first block of the table, the dependent variable is productivity as measured by forward citations over a 3-year horizon. As shown in the first row, an increase in the history of omissions (in the total number of past omissions) reduces the number of future citations. In particular, an increase from low to high omissions is associated with an average decline of 12% in the average number of citations over the next 3 years. In the second block of the table, the dependent variable is productivity as measured by quality of future papers over a 3-year horizon. As shown in the first row, an increase in the history of omissions has a negative but mostly insignificant effect on the quality of the future publications. Controlling for cross-effects between gender and the history of omissions does not substantially change the results across men and women. This could be due to a selection effect, as women in this sample have at least two top publications. Hence, it is more likely that senior ranking women will

Table 11. History of omissions and future publications

	Outcome Variable:					
	Forward citations			Quality		
	(1)	(2)	(3)	(4)	(5)	(6)
H_Omissions	**-0.081*****	**-0.121*****	**-0.077*****	**-0.002**	**-0.004****	**-0.002**
	(0.009)	(0.011)	(0.010)	(0.002)	(0.002)	(0.002)
gender	0.049	0.048	0.010	-0.002	-0.008	-0.012
	(0.033)	(0.042)	(0.044)	(0.006)	(0.008)	(0.008)
H_Citations	0.156***	0.214***	0.171***	-0.004**	-0.001	-0.003
	(0.009)	(0.010)	(0.009)	(0.002)	(0.002)	(0.002)
H_Quality	-0.001	0.000	-0.001	0.000	0.000	0.000
	(0.002)	(0.002)	(0.002)	(0.000)	(0.000)	(0.000)
H_top5	-0.005	0.008*	0.009**	0.002***	0.003***	0.003***
	(0.003)	(0.005)	(0.004)	(0.001)	(0.001)	(0.001)
Gend*H_Omission		0.022	0.024		0.007	0.007*
		(0.030)	(0.029)		(0.004)	(0.004)
Current paper gender	Y		Y	Y		Y
Past publications	Y		Y	Y		Y
Time lag	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Affiliation FE	Y		Y	Y		Y
Field FE	Y	Y	Y	Y	Y	Y
N	22521	22957	22521	22502	22938	22502
R-sqr	0.233	0.184	0.234	0.501	0.500	0.501

This table presents the relationship between the history of past omissions and the number of future citations or the quality of future publications. The future horizon is 3 years. The history of omissions, H_Omissions, is measured as the cumulative number of past omissions. The control variables include author-specific variables, such as gender, characteristics of past publications (quality, citations, number of top 5), paper characteristics such as field, year of publication, time length (lag) between two publications. Linear probability model. Standard errors clustered at author level and reported in parentheses.
$(* = p < 0.10, ** = p < 0.05, *** = p < 0.01)$

be over-represented in the sample. Additionally, omission seems to increase the quality of future female publication, while it reduces the quality of male future publications.

In conclusion, a history of omitted articles is associated with lower future productivity, as measured by future citations, future top 5 publications, and, for the men, by future publication quality.

2.9. Role of editor's gender

Editors play a central role in the article review process. The editors of economic journals, especially at the top 5, are predominantly male. In terms of editorial structure, AER and EMA designate one person as the main editor, while JPE, QJE and REStud traditionally relied on editorial boards. More recently, however, the JPE designated a lead editor on the front page of the journal. Over the sample period 1991-2019, EMA has had no female editors, and QJE has had no females on its editorial board. REStud has the largest editorial board, with editors located all over the world. Thus, across top 5 journals, the editorial structure is most similar across AER and EMA. Next, in 2011, the AER appointed a female editor, and this tradition has continued until 2019. By contrast EMA has not had gender changes in its editorial structure. This allows for an estimation of the impact of a (possibly exogenous) change in editorial policy on odds of omitting relevant female papers, compared to male papers.

To estimate the effects of the change in the gender of the AER editor, we track the probability for an AER publication to omit relevant prior papers written by females before and after the change in the editor's gender. We then compare these changes with the corresponding ones for papers published in EMA. If the change in the gender of the editors plays a role for publication standards, the coefficient of interest should be negative and statistically significant. In other words, a change from a male to a female editor reduces the odds of female authors being omission from the relevant references:

$$\Delta^2 = [Pr(Omit)_{aer}^{>=2011} - Pr(Omit)_{aer}^{<2011}] - [Pr(Omit)_{eco}^{>=2011} - Pr(Omit)_{eco}^{<2011}] \quad (2.9.1)$$

Because the change in the editor's gender could possibly affect the propensity to cite female papers, we examine changes in the probability of omission specifically for papers written by women (at least one female). Therefore, the associated coefficient will capture the extent to which changes in the likelihood of omission for female-authored papers by other papers published in the AER (relative to papers published in EMA) differs from the corresponding changes experienced by similar male authors:

$$\Delta^3 = \Delta_F^2 - \Delta_M^2 \quad (2.9.2)$$

Table 12. Effect of editors' gender

	Outcome variable: Omission					
	2011		2012		2013	
	(1)	(2)	(1)	(2)	(1)	(2)
$A1f_j$	0.088	0.072	0.087	0.018	0.086	-0.005
	(0.077)	(0.156)	(0.077)	(0.147)	(0.076)	(0.139)
After change	-0.105	-0.072	-0.091	-0.071	-0.291	-0.292
	(0.242)	(0.244)	(0.236)	(0.239)	(0.225)	(0.228)
AER	-0.014	-0.056	0.037	-0.013	0.081	0.031
	(0.084)	(0.088)	(0.081)	(0.085)	(0.078)	(0.082)
(After change)$\cdot$ AER	0.474***	0.530***	0.410***	0.485***	0.346***	0.428***
	(0.110)	(0.120)	(0.112)	(0.122)	(0.115)	(0.126)
$(A1f_j)$$\cdot$(after change)		-0.081		0.003		0.050
		(0.215)		(0.215)		(0.222)
$(A1f_j)$$\cdot$(AER)		0.298		0.344*		0.341*
		(0.204)		(0.192)		(0.182)
$(A1f_j)$$\cdot$(AER)$\cdot$(After change)		-0.338		-0.442		-0.468*
		(0.279)		(0.278)		(0.282)
N	15598	15598	15598	15598	15598	15598
R-sqr	0.180	0.180	0.180	0.180	0.179	0.180

This table presents the difference in difference and the triple difference estimation to assess the effect of a change in the gender of the editor. The model aims to look at if the probability to omit a paper is influenced by the change in editorial policy. American Economic Review is used as the treated group. Econometrica is the control group. The change occurs in 2011. To take into account delays of the effects, the table presents the effect for the year 2012 and 2013. The variable of interest are *(After change)*AER* for the difference in difference and *(at least one female j)*(AER)*(After change)* for the triple difference. The table shows the odds ratios. An odd greater than 1 means that the policy increases marginally the likelihood of omission. An odd ratio lower than 1 means that the policy reduces the likelihood of omission. The control variables include the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$* * * = p < 0.01$)

where Δ^2_g, $g \in \{F,M\}$ is Δ^2 defined for gender F or female, and gender M or male. A negative coefficient implies that the change in gender reduced the odds of omission for female papers published in AER (versus EMA), compared to male papers. The analysis controls for a set of characteristic of the citing, cited, and omitted papers.

The results are presented in Table 12. Δ^2 is positive and significant meaning that the change in editorial policy tends to increase the overall probability to be omitted. The marginal probability increase can go up to 5% in terms of the omission propensity. Interestingly,

Δ^3 is negative (in all specification) and significant (in some). It means that the policy change tends to be beneficial for mixed gender team, somehow at the expense of only male teams (solo or coauthors). Concretely, the odd of being omitted is 30% to 50% less for mixed author teams than for men author teams after the policy change. Nevertheless, there seems to be a delay for this policy change to have a substantial impact.

2.10. Conclusion

Women are still under-represented in math-intensive fields. But very few studies have tried to analyse whether the potential problem lay in the lack of recognition of their work. This paper has, therefore, addressed the issue using data on Economics. It shows that women have higher probabilities of being omitted from references. This problem is persistent, even for women publishing in top-journal in the same way as men. However, the most vulnerable population appears to be women of mid-tier institutions. Indeed, for women at a top-tier institution, the system seems to internalize the information on their potentialities and updates its prior. Consequently, the gender bias in omission is smaller. For women at low-tier institution, the affiliation effect is much more important than the gender effect. That being said, the situation can, however, be greatly improved with an increase in the number of women in faculties and at decision-making positions. Finally, the conclusion of this study is not only restricted to the economic field. It aims at a more general horizon by explaining how the presence of discriminating factors can influence the perception of individual works even when they are more deserving than others. In addition, recognizing how issues related to inequality can affect socio-economic factors, the current paper advocates for greater diversity and inclusion of minorities for better productivity.

Chapter 3

Time-varying investment risk in financially integrated economies

3.1. Introduction

Entrepreneurs play a crucial role in the innovation landscape and the economic growth of a country.[1] However, entrepreneurship remains very risky. In fact, those innovative agents are exposed to a largely undiversifiable investment risk that could take various forms: difficulty of access to credit, negative profits, technology shock, ability shock, depreciation shock, etc. Moreover, this market failure implies also a precautionary saving motive and creates a wedge between the marginal product of capital and the risk-free rate (Angeletos, 2007). Therefore, in crisis event, it affects not only their ability to invest, but also their willingness to invest. Recent studies have suggested that in a closed economy, a time-varying investment risk could create very deep and detrimental recession because agents continue to bear a substantial fraction of the aggregate risk (DiTella, 2017). Thus, in such a situation, the presence of a relatively less risky asset is greatly valued. But so far, little or no attention has been given to the possibilities of sharing time-varying innovation risk through financial integration.

This paper investigates the insurance against time-varying idiosyncratic investment risks for financially integrated economies.[2] Specifically, it studies the costs and benefits of financial integration for countries with different degrees of financial development. Traditional business-cycle models with first-moment or aggregate TFP shocks cannot match the observed direction of global capital flows upon integration. By contrast, models with idiosyncratic investment risk, such as the one in Angeletos and Panousi (2012), can explain global imbalances but

[1]Joseph Schumpeter (1911/1934) was one of the pioneers defending this idea. See among others King and Levine (1993), Wong et al.(2005), Audretsch et al. (2006). According to an OECD report of 2017, SMEs (small and medium-sized enterprises) account for 60 to 70 per cent of jobs in most OECD countries.

[2]An uncertainty shock is also designed in this paper by an uninsurable time-varying idiosyncratic investment or capital income risk.

they cannot give the implications in terms of aggregate fluctuations.[3] Then, the idiosyncratic investment risk is introduced as a second-moment shock which is coupled to a first moment shock (a technology shock). In line with the empirical literature, the correlation between both shocks captures the countercyclical behavior of volatility.[4] For instance, recessions are periods when the dispersion of firms' sales increases. These two types of shocks capture how individual agents react to shocks that affect not only their ability to invest but also their willingness to invest. Therefore, this paper sheds light on the role of a covariance structure of shocks in an international finance framework. In addition to empirical relevance, this paper contributes to the ongoing debate about the costs and benefits of globalization and financial integration. In particular, it highlights a previously unexamined source of benefits, namely the insurance against investment risk shocks provided by a financial union. In this line, one of the key contributions of this paper is to rationalize two main empirical findings: the decrease in aggregate stock return volatility following financial liberalization (Umultu, Akdeniz and Altay-Salih, 2009) and the perverse effect of capital control when introducing firm level data (Alfaro, Chari and Kanczuk,2017). This novel theoretical possibility emerges due to asset reallocation, increases in production and increases in the willingness to undertake risks.

Preview of the model: The model is conceptually related to Angeletos and Panousi (2011) and technically extends the methods and Brunnermeir and Sannikov (2015) and Di Tella (2017). This is a continuous time, two-country model with heterogeneous agents and an aggregate shock. In each country, there is a continuum of firm-household (investor) whose production is subject to an idiosyncratic shock. At each point in time, the idiosyncratic shock is drawn from a distribution whose mean is 0 and its variance is time-varying. The variance of this distribution could enter in a high regime (above the long-run value) or a low regime (below the long run value) depending on the aggregate shock. Investors have the choice between investing in the capital (the asset subject to the idiosyncratic shock) or investing in a one-period bond. Hence, the economy features a market incompleteness as the only asset available to diversify away the risk is a one period bond. When an investor chooses the capital asset, she runs a private firm. There is a free international market for capital and a free international market for consumption good. The financial autarchy state is characterized by a closed capital account for bonds and agents have only access to the domestic bonds. In financial integration, bonds can be traded internationally. The countries differ in their level of financial development which is assessed by the level of the long run value of uninsurable idiosyncratic risk. In fact, uninsurable idiosyncratic investment

[3]See Gourinchas and Jeanne (2013) for the fact that capitals fail to flow from rich to poor.
[4]Bloom(2014), Bachmann and Moscarini(2012), Kehrig(2015), Decker et al. (2016).

risk captures market imperfections as it impedes investors to diversify away the risk they encounter. When the investors could diversify a higher fraction of risk (financially developed country), the long run value of idiosyncratic risk will be lower. At the opposite, when the degree of diversification is lower, this long run value will be higher (financially underdeveloped country).

Preview of the results: The main results are as follows. First, under financial autarchy, each country bears its own specific risk. The relatively risky country experiences a strong demand for precautionary saving which is not met because of the limited supply of the safe asset. This depresses the interest rate. In recessions where the level of idiosyncratic risk is higher, agents are less willing to invest and require a higher premium to invest in capital because of the greater risk they will encounter. This increases the wedge between the marginal product of capital and the interest rate, leading to a drop in asset prices. Because of the lack of diversification, an uncertainty shock in a given country increases unambiguously the volatility of asset prices and thus, the endogenous volatility in the system. Then balance sheets are hit tremendously.

However, with the integration, the availability of the foreign asset is greatly valued in the presence of a time-varying idiosyncratic investment risk. Because investors are reluctant to invest, having another asset enables a portfolio reallocation between capital and bond. The relatively risky country runs a current account surplus, which at the same time reduce its exposure to the risk on capital. Then, this mitigates the increase in endogenous volatility. This mechanism that aimed at reducing the negative impact of the shock is absent from model without a time-varying investment risk. Most surprisingly though, we could observe a decrease in aggregate volatility with an increase in the domestic idiosyncratic risk. In general, the volatility of asset prices is lower (two to five-fold lower) in the integration state compared to the autarchy state. Hence, the net worth of agents increases in the relatively risky country in the global market as opposed to financial autarchy. The country then experiences a wealth effect, that will increase its risk-taking and have a better growth path compared to the autarchy situation. The key mechanism here is the sensitivity of investment to a variation in the investment risk (specific to the presence of the time-varying investment risk) and the possibility of diversification offered by the openness. Overall, the Portofolio reallocation is more than two times higher in an economy with time-varying investment risk compared to an economy without.

Considering the cycle frequencies, aggregate variables are better at each point in time in the financial integration state. But the side effect is the higher exposure of the riskiest country to foreign shock in the safest country (two times more exposed to foreign shock).

Consequently, with a higher net worth and a lower exposure to domestic shock compared to autarchy, financial integration even at the cost of exposure to foreign shock, yields sizeable welfare gains. The pure gains of the diversification of idiosyncratic risk for the riskiest country are on average 22% versus 8% for the safest country. Those values are substantially greater when compared to the existing values of the literature.[5] To summarize, three main effects are guiding the results: a portfolio reallocation effect, a price effect and a wealth effect. Because of those three effects, a less developed country can implement a stabilization policy by joining a financial union and benefit from the risk-sharing spillover with the developed country. However, as the long run value of idiosyncratic risk increases, this induces an endogenous borrowing constraint that limits the participation to the international financial market. And if the country with the higher long-run value of idiosyncratic risk is the one with a higher relative share of world GDP, this will decrease the global world capital flows and lower the price of capital. Finally, we find evidence that our main theoretical channel is supported by the data.

3.2. Literature review

This paper is primary related to two strands of the literature: the literature that emphasizes the role of idiosyncratic investment risk in integrated economies and the vast literature on the role of financial integration.

First, the idea that idiosyncratic investment risk can create frictions and impacts the decision of the agents started with Angeletos(2007). He showed that at the opposite of Bewley-type model, the aggregate economy can display a lower capital and a lower interest rate (compared to the complete market) when there is idiosyncratic investment risk. In this regard, the closest papers to ours is Angeletos and Panousi (2011). Even if we share the common set-up of introducing capital income risk in a financially open economy with Angeletos and Panousi (2011), the fact that we additionally add a time-varying dimension brings important changes. The first one is on a quantitative order. Looking at the transitional dynamics, the introduction of uncertainty shock reduces by two compared to Angeletos and Panousi (2011) the share of capital hold by the relatively risky country in the integration states. But, the investment to output ratio in integration increases by 1% in Angeletos and Panousi (2011) and 27% in our set-up. The wealth effect considered at business cycle frequencies is greater. Moreover, there is a noticeable qualitative difference between this

[5] A similar model as Angeletos and Panousi (2011) or Coeurdacier et al.(2019) with no time-varying risk generates a welfare gain more than 50% lower.

paper and Angeletos and Panousi (2011). In fact, in the latter, we cannot give the implication of idiosyncratic risk for business cycle fluctuations and talk about stabilization policy. Angeletos and Panousi (2011) framework have been extended to study the positive relation between capital outflows and growth or savings and asset price boom (see for e.g. Benhima(2013), Sandri (2014), Feng (2017)). However, all those studies have focused on a long run perspective and not business cycles implications.

Second, this paper fits well within the large literature on the effect of financial globalization.[6] Similar to this paper, Coeurdacier et al. (2019) combine in the same framework risk sharing and capital allocation in a two-country neoclassical growth model. [7] But in their set-up, the gains from risk sharing are offset by the losses from capital allocation. The gains from integration are therefore elusive when one considers both capital allocation and risk sharing, even for extreme values of risk aversion (risk aversion of 40). By contrast, this paper shows that taking into account uncertainty shocks reverses this conclusion. In fact, in presence of an uncertainty shock, agents value differently the diversification opportunities. Therefore, the gains from integration are substantial, even with a trade-off between risk sharing and capital allocation. In Buera and Shin (2017), financial integration comes with a reform that reduces exogenously the effect of country idiosyncratic risk. At the opposite, we propose a mechanism through the reallocation, the price and the wealth channel that reduces endogenously the effect of idiosyncratic risk with the openness and this. Mendoza et al. (2009) showed in a model with idiosyncratic labor income shock, that welfare findings are the consequences of changes in the interest rate after financial integration. First, they ignored capital income risk induced by investment decision and capital accumulation. Second, they abstract from aggregate shock and the possible link between aggregate shock and idiosyncratic risk. Therefore, the changes in the interest rate is not coupled with the efficiency gains of capital allocation as in our case. This gives rise to very different quantitative implications: the less developed country could benefit as well from financial integration. In addition, Brunnermeier and Sannikov(2015) study how a pecuniary externalities can arise in two integrated economies and the necessity of capital controls in such a situation. This paper departs from theirs by using a micro-founded model with heterogeneous agents, where the role of a financial union is particularly emphasized.

[6]Among others: Agénor (2003), Broner and Ventura (2016), Devereux and Yu (2014), Gourinchas and Jeanne (2006), Hoxhaa et al.(2013), Kray and Ventura (2000), Lewis and Liu (2015)

[7]Devereux and Sutherland(2009) also incorporate risk sharing in DSGE two-country model. But the goal of their paper was to analyse the effect of alternative portfolio choice (equity, bonds) on the risk sharing opportunities.

Baxter and Crucini (1995) and Colacito and Croce (2010) have also showed that the effect of financial openness is very small in the presence of transitory shocks. Only big shocks (permanent shock) are relevant to find higher welfare gains from financial integration. This paper shows clearly that it really depends on the type of shocks. If this conclusion is right for a first moment shock, it changes when it is a second moment shock. The effect of an uncertainty shock on the economy can be substantial even if it is a short-lived shock.

Looking at how financial openness affects the macroeconomic volatility, Aghion et al. (2004) concluded that economies at an intermediate level of financial development are more unstable. In the same vein, Buch et al. (2005) have found that the volatility of real economic variables upon financial integration depends on the underlying policy shocks. We depart from those papers using first heterogeneous agents, but also in the results. In fact, it is shown that the price effect resulting in a financial union can dampen the traditional impact of financial globalization on volatility even if the openness increases the exposure to foreign shock. Bai and Zhang (2011) point out the default risk as a potential explanation of limiting risk international sharing. Bengui et al. (2013) show that international portfolio rigidities can reduce the possibility of risk sharing. Complementary to their studies, we propose here a financial union to take advantage of financial integration in an environment where risk averse investors face an uncertainty shock.

Moreover, this paper stands in the new literature that tried to understand foreign asset dynamics with volatility shock. Fogli and Perri (2015) showed how time-varying volatility shock can explain the medium and long run dynamic of net foreign asset (NFA) positions of OECD countries. Indeed, when there is a bad volatility shock, agents save more, and part of these savings are in international bonds. But they fail to have the drop in output growth following a bad volatility shock. Gourio et al. (2015) explained how stock market return volatilities forecast capital flows in emerging countries using the “expropriation risk” (a set of political risk and legal disputes). With a higher country-specific volatility, foreigners pull their capital out because of the increase in expropriation risk and domestic residents of those countries with high volatility sell more assets than they buy. The present model gives a new channel to assess the dynamic of NFA (channel that is true whether we consider a developed country or an emerging market). This paper uses idiosyncratic volatility shock instead of aggregate volatility shock. Indeed, an increase in NFA following a bad uncertainty shock is due to a portfolio reallocation that investor made, choosing to invest in international bond instead of the risky capital.

In addition, the interest in second-moment shocks has been raised by the seminal paper by Bloom (2009). Indeed, volatility shocks (second-moment shocks) appear to be key drivers of business cycles with the specificity to create particularly harmful recession periods. Moreover, following Bachmann and Bayer (2013), Bachmann and Moscarini (2011), Bloom (2009, 2014), Christiano et al. (2014), Di Tella(2017) Fajgelbaum et al. (2016), Fernandez-Villaverde et al. (2011), we introduce time varying volatility. Similar to Di Tella(2017), we specifically consider time varying idiosyncratic risk (not for the aggregate risk). Whereas Di Tella focused on risk sharing between household and experts in the presence of an uncertainty shock, we focus instead on risk sharing between two countries.

Furthermore, this paper sheds new light on the age-old question about the welfare effects of eliminating business cycles. Lucas(1987) estimated that the cost of eliminating business cycles is negligible, in a world with homogeneous agents and without market frictions. Krussel and Smith (1999,2009) revisited the topic in a model where agents face uninsurable idiosyncratic wage-income risk. This paper is revisiting the question in an environment where investors face undiversifiable return risk. Here, the presence of countercyclical risk increases economic fluctuations in autarchy. This force tends to argue towards bigger welfare benefits from eliminating cycles, compared to the standard neoclassical model. In addition, the effects are quantitatively larger than in Krussel and Smith (1999,2009).

Finally, the paper is rationalizing some empirical findings. Despite the fact that we find substantial support using the data of our main theoretical channel, the paper could rationalize a couple of other empirical findings in the literature. Umultu et al. (2009) have found that the degree of financial liberalization reduces the aggregate stock return volatility. We argue in this paper that cross country differences in idiosyncratic risk could explain this fact. A rise in the capital income risk increases precautionary saving that enables the aggregate price of capital to depend mostly on the state of the safest country. Hence, this reduces the volatility of the endogenous aggregate risk. A recent paper by Alfaro et al. (2017) shows that using firm level data in Brazil to assess the effect of capital controls gives results at odd with the aggregate data assessment. We provide a theoretical channel that supports their empirical findings.

The rest of the paper is organized as follows. First, the model will be described. Second, the equilibrium conditions at individual and aggregate level will be given. The third part will focus on the resolution of the model and after that a quantitative analysis will be made. Then, we finish with a conclusion.

3.3. Model

Environment. There are two countries, indexed by $j \in \{A,B\}$, and a single good, which can be used for either consumption or investment purposes. Each country is populated by a continuum of infinitely lived households, indexed by i and distributed uniformly over [0,1]. Each household owns and runs a firm which can only use the capital stock invested by the particular household. Time is continuous, indexed by $t \in [0,\infty)$.

Preferences. Preferences are Epstein-Zin, so that, for a household i in country j, they are defined as the limit, for $\Delta t \to 0$, of the solution to the following recursive specification:

$$U_{ijt} = \left\{ (1 - e^{-\beta \Delta t})\, c_{ijt}^{1-\theta} \; + \; e^{-\beta \Delta t} \; \left(E_t \;\left[\, U_{ij\, t+\Delta t}^{1-\gamma} \, \right] \right)^{\frac{1-\theta}{1-\gamma}} \right\}^{\frac{1}{1-\theta}} , \tag{3.3.1}$$

where $\beta > 0$ is the discount rate, $\gamma > 0$ is the coefficient of relative risk aversion, and $\theta^{-1} > 0$ is the elasticity of intertemporal substitution.

Risk and asset structure. There are two assets and there is aggregate as well as idiosyncratic risk. The assets differ according to the type of risk they are subject to.

In each economy, a household can invest in those two assets. It can freely save or borrow in a bond—up to a natural borrowing constraint—and it can accumulate physical capital within its own family business. Let $k_{i,t}^{j}$ denote investment in capital and $b_{i,t}^{j}$ denote investment in the bond.

The return of the bond is subject to only the aggregate risk. Aggregate risk is introduced through dZ_t^j, a standard Wiener process, and it can be interpreted as an aggregate TFP shock. For example, the aggregate economy may be in a recession or in an expansion.

Capital investment is subject to the aggregate risk as well as to an undiversifiable idiosyncratic risk. Idiosyncratic risk is introduced through $dW_{i,t}^{j}$ and it is uninsurable because markets are incomplete. The reason for such incompleteness is left exogenous but one can think about it as resulting from frictions in financial markets. Literally taken, $dW_{i,t}^{j}$ represents a stochastic depreciation shock but it can actually be interpreted more broadly as encompassing various sources of idiosyncratic risk in the entrepreneurial activity and, more generally, in the returns to private investment. The idiosyncratic risk washes away in the aggregate. Private risk premium is the excess return needed for the household to invest in physical capital, rather than in the bond.

The household can change its risk exposure endogenously, via portfolio allocation and investment scale decisions.

Technology. Denote by $k_{i,t}^j$ the individual capital holdings of a household i in country j, and $\iota_{i,t}^j$ the average growth rate of $k_{i,t}^j$. $k_{i,t}^j$ is the "efficiency units" of capital for an individual agent at time t.[8] Then, a household uses capital to produce a flow of output, $y_{i,t}^j$, over a short period as follows:

$$y_{i,t}^j = (a - \Gamma(\iota_{i,t}^j))k_{i,t}^j \tag{3.3.2}$$

where A is the average value of the TFP shock and Γ is a standard convex-adjustment-cost technology with $\Gamma' > 0$ and $\Gamma'' > 0$.

The growth rate of $k_{i,t}^j$ is random and subject to the aggregate shock, introduced through the Wiener process dZ_t^j, and to the idiosyncratic shock introduced through the Wiener process $dW_{i,t}^j$:

$$\frac{dk_{i,t}^j}{k_{it}^j} = \iota_{i,t}^j dt + \sigma^j dZ_t^j + v_t^j dW_{it}^j \tag{3.3.3}$$

The assumption here is that the variance of the distribution of aggregate shocks is constant over time, at σ^j, whereas the variance of the distribution of idiosyncratic shocks, v_t^j, is time varying according to the process:

$$dv_t^j = \kappa(\bar{v}^j - v_t^j)dt + \sigma_v^j \sqrt{v_t^j} dZ_t^j \tag{3.3.4}$$

where $\bar{v}^j$ is the long-run value of this risk in country j, κ is a mean-reversion parameter and $\sigma_v^j < 0$ is the loading on the aggregate shock. So here, the variance of the distribution of idiosyncratic shocks is assumed to increase in recessions. This assumption leads to a countercyclical volatility which is in line with the strong evidence that micro uncertainty tends to increase sharply in recessions and falls in booms.[9] This formulation is quite general and enables a discussion for different cases. If there is no aggregate shock, then the variance of the idiosyncratic shock is constant. If the loading of the aggregate shock is 0, then the idiosyncratic shock is constant. In the remaining of the paper $\sigma(v^j) = \sigma_v^j \sqrt{v_t^j}$

Prices, returns and wealth. There is a competitive international market for capital. The price of capital, $P_t > 0$, evolves according to a process with drift $\mu_{p,t}$ and with diffusion that depends on the aggregate shocks of both countries:

$$\frac{dP_t}{P_t} = \mu_{p,t} dt + \sigma_{p,t}^A dZ_t^A + \sigma_{p,t}^B dZ_t^B \tag{3.3.5}$$

---

[8]The "efficiency unit" of capital $k_{i,t}^j$ could be viewed as: $k_{i,t}^j = A_t \kappa_{i,t}^j$ where $\kappa_{i,t}^j$ is the physical capital and A_t the TFP shock. Then aggregate shock to $k_{i,t}^j$ can be interpreted as persistent TFP shocks.

[9]See for example, Bloom (2014), Campbell et al. (2001), Kehrig (2015), Bloom et al. (2018), Bachmann and Bayer (2011) , Storesletten et al. (2004).

where $\sigma_{p,t}^A$ is the loading on the aggregate shocks in country A and $\sigma_{p,t}^B$ is the loading on the aggregate shocks in country B. These loadings are endogenous and will be determined in equilibrium.

Let $dR_{i,t}^j$ be the return from investing on unit of wealth in capital for an agent i in country j at time t. The return in capital holdings includes the profit gains ie the yield from output remaining after internal investment and the gains from the changes in the value of capital. Let $\bar{j}$ be the other country. This return evolves according to:

$$dR_{i,t}^j = \underbrace{\frac{a - \Gamma(\iota_{i,t}^j)}{P_t} dt}_{\text{Profit Gain}} + \underbrace{\left(\mu_{p,t} + \iota_{i,t}^j + \sigma_{p,t}^j \sigma^j\right) dt + \overbrace{(\sigma_{p,t}^j + \sigma^j)}^{\substack{\text{Domestic} \\ \text{Aggregate risk}}} dZ_t^j + \overbrace{\sigma_{p,t}^{\bar{j}}}^{\substack{\text{Foreign} \\ \text{Aggregate risk}}} dZ_t^{\bar{j}} + \overbrace{v_t^j}^{\substack{\text{Domestic} \\ \text{Idiosyncratic risk}}} dW_{i,t}^j}_{\text{Capital Gain}}$$

For each unit of capital they invest in, investors' returns are subject to an exogenous domestic aggregate risk σ^j, an endogenous domestic aggregate risk due to the market price of capital $\sigma_{p,t}^j$, an idiosyncratic risk v_t^j and an endogenous foreign risk $\sigma_{p,t}^{\bar{j}}$. Then, this specification helps also in capturing some spillover effects.

Let r_t the average return from investing in the bond.

Let $n_{i,t}^j$ denote household net worth.

The law of motion for net worth is:

$$dn_{i,t}^j = r_t b_{i,t}^j dt + P_t k_{i,t}^j dR_{i,t}^j - c_{i,t}^j dt \tag{3.3.6}$$

3.4. Equilibrium

The equilibrium is given by individual utility maximisations and a general equilibrium where agents prices and aggregate quantities are found.

3.4.1. Investor maximization

Investors choose optimally their level of investment, the share of their wealth they want to allocate to capital $\hat{k}_{i,t}^j$ and the share of their net worth they would like to allocate to consumption $m_{i,t}^j$. [10]

[10] This is similar as doing a guess (and later verify) that the policy functions for consumption and capital investment are proportional to net worth:

$$c_{i,t}^j = m_{i,t}^j \cdot n_{i,t}^j \tag{3.4.1}$$

$$k_{i,t}^j = \hat{k}_{i,t}^j \cdot n_{i,t}^j \tag{3.4.2}$$

The maximization problem for an investor i in country j, where $\bar{j}$ is the other country, is:

$$\max_{m^j_{i,t},\iota^j_{i,t},\hat{k}^j_{i,t}} U(c) \tag{3.4.3}$$

$$\frac{dn^j_{i,t}}{n^j_{i,t}} = (\mu^j_{i,n^j,t} - m^j_{i,t})dt + \sigma^j_{i,n^j,t}dZ_{jt} + \sigma^{\bar{j}}_{i,n^j,t}dZ_{\bar{j},t} + \tilde{\sigma}_{i,n^j,t}dW_{ijt} \tag{3.4.4}$$

where:

$$\mu^j_{i,n^j,t} = r_t + P_t\hat{k}^j_{i,t}(E_t[dR^j_{i,t}] - r_t) \tag{3.4.5}$$

$$\sigma^j_{i,n^j,t} = P_t\hat{k}^j_{i,t}(\sigma^j + \sigma^j_{p,t}) \tag{3.4.6}$$

$$\sigma^{\bar{j}}_{i,n^j,t} = P_t\hat{k}^j_{i,t}\sigma^{\bar{j}}_{p,t} \tag{3.4.7}$$

$$\tilde{\sigma}_{i,n^j,t} = P_t\hat{k}^j_{i,t}v^j_t \tag{3.4.8}$$

with a solvency constraint $n^j_{i,t} > 0$.

Investors receive a return from bond holdings and they require a premium $E_t[dR^j_{i,t}] - r_t$ to invest in the risky capital. Their net worth is affecting by the idiosyncratic shock as well as the aggregate shocks in the domestic country and the foreign country. But their exposure to those shocks will depend on the fraction of their net worth investing in capital. Following their leverage positions, they could borrow to raise funds and invest in capital or substitute the two assets by saving in bonds.

3.4.2. General equilibrium

Let Y^j_t, C^j_t, K^j_t, and B^j_t denote the aggregate levels of output, consumption, capital, and bond holdings in country j at date t (that is, the cross-sectional averages of $y^j_{i,t}$, $c^j_{i,t}$ and so on). We consider two policy regimes. In the first, countries are in financial autarchy: the bond cannot move across borders. In the second, the countries are “financially integrated": they can borrow and lend to one another using the bond. We define the corresponding equilibrium concepts as follows. Therefore, we will basically consider two situations: a closed capital account for bonds (financial autarchy) and an open capital account for bonds (financial integration).

Definition 1: Given K_0 an initial capital stock in the whole world, a competitive equilibrium in the *autarchy* regime is a sequence of country-specific interest rates and macroeconomic quantities, $\{r^j_t, P_t, Y^j_t, C^j_t, K^j_t, \chi^j_t\}_{t\in[0,\infty)}$ for $j \in \{A,B\}$, and a collection of individual contingent plans, $\{c^j_{i,t}, k^j_{i,t}, b^j_{i,t}, , y^j_{i,t}\}_{t\in[0,\infty)}$ for $i \in [0{,}1], j \in \{A{,}B\}$, such that the following are true: (i) individual plans are optimal given the sequences of prices; (ii) macroeconomic quantities are obtained by aggregating individual plans; (iii) all markets clear, i.e. $\int_i k^A_{i,t}di + \int_i k^B_{i,t}di = K_t$, $\int_i c^A_{i,t}di + \int_i c^B_{i,t}di = \int_i y^A_{i,t}di + \int_i y^B_{i,t}di$; (iv) bond markets clear at the country level, namely $\int_i b^j_{i,t}di = 0\ \forall j$, that is $B^j_t = 0$ for all $j{,}t$.

Definition 2: Given K_0 an initial capital stock in the whole world, a competitive equilibrium in the *integrated* regime is a sequence of world-wide interest rates, $\{r_t\}_{t\in[0,\infty)}$, and macroeconomic quantities, $\{P_t, Y^j_t, C^j_t, K^j_t, \chi^j_t\}_{t\in[0,\infty)}$ for $j \in \{A{,}B\}$, and a collection of individual contingent plans, $(\{c^j_{i,t}, k^j_{i,t}, b^j_{i,t}\}_{t\in[0,\infty)})_{i\in[0,1]}$ for $i \in [0{,}1], j \in \{A{,}B\}$, such that the following are true: (i) individual plans are optimal given the sequences of prices; (ii) macroeconomic quantities are obtained by aggregating individual plans;(iii) all markets clear, i.e. $\int_i k^A_{i,t}di + \int_i k^B_{i,t}di = K_t$, $\int_i c^A_{i,t}di + \int_i c^B_{i,t}di = \int_i y^A_{i,t}di + \int_i y^B_{i,t}di$; (iv) the bond market clears at the world level, namely $\int_i b^A_{i,t}di + \int_i b^B_{i,t}di = 0$ for all t.

The law of motion of the aggregate capital in country j:

$$\frac{dK^j_t}{K^j_t} = \int_i \iota^j_{i,t}dt + \sigma^j dZ^j_t \tag{3.4.9}$$

The remaining of this section first characterizes the individual household's problem for a given sequence of prices, and it then proceeds to characterize the general equilibrium under both regimes.

3.5. Solving the model

The solution will determine how the equilibrium price P_t, the allocation of capital and the agents' consumption decisions depend on the history of aggregate shocks. The procedure to solve for the equilibrium has two steps. First, use the conditions for utility maximization and for market clearing to derive the properties of the equilibrium processes. Second, show that the equilibrium dynamics can be characterized by a vector of three state variables and derive a system of equations that determine equilibrium variables as functions of the state vector.

3.5.1. Individual Quantities

We will Focus on country A. The Results will be analogous for country B. Individual problem is solved by using a Guess and verify method, assuming that the value function takes the following form:

$$V_t^A(n_{i,t}^A) = \frac{(\chi_t^A n_{i,t}^A)^{1-\gamma}}{1-\gamma} \tag{3.5.1}$$

Here, χ_t^A is positive and represents a net-worth multiplier that describes the marginal utility of wealth. We conjecture that χ_t^A follows an Ito process with drift $\mu_{\chi_t^A,t}^A$ and with loadings $\sigma_{\chi^A,t}^A$ and $\sigma_{\chi^A,t}^B$, respectively, on the aggregate shocks of country A and B:

$$\frac{d\chi_t^A}{\chi_t^A} = \mu_{\chi^A,t}^A dt + \sigma_{\chi^A,t}^A dZ_t^A + \sigma_{\chi^A,t}^B dZ_t^B \tag{3.5.2}$$

Proposition 1: *Like in Samuelson's and Merton's classic portfolio analysis, the investor consumption-investment problem reduces to optimal rules linear in wealth:*

$$c_{i,t}^A = m_t^A n_{i,t}^A \tag{3.5.3}$$

$$k_{i,t}^A = \hat{k}_t n_{i,t}^A \tag{3.5.4}$$

Where

$$m_t^A = \rho^{\frac{1}{\theta}} (\chi_t^A)^{\frac{\theta-1}{\theta}} \tag{3.5.5}$$

and $\hat{k}_t$ is determined using the risk premium condition:

$$E(dR_t^A) - r_t = \underbrace{(\sigma^A + \sigma_{p,t}^A)(\gamma\sigma_{n^A,t}^A - (1-\gamma)\sigma_{\chi^A,t}^A)}_{\substack{\textit{Domestic aggregate}\\ \textit{risk premium}}} + \underbrace{\sigma_{p,t}^B(\gamma\sigma_{n^A,t}^B - (1-\gamma)\sigma_{\chi^A,t}^B)}_{\substack{\textit{Foreign aggregate}\\ \textit{risk premium}}} + \underbrace{\gamma P_t \hat{k}_t^A (v_t^A)^2}_{\substack{\textit{Domestic idiosyncratic}\\ \textit{risk premium}}} \tag{3.5.6}$$

Moreover, the investment per level of capital is the same for all individual investor and is a function of the price of capital:

$$\Gamma'(\iota_t^A) = P_t \tag{3.5.7}$$

3.5.4 and (3.5.3) implies that investors decided to invest, save and consume the same share of their net worth in each country: the marginal propensity to consume and the share of net worth invested in capital did not depend on the level of individual specific net worth. Then, the economy could be properly written as function of net worth. In addition, with (3.5.7) the growth rate of the economy is given by the price of capital. Hence, the dynamics of the financial market will determine the dynamics of the real economy. Looking at (3.5.6), the risk premium in country A depends on the exposure of country A to its own fundamental

(σ^A) and endogenous aggregate shock (σ_p^A) times the term $\gamma\sigma^A_{n^A,t} - (1-\gamma)\sigma^A_{\chi^A,t}$. The later expression could be viewed as the market price of the aggregate risk in country A which is a weighted function of how the net worth and the stochastic investment opportunity sets are driven by the aggregate shock in country A. Similarly, the risk premium depends also of the endogenous aggregate shock in country B with a price $\gamma\sigma^B_{n^A,t} - (1-\gamma)\sigma^B_{\chi^A,t}$. On the top of that, there is the idiosyncratic risk that is also valued by a price $\gamma P_t \hat{k}^A_t v^A_t$. In an equilibrium in which the capital is a small share of the aggregate risk, this will tend to reduce the idiosyncratic premium and thus reducing the total risk premium, all else equal. From the risk premium condition, the share of net worth invested in capital is:

$$P\hat{k}^A = \frac{E(dR) - r + (1-\gamma)\sigma^A_{\chi^A,t}(\sigma^A + \sigma^A_{p,t}) + (1-\gamma)\sigma^B_{\chi^A,t}\sigma^B_{p,t}}{\gamma((v^A)^2 + (\sigma^A + \sigma^A_{p,t})^2 + (\sigma^B_{p,t})^2)} \tag{3.5.8}$$

All else being equal, a high level of risk aversion or a high level of the volatilities tends to decrease the fraction net worth invested in capital. Risk neutral agents will prefer to have high net worth when the χ^A is high. At the opposite, risk averse agents will prefer to have higher net worth when χ^A is low. Then for $\gamma < 1$, investors tend to be less risk averse and will have a high level of capital share (relatively high premium gives incitations to invest) and the opposite will be true for $\gamma > 1$. In the absence of time-varying shock, $P\hat{k}^A = \frac{E(dR)-r}{\gamma((v^A)^2}$. Each additional risk distorts the incentive to invest and agents are compensated by an additional term in the risk premium.

3.5.2. Aggregate quantities

The homothetic preferences and linearity of budget constraints induce an aggregation of quantities. The dynamic of the capital for the whole economy and the aggregate net worth in each country could be derived as ($j = A$,B):

$$\frac{dK_t}{K_t} = g_t dt + \sigma^A \frac{K^A_t}{K_t} dZ^A_t + \sigma^B \frac{K^B_t}{K_t} dZ^B_t \tag{3.5.9}$$

$$\frac{dN^j_t}{N^j_t} = (\mu_{n^j,t} - m^j_t)dt + \sigma^A_{n^j,t} dZ^A_t + \sigma^B_{n^j,t} dZ^B_t \tag{3.5.10}$$

Moreover, there is no need to keep track of the distribution of wealth across investors. Instead, what matters is the relative wealth of agents, the share of total wealth that belong to country A. Let x_t be the wealth share of country A:

$$x_t \equiv \frac{N^A_t}{P_t K_t} \tag{3.5.11}$$

$0 < x_t < 1$ and we guess that x_t follows the process:

$$dx_t = \mu_{x,t}dt + \sigma^A_{x,t}dZ^A_t + \sigma^B_{x,t}dZ^B_t$$

Then, there are four state variables from which two endogenous state variables: K_t the worldwide aggregate capital, x_t the share of aggregate wealth of country A; and two exogenous state variables: v^A_t the idiosyncratic risk in country A and v^B_t the idiosyncratic risk in country B.

Let s^A_t be the share of total capital K_t that belongs to country A: $s^A_t = \frac{K^A_t}{K_t}$. Using s^A_t, the system exhibits a scale invariance property. Indeed all the variables in the economy are de-trended with respect to aggregate net worth. Hence, the system is solved by looking for a markov equilibrium with the three state variables v^A_t, v^B_t and x_t:

$$P_t = P(v^A_t, v^B_t, x_t) \quad \chi^A_t = \chi^A(v^A_t, v^B_t, x_t) \quad \chi^B_t = \chi^B(v^A_t, v^B_t, x_t) \quad r_t = r(v^A_t, v^B_t, x_t) \quad s^A_t = s^A(v^A_t, v^B_t, x_t)$$

Proposition 2: *Using the evolution of the aggregate variables, the dynamic of the state variable x_t has the following components:*

- *In Financial Openness*

$$\mu_{x,t} = x_t[\mu_{n^A,t} - m^A_t - \iota_t - \mu_{p,t} - \sigma^A_{p,t}\sigma^A s^A_t - \sigma^B_{p,t}\sigma^B s^B_t + (\sigma^A_{p,t} + s^A\sigma^A)^2 +$$
$$(\sigma^B_{p,t} + s^B\sigma^B)^2 - \sigma^A_{n^A,t}(\sigma^A_{p,t} + s^A\sigma^A) - \sigma^B_{n^A,t}(\sigma^B_{p,t} + s^B\sigma^B)] \quad (3.5.12)$$

$$\sigma^A_{x,t} = x_t[\sigma^A_{n^A,t} - s^A\sigma^A - \sigma^A_{p,t}] \quad (3.5.13)$$

$$\sigma^B_{x,t} = x_t[\sigma^B_{n^A,t} - s^B\sigma^B - \sigma^B_{p,t}] \quad (3.5.14)$$

- *In Financial autarchy*

$$\mu_{x,t} = x_t(1 - x_t)[-x_t(\sigma^A)^2 + (1 - x_t)(\sigma^B)^2] \quad (3.5.15)$$

$$\sigma^A_{x,t} = x_t(1 - x_t)\sigma^A \quad (3.5.16) \qquad \sigma^B_{x,t} = -x_t(1 - x_t)\sigma^B \quad (3.5.17)$$

$\sigma^j_{x,t}$ is the loading of x_t with respect to aggregate shock in country j. If $\sigma^j_{x,t} > 0$, a negative shock in country j is associated with a reduction of aggregate wealth of country A: x_t will go down. x_t is then an important quantity that would determine the degree of the amplification of the shock. Using the state variables, we can define a Markov equilibrium.

Definition 3: A Markov Equilibrium in (v^A_t, v^B_t, x_t) in the financial autarchy state is a set of aggregate functions for price P_t, r^A_t, r^B_t quantities $\hat{k}_t$, ι_t, s^j_t χ^j_t, m^j_t, and a law of motion for the endogenous aggregate state variable $\mu_{x,t}$ and $\sigma^j_{x,t}$ such that:(i) χ^j solves the Hamilton-Jacobi-Bellman equation of investors in country j given P_t, r_t, $\mu_{x,t}$ and $\sigma^j_{x,t}$

; (ii) Good market clearing: $m_t^A x_t + m_t^B(1-x_t) = \frac{a-\Gamma(\iota_t)}{P_t}$; (iii) Capital market clearing: $s_t^A + s_t^B = 1$ and $p_t \hat{k}_t = 1$; (iv) Bond market clearing: $x_t = s_t^A$.

Definition 4: A Markov Equilibrium in (v_t^A,v_t^B,x_t) in the financial integration state is a set of aggregate functions for price P_t, r_t, quantities $\hat{k}_t$, ι_t, s_t^j χ_t^j, m_t^j, $\hat{B}_t^j = \frac{B_t^j}{K_t}$, and a law of motion for the endogenous aggregate state variable $\mu_{x,t}$ and $\sigma_{x,t}^j$ such that:(i) χ^j solves the Hamilton-Jacobi-Bellman equation of investors in country j given P_t, r_t, $\mu_{x,t}$ and $\sigma_{x,t}^j$; (ii) Good market clearing: $m_t^A x_t + m_t^B(1-x_t) = \frac{a-\Gamma(\iota_t)}{P_t}$; (iii) Capital market clearing: $s_t^A + s_t^B = 1$ and $p_t \hat{k}_t = \frac{s_t^A}{x_t}$; (iv) Bond market clearing: $\frac{\hat{B}_t^A}{P_t} = x_t - s_t^A$.

$\frac{s_t^A}{x_t}$ is the ratio of the asset of investors in country A over their net worth. Hence, $\frac{s_t^A}{x_t}$ have an interpretation in terms of balance sheets of the investors and could capture the strength of the balance sheet. Interestingly, in financial autarchy, this ratio is constant and equals to 1. However, in financial integration $\frac{s_t^A}{x_t}$ is a time varying object that depends on the level of the risk in the economy and the relative wealth share in the economy. Therefore, two degrees of freedom are present in financial integration that help to mitigate the risk: the share of aggregate wealth x_t and the strength of the balance sheet $\frac{s_t^A}{x_t}$. At the opposite, agents can only play around with the share of aggregate wealth x_t in financial autarchy. But, this is general feature of autarchy versus integration. The most important fact is the change in the sensitivity of those variables, that we will study in detail in the next section.

Exposure to domestic shock. Using Ito lemma, we can obtain an expression of $\sigma_{p,t}^A$ and $\sigma_{p,t}^B$:

$$\sigma_{p,t}^A = \frac{P_{vt}}{P_t}\sigma(v^A) + \frac{P_{xt}}{P_t}\sigma_{x,t}^A \qquad \sigma_{p,t}^B = \frac{P_{vt}}{P_t}\sigma(v^B) + \frac{P_{xt}}{P_t}\sigma_{x,t}^B$$

P_{vt} is the derivative of price with respect to the idiosyncratic shock and P_{xt} the derivative of price with respect to the relative wealth share.

Substituting the expressions of $\sigma_{p,t}^A$ and $\sigma_{p,t}^B$ give the following expression for $\sigma_{x,t}^A$:

$$\frac{\sigma_{x,t}^A}{x_t} = \underbrace{\frac{(1-x_t)\frac{s_t^A}{x_t}}{1-(s_t^A - x_t)\frac{P_{xt}}{P_t}}}_{\text{Aggregate Amplification}} \sigma^A + \underbrace{\frac{(\frac{s_t^A}{x_t}-1)\frac{P_{vt}}{P_t}}{1-(s_t^A - x_t)\frac{P_{xt}}{P_t}}}_{\text{Idiosyncratic Amplification}} \sigma(v^A)$$

If $\sigma_{x,t}^A > 0$, a bad shock on Z^A is translated into lower relative aggregate wealth and this amplify the adverse shock. Let's call *Agg* the aggregate amplification and *Idio* the idiosyncratic amplification. *Agg* is the amplification effect specific to the fundamental shock in the economy. At the opposite, *Idio* is exclusively due to the idiosyncratic risk. Both amplification factor contains the effect of the sensitivity of price to x_t (the “loss spiral" via x_t).[11] If this sensitivity is high, then additional amplifications occur via price adjustment. At the opposite, if it goes to 0, this reduces the amplification effect. *Idio* contained the direct effect of idiosyncratic risk on price (the “loss spiral" via v_t^A). But, the level of the leverage ratio $\frac{s_t^A}{x_t}$ could mitigate the effect of the losses and diminish the shock especially the idiosyncratic risk.

If there is no idiosyncratic risk or no time-varying idiosyncratic risk, then $\frac{\sigma_{x,t}^A}{x_t} = Agg * \sigma^A$ and all the amplification effects on the balance sheet of investors are coming from σ^A. This is the case in an economy with financial integration and only an aggregate shock. In that case, the sign of $\sigma_{x,t}^A$ and $\sigma_{1-x,t}^B$ are unambiguously positive or negative. $\sigma_{x,t}^A > 0$ if $1-(s_t^A - x_t)\frac{P_{xt}}{P_t} > 0$ and $\sigma_{x,t}^A < 0$ if $1-(s_t^A - x_t)\frac{P_{xt}}{P_t} < 0$. If there is a time-varying idiosyncratic risk, then the sign of *Idio* matters as well as the nature of the country. In financial autarchy, the sign of $\sigma_{x,t}^A$ and $\sigma_{1-x,t}^B$ is known, positive and constant. We have the following proposition.

Proposition 3 : *If P is increasing and sufficiently concave in x, then* $0 < 1-(s_t^A - x_t)\frac{P_{xt}}{P_t} < 1$.

With the quadratic adjustment cost function, the fact that P is increasing and concave in x is a natural outcome. Thus, this is a plausible assumption. With proposition 3, the sign of *Agg* is known but the sign of *idio* depends on the leverage effect.

Proposition 4: *In presence of time-varying capital income risk, the economy exhibits a mitigating effect, i.e., a cut in the concentration of aggregate risk on the balance sheet of investors in the riskiest country.*

For $\frac{s_t^A}{x_t} > 1$, $Idio > 0$ and additional amplification occurs via the uncertainty shock. For $\frac{s_t^A}{x_t} < 1$, $Idio < 0$ and the effect of the uncertainty shock is counteract by the adjustment of the balance sheet. The analysis is the same for $\sigma_{1-x,t}^B$. Hence, as long as country B is the relatively risky country, a bad uncertainty shock will be cut down by a reduction of the leverage effect in financial integration.

Therefore, we see that there is a mitigating effect for B, absent in the case of A. The sole presence of idiosyncratic risk is not sufficient to have the mitigating effect. The

[11] The concept of “loss spiral" was introduced by Brunnermeier and Pedersen (2008) and designed a two-way feedback loop where a reduction in net worth induces a drop in asset prices that further reduces the net worth by more and so on.

time-varying idiosyncratic risk is important for the riskiest country to counteract the effect of the bad shock over the business cycles. In financial integration, $\sigma^A_{x,t} > 0$ for lower values of v^A_t and $\sigma^A_{x,t} \to 0$ for high value of v^A_t. Indeed, $\sigma^A_{x,t}$ is the exposure of the relative wealth share of country A to the aggregate shock in country A. When it is positive, a bad shock in country A will tend to reduce the total wealth share of that country. But because of the opportunity to save via foreign bond holdings, when v^A_t is sufficiently high, agents in country A tend to have a great part of their wealth in bond holdings. As they become net saver, an additional increase in v^A_t, let them to save by more and so the direct effect of a bad shock on the balance sheet will be to mitigate the decrease of x_t. By contrast, when the economies are closed $s^A = x$, $\sigma^A_{x,t}$ is always positive and does not vary directly with v^A. There is no possible mitigation of a bad shock. Hence, the mitigation is only a property of the time-varying idiosyncratic risk in financial integration and is absent from all the other set-up (financial autarchy, integration and aggregate shock with or without fixed idiosyncratic risk).[12] In presence of idiosyncratic risk, openness helps to diversify away the risk with a reduction in the dependence of volatility of aggregate to idiosyncratic capital income risk. This is possible using the reallocation effect.

Proposition 5: *Reallocation effect: s^j_t decrease with v^j_t and the sensitivity of s^j_t to v^j_t increases as v^j_t increases:* $\frac{\partial \sigma^j_{s^j,t}}{\partial v^j_t} > 0$, $j \in \{A,B\}$.

Following a shock what is determinant is the value of $\sigma^j_{s^j,t}$. The fact that the exposure of s^j to the aggregate shock Z^j_t increases as the idiosyncratic risk increases means that there is a capital reallocation that occurs in the economy. The risky capital is allocated to the less risky country where investing in capital incurs less drawbacks. Investors are not willing to invest and thus proceed to a portfolio substitution between the risky asset and the safest asset. The presence of the uncertainty is so distortive that in the absence of an alternative foreign asset, investor avoid to invest and have no other diversification choice. In financial autarchy, where $s^A = x_t$, the direct effect of v^j_t on s^j_t disappears and the remaining effect is only the effect via x_t. Then, there is drop in the share of capital without a substantial portfolio reallocation. This has detrimental consequences on the net worth of investors.

[12] Even if we discuss the sign of $(s^A_t - x_t)\frac{P_{xt}}{P_t}$, the mitigation property is still a feature of the time-varying risk in an integration state. The only change is that the positive effect of the balance sheet in switching off the uncertainty shock is seen with the safest country.

Let us consider the following expression for $\sigma^A_{x,t}$:

$$\sigma^A_{x,t} = [(s^A_t - x_t)\sigma^A_{p,t} - x_t s^A_t \sigma^A]$$

Then we can express the direct effect of v^A_t on $\sigma^A_{p,t}$ keeping x_t constant:

$$\frac{\partial \sigma^A_{p,t}}{\partial_{v^A_t}} = \frac{\frac{\partial \sigma^A_{x,t}}{\partial v^A_t} - \frac{\partial s^A_t}{\partial v^A_t}\left(\sigma^A_{p,t} + (1-x_t)\sigma^A\right)}{s^A_t - x_t} \tag{3.5.18}$$

For high value of v^A_t ($s^A_t - x_t < 0$), $\sigma^A_{p,t}$ could decrease with v^A_t if $\frac{\partial s^A_t}{\partial v^A_t}$ is sufficiently high. That is the share of capital of country A must be very sensitive to the idiosyncratic risk in country A. This is what we proved in proposition 5. This type of risk affects the willingness to invest, so the level on investment is very sensitive to the idiosyncratic risk for risk averse agent. When the idiosyncratic risk increases, $\sigma^A_{p,t}$ decreases as a result of capital reallocation between both countries. Then, the country with the higher level of idiosyncratic risk will not affect as much the evolution of the price of capital because it holds a very low share of aggregate capital. Overall, the level of the price of capital will be higher compared to the autarchy state because of better risk sharing opportunities in openess. The portfolio reallocation that occurs when idiosyncratic risk evolves enables to diversify the risk and the agents are less subject to an increase in the idiosyncratic risk. A bad shock becomes less harmful compared to a the autarchy situation. Finally, they can reach a higher level of growth rate and a better investment.

Looking at the bonds, we can write them as the following:

$$B^A_t = (x_t - s^A_t)P_t K_t \tag{3.5.19}$$

First, when v^A_t increases, $x_t - s^A_t > 0$ and B^A_t increases in bad times. Second, this equation constitutes an endogenous borrowing contraint or a “natural borrowing contraint". In absence of this equality, agents will tend to increase their savings by more in bad times to mitigate the drop they could faced with the risky capital. This “constraint” is a function of their current net worth. When the idiosyncratic risk is high, the value of the collateral drop, and this fictitious constraint becomes more tighter. In economy where the idiosyncratic risk is naturally high, the constraint is tight because P_t and K_t drop. This create an endogenous constraint that limits the ability of country B to save in bonds as much as it will be willing to. This could increase the severity of a bad shock when comparing two integrated economies with different level of long run idiosyncratic risk.

Exposure to Foreign shock. Another interesting property of the model is the exposure to foreign shock. Consider the exposure of the balance sheet x_t to foreign shock.

$$\frac{\sigma^B_{x,t}}{x_t} = \underbrace{\frac{-s^B_t}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}}_{\text{Aggregate Amplification}} \sigma^B + \underbrace{\frac{(\frac{s^A_t}{x_t}-1)\frac{P_{vB_t}}{P_t}}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}}_{\text{Idiosyncratic Amplification}} \sigma(v^B)$$

Here again, the exposure to the aggregate foreign shock is composed of an aggregate amplification and an idiosyncratic amplification. for $1-(s^A_t-x_t)\frac{P_{xt}}{P_t} > 0$, the aggregate amplification is negative and the idiosyncratic risk is positive. The exposure of the safest country to the idiosyncratic risk of the riskiest country is reduced. But the opposite is true for the riskiest country: in fact, the time-varying risk will increase its exposure to the foreign shock. If the economy is in autarchy or in the integration stage with no time-varying capital income risk, this property disappears and all the countries are exposed to the same way to the foreign shock.

3.6. Computational

In this section, we will present the steps to solve the model, the choice of the parameters as well as the experiment we conduct.

3.6.1. Algorithm

All the variables in the model can be written as a function of: s^A, r_t, P_t, χ^A_t and χ^B_t in financial integration and s^A, r^A_t, r^B_t, P_t, χ^A_t and χ^B_t in financial autarchy. The economy is characterized by the following equations: HJB in country A, HJB in country B, risk premium equation in country A, risk premium equation in country B, the capital good market and the consumption good market clearing condition.

With the capital good market, the share of aggregate capital in country B is expressed as a function of the share of aggregate capital in country A. In financial openness, we have to find s^A. Mixing the risk premium equations in both countries, this gives a polynomial equation of degree 5. From Abel theorem in algebra, we know that there is no analytical solution for this type of equations. To solve it, a numerical method combining both Newton method and bisection method is implemented in order to improve at the same time the speed and the accuracy. Then, one of the risk premium equation is used to compute the interest rate. In financial autarchy, s^A is determined by x. So the two risk premium equations will be used to determine their respective interest rate r^A and r^B. Plug in the expressions of s^A and r or r^A and r^B in the remaining equations, the whole system could be summarized as a set of three second order stochastic partial differential equations in v^A, v^B, x. The unknown are P_t, χ_t^A and χ_t^B and the equations are the HJB in country A, the HJB in country B and the good market clearing condition. Those equations are highly non linear and cannot be solved analytically, or even as usual parabolic, elliptic and hyperbolic partial differential equations. Then, they are transformed into a system with false transient equations. This is done by adding a fictitious final time step T so that the time dimension becomes an additional state variable.

Finally, the three variables are written as function of the state variables: $P(v^A,v^B,x,t)$, $\chi^A(v^A,v^B,x,t)$ and $\chi^B(v^A,v^B,x,t)$. Computing the drifts and the variances of the different objects using Ito Lemma, we have to deal with three spatial derivatives in v^A, v^B and x_t and one temporal derivative.[13] The spatial derivatives are discretized by a finite difference method of level 2. Finally, we have a system with only a time derivative that we can solved using a robust method for differential equation, for instance Runge Kunta 4.[14] Notice that the good market clearing condition is an algebraic equation which differentiated taking the derivative with respect to time. The solution is found by iterating backward until we reach a step where the time derivatives shrank to zero. One can start with an arbitrary final point, the most important requirement is that the consumption good market clearing conditions should be satisfied. After finding the evolution of all the aggregates with respect to the state variables, we could simulate the dynamic of the model.

3.6.2. Numerical example

The calibration is based on standard parameter values of the literature. The discount rate is set to 5% to match an average risk free rate of 1%. The EIS is set to be greater than 1, quite common in model with stochastic volatilities (Bansal et al. (2014), Campbell and Beeler (2009)). We use a conservative value of 1.5. A bench of studies has found values for relative risk aversion in OECD countries ranging from 1-10 (Vissing-Jørgensen and Attanasio (2003), Barsky et al. (1997), Cohen and Einav (2005), Guiso and Paiella (2005), Dohmen et al. (2005), Palsson (1996)). We choose 5 as a middle value that helps also to generate relevant risk premia. The volatility of TFP shock is 1.25%. The adjustment cost function have the simple convex form: $\Gamma(\iota) = A(\iota+\delta)^2 + B(\iota+\delta)$. The depreciation rate is set to 0.05. The parameter of the function enables to match the average investment to GDP ratio and the average growth rate of GDP.

[13]The spatial derivatives include first order derivative, second order derivatives and cross derivatives.
[14]Alternative method could be collocation algorithm, fine element method.

The calibration of the process of idiosyncratic risk is the same as in DI Tella (2017) with a mean reversion parameter of 1.38 and an exposure to aggregate shock of -0.17. The long run value of idiosyncratic risk will be set to 0.25 for the country with the lower level of risk and 0.40 in the alternative specification to characterize the country with the highest level of idiosyncratic risk. Notice that there is no empirical measure of the uninsurable idiosyncratic risk. But based on some studies Koren and Tenreyo (2007) for example, we know that less developed countries have higher idiosyncratic risk compared to developed one. Also, the idiosyncratic volatilities of stock market return can contained a part of risk which is diversifiable, or even a “good volatility” and another part which is undiversifiable.[15] Alternative parameter values will be discussed in the next section for more robustness to the results.

3.6.3. Aggregates and dynamics

We conduct the following experiment. We first consider a closed bond market (financial autarchy) and an open bond market (financial openness). We look at the changes in the results by taking into account different specifications. A first case with no idiosyncratic capital income risk and just volatility of aggregate shock, a second case with fixed idiosyncratic capital income risk and a third case with the time-varying idiosyncratic capital income risk as in our benchmark.

We make the following difference calculation: we compute the variation of the value of the variable in financial openness compared to the value of the variable in financial autarchy for each type of specification. This enables to show how taking into account time varying volatility could create a difference between financial openness and financial autarchy. We calculate the following: $\frac{Var^{FO}-Var^{FA}}{Var^{FA}}$.

When there is no idiosyncratic risk, the model show a very small difference between autarchy and openness. But the fact of adding this risk even when it is fixed, increased the difference between open and closed economy. With time-varying capital income risk, the value of the asset price in openness reach until 12% of the value in the closed economy. Then, in presence of time-varying idiosyncratic risk, there is a sizeable difference moving from closed to openness. The idea is that, if the gains from openness in previous studies were so small, maybe this was because both environment where very closed to each other, so no great difference is observed, there were almost no gains from moving to one. Consequently,

[15]Bartram, Brown and Stulz (2012) explained that the observed high idiosyncratic volatility in US stock market return compared to similar economy is due to the investor protection, stock market development, new patents, and firm-level investment in research and Development.

Fig. 1. Comparision of price and growth dynamics

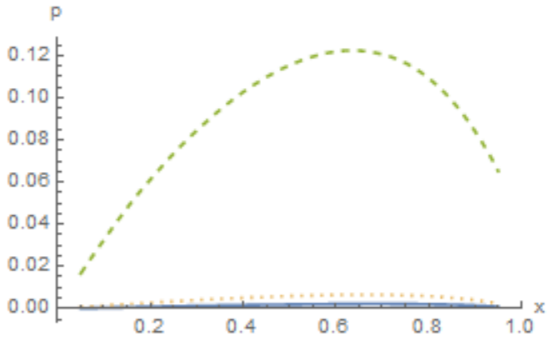

This figure shows the difference of the integrated economy versus the closed economy; Blue line: No idiosyncratic risk and just a difference in TFP volatility; Dotted line: Idiosyncratic risk but not time varying and no TFP volatility; Dashed line: Time varying idiosyncratic risk.

this new theoretical finding could bring new insights in assessing the gains or the losses of financial openness.

Figure 2 shows the evolution of different volatilities as the idiosyncratic risk changes.

Fig. 2. Exposure of relative wealth share and price of capital to aggregate shock in A

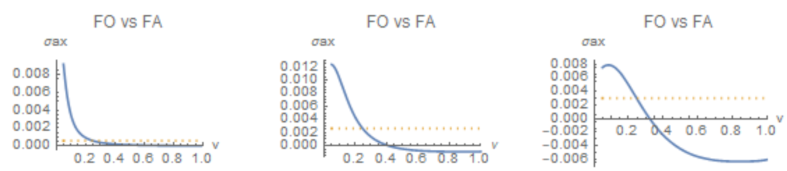

(a) $\sigma^A_{x,t}$

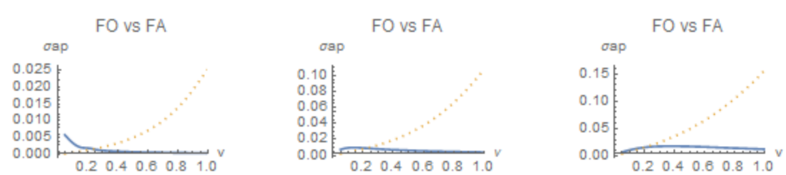

(b) $\sigma^A_{p,t}$

Those figures show the exposure of relative wealth share to aggregate shock in A and Volatility of Price of capital with time varying idiosyncratic risk for different values of the relative wealth share in country A: left X=0.05; Middle X=0.3; Right X=0.6. Blue line: Financial openness; Dotted line: financial autarchy.

Openness reduces the volatility of asset prices, and then the exposure to domestic shock. In fact, the decision of investing in capital is very sensitive to the presence of the idiosyncratic risk. So when the later increases, there is a reallocation of capital from the more risky to the less risky country which occurs at the same time with a positive net bond position for the riskiest country. Thus, the volatility of price of capital is less subject to the variation in country A. In other words, that enables to reduce the exposure of the price of capital to the risky country. When the idiosyncratic risk increases in openness, the sensitivity of the price of capital to that risk decreases, and at the opposite, it increases in autarchy. This reallocation helps to diversify the risk. The country end up with a better position in financial openness. The mechanism can be summarized as followed:

- Openness: $s_t^A \neq x_t$, $\forall t$

$$\downarrow Z^A \Rightarrow \begin{cases} \uparrow v^A \Rightarrow \underbrace{\uparrow \sigma^A_{s^A} \Rightarrow \downarrow s^A}_{\text{Portfolio reallocation}} \Rightarrow \downarrow \sigma_p^A, \sigma_n^A, \sigma_x^A \\ \\ \text{Effect on x} \end{cases}$$

- autarchy: $s_t^A = x_t$, $\forall t$

$$\downarrow Z^A \Rightarrow \begin{cases} \uparrow v^A \Rightarrow \text{No direct effect on } s^A \\ \\ \text{Effect on x} \end{cases}$$

In integration, the reallocation channel is very high. This effect triggers the decline in the volatilities and the exposure of investor in country A net worth to domestic shock is reduced. In autarchy, the reallocation channel almost disappears and the volatilities of the aggregate variable are still increasing with the risk.

Further, looking at the risk premium equation:

$$\mu_p + \iota + \sigma^A \sigma_p^A + \frac{a-\Gamma}{P} - r = (\sigma^A + \sigma_p^A)(\gamma \sigma^A_{n^A} - (1-\gamma)\sigma^A_{\chi^A}) + \sigma_p^B(\gamma \sigma^B_{n^A} - (1-\gamma)\sigma^B_{\chi^A}) + \gamma \frac{s^A}{x}(v^A)^2 \tag{3.6.1}$$

When v^A increases, all else equal, agents requires a higher risk premium to invest in capital so that the price of capital falls. But the strength of the idiosyncratic risk is altered or amplified by the leverage ratio $\frac{s^A}{x}$. For a rising v^A, the leverage ratio drops in openness and the effect of idiosyncratic risk is reduced :

$$\underbrace{\frac{s^A}{x}}_{\downarrow} \overbrace{v^A}^{\uparrow}$$

This will reduce the drop of the price of capital as v^A increases. But in financial autarchy, $s^A = x$ and as the reallocation channel disappears, the leverage effect disappear too. Then, the risk premium increases unambiguously and the effect on asset prices is higher.

$$\underbrace{\frac{s^A}{x}}_{=} \overbrace{v^A}^{\uparrow}$$

All those effects lead to figure 4.

Fig. 4. Price of capital with time varying idiosyncratic risk

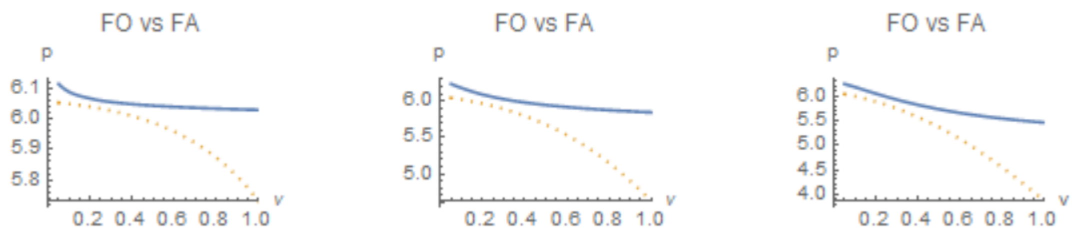

Price of capital with time varying idiosyncratic risk for different values of the relative wealth share in country A: left X=0.05; Middle X=0.3; Right X=0.6. Blue line: Financial openness; Dotted line: financial autarchy.

Therefore, the presence of foreign bonds in openness triggers a reduction in the exposure to domestic shock and also a wealth effect that helps the country to invest more than it will do in financial autarchy. This is why the price of capital will be higher in openness compared to autarchy and the country ends up with a greater level of growth in financial openness.

Consider the other specifications with no time-varying idiosyncratic risk. The elasticity of capital share to idiosyncratic risk which produced this strong portfolio restructuring effect is shut down when there is no more time-varying investment risk. The volatility of s^A is equalled to: $\sigma^A_{s^A} = (1 - s^A)\sigma^A$. And the effect of x on s^A is almost the same whenever we are in open economy or in autarchy [16]. Moreover, as we omit the aggregate shock, all the volatility effect disappears.

But it is worth noticing that if openness reduces the exposure to domestic shock, nonetheless it raises the exposure to foreign shock as shown by figure 5.

16 $\frac{\partial \sigma^A_{s^A}}{\partial x} = -\sigma^A \frac{\partial s^A}{\partial x}$ and $\frac{\partial s^A}{\partial x} = 1$ in autarchy

Fig. 5. Exposure of relative wealth share to aggregate shock in B

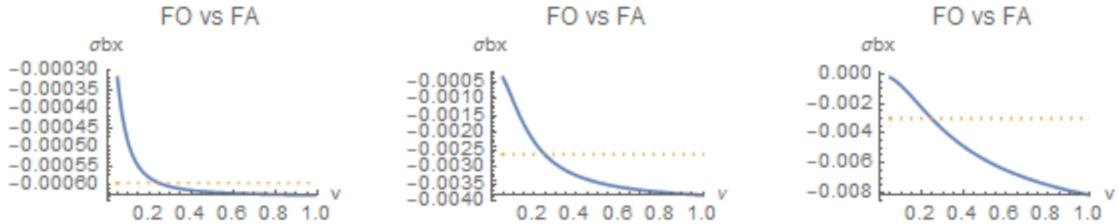

Exposure of relative wealth share to aggregate shock in B with time varying idiosyncratic risk for different values of the relative wealth share in country A: left X=0.05; Middle X=0.3; Right X=0.6. Blue line: Financial openness; Dotted line: financial autarchy.

The inter-linkage of economies brought by financial openness increases the reaction of a domestic economy to foreign shock, and then increases the risk of systematic crises. But the whole volatilities of the shocks are reduced.

3.6.4. Distributions and welfare

We look at the welfare of being in a financially integrated regime compared to the one of being in a financial autarchy regime. More specifically, we seek to answer the question, what is the minimal compensation an agent in the financial autarchy regime would need to be as well as an agent in the financial integration regime?

Fig. 6. Welfare evolution with wealth share

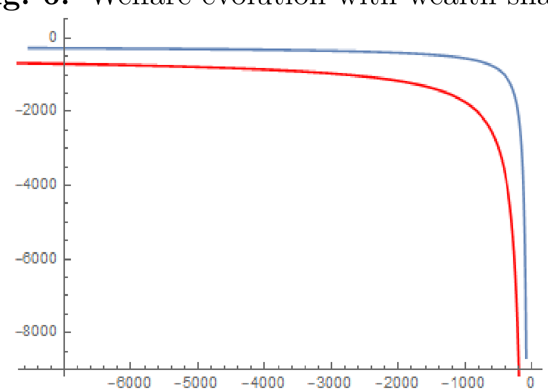

Welfare of More risky country (horizontal axis) versus Less risky country (vertical axis) with wealth share for $v^A = 0.2$, $v^B = 0.25$. Blue line: Financial integration; Red line: financial autarchy.

Figure 6 plots the value functions for agents in country A on the horizontal axis and for agents in country B on the vertical axis. The red line is for the financial autarchy and the blue line for financial integration. Both value functions are expressed as a function of

x for v^A and v^B at their long run value. The two curves are very distinct and the financial integration value functions outweigh the one in autarchy for every value of x. Agents are better off in integration.

Fig. 7. Sum of the Welfare evolution with wealth share

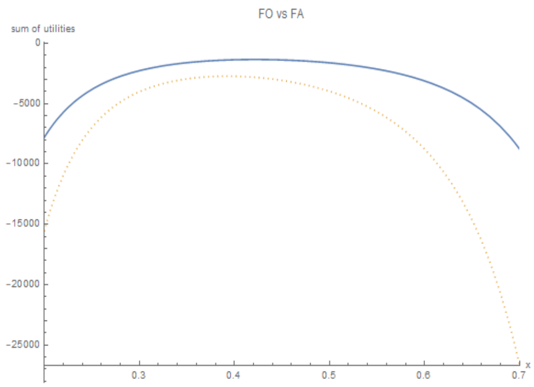

Sum of the Welfare of More risky country and Less risky country (vertical axis) with wealth share for $v^A = 0.2$, $v^B = 0.25$. Blue line: Financial integration; Red line: financial autarchy.

Figure 7 gives a more precise view of the cumulative gains. Financial integration is welfare improving for the whole economy, for each value of the relative wealth share. Absolute capital controls are welfare reducing in the presence of time-varying risk. A quantification of this welfare is presented in figure **??**.

Fig. 8. Percentage compensation

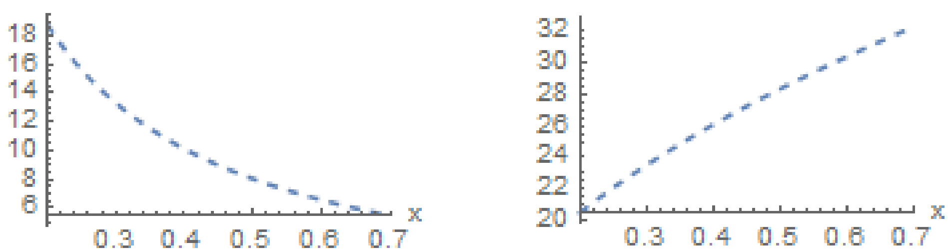

Percentage compensation for Less risky country (left graph) and More risky country (right graph) with wealth share for $v^A = 0.2$, $v^B = 0.25$.

The riskiest country is the one who benefits the most from the integration: the gains ranging from 20% to 30%. The relatively safe country gains from 6% to 16%. Those values are quite high but it is important to keep in mind that they captured the sole effect of uncertainty shock. Alternative specification and taking into account other macroeconomics concepts will reduce the overall effect of integration. The main message here is that in the

Table 1. Welfare comparison

Relative value	Investment risk-no aggregate	Investment risk-aggregate	Investment risk-Aggregate-link
Less risky	1	1.5	-1.12
More risky	1	1.03	1.5

This table shows the welfare comparison for the less risky and the more risky country for different states of the economy: model with investment risk and no aggregate risk; model with Investment risk and aggregate risk but no time-varying idiosyncratic risk; model with Investment risk, Aggregate risk and a link between the Investment risk and the aggregate risk. The model without the time varying idiosyncratic risk is taken as a reference.

presence of uncertainty shock, financial integration improve risk sharing between agents and is beneficial especially for the more risky country. The investors are enable to diversify away their risks and reach a higher level of utility.

Because we want to exhibit the particular role of time varying idiosyncratic risk, the results are expressed in comparison with model without the time varying idiosyncratic risk.

The usual measurement of the welfare gains or losses of financial integration ignore the heterogeneous risks that agents faces. Although, the economy with a high level of risk will face a low level of capital and production, the risk sharing opportunity is still highly valuable in presence of an uncertainty shock.[17]

3.7. Empirical Evidence

As demonstrated above, the main theoretical results of this paper stem from three essential mechanisms: the price effect, the reallocation effect, and the wealth effect. In this section, we will test with data whether the previous channels are actually at work in the presence of idiosyncratic risk and financial openness. Namely, we will test the relationship between aggregate price volatility and idiosyncratic risk in the presence of financial openness. In a second step, we will empirically analyze the relationship between the capital reallocation ratio and the idiosyncratic risk in the presence of financial openness. Finally, we will see

[17] Additional robustness analyzes were conducted to see the effect of varying parameters. Among other things, an increase in idiosyncratic risk in the riskiest country increases the gains from the financial integration described above. An increase in the correlation between idiosyncratic risk and the aggregate shock amplifies the reallocation effect in financial integration and further distorts the price of capital in financial self-sufficiency. An increase in the risk aversion parameter reduces the willingness to invest in risky assets. Investors value investment in risk-free assets much more. In the same vein as Di Tella (2017), a coefficient of intertemporal elasticity closer to 1 tends to make the price effect disappear.

if the idiosyncratic risk favours relatively higher capital growth in integration compared to autarchy.

3.7.1. Data Description and Methodology

The data used comes from the Orbis database of Bureau van Dijk which generates data on companies from around the world. These data include firms publicly but also private firms. The data are monthly and cover the period from January 2002 to December 2017. We exclude firms that are not listed on the stock exchange since data on stock prices are unavailable and countries with less than 50 firms present in the database. data. There are therefore a total of 6,368,128 observations (Country-Companies-Months-Years). Inventory prices are denominated in US dollars and correspond to closing prices. 72 countries are therefore displayed in the sample as shown in Table XX.

Returns are calculated according to the method of Umultu et al. (2009) and Campbell et al. (2001). Let i be a stock, j a country and t is for the time. The return on the world market portfolio $R_{w,t}$ is the weighted (the weights are given by $w_{j,t}$) average return of the country level market portfolio $R_{j,t}$: $R_{w,t} = \sum_j w_{j,t} R_{j,t}$. Similarly, the return on the country market portfolio is the weighted average return of individual stocks in that country: $R_{j,t} = \sum_{i \in j} w_{i,t} R_{i,j,t}$, where $R_{i,j,t}$ is the individual stock return i in country j at time t and $w_{i,t}$ is the weight of the individual stock return in country j. Therefore, according to Umultu et al. (2009) and Campbell et al. (2001), the market model in an international framework is:

$$R_{i,j,t} = R_{j,t} + \epsilon_{i,j,t} \tag{3.7.1}$$

and

$$R_{j,t} = R_{w,t} + \eta_{j,t} \tag{3.7.2}$$

Since our study focuses on idiosyncratic risk, it is not necessary to estimate the equation 3.7.2. In fact, the country level risks and the world wide risk are of a minor interest for our study.

The aggregate idiosyncratic risk will be:

$$v_{j,t}^2 = \sum_{i \in j} w_{i,t} (\sum_{\tau \in t} \epsilon_{i,j,t}^2) \tag{3.7.3}$$

Further, we follow Fogli and Perri (2015) in computing the relative value of the aggregate idiosyncratic risk as follows:

$$\tilde{v}_{j,t}^2 = v_{j,t}^2 - \frac{1}{N} \sum_{k \neq j} v_{k,t}^2 \tag{3.7.4}$$

The endogenous aggregate risk evaluated by the growth rate of the price of capital will be the aggregate return risk:

$$\sigma_{j,t}^2 = \sum_{i \in j} w_{i,t} (\sum_{\tau int} (R_{i,j,t} - \bar{R}_{i,j,t})^2) \tag{3.7.5}$$

Where $\bar{R}_{i,j,t}$ is the mean return of stock i in country j at time t.

To compute the weights, we use market capitalisation. For example, $w_{i,t}$ will be the market capitalisation of stock i in country j at time t over the total market capitalisation of country j.

Due to data restrictions, we will be unable to use the actual liberalization dates of individual countries to measure the date of financial liberalization. Moreover, because of the uncertainties about the precise period for dating those events, using dates remains the least preferred method in the literature. Thus, as a proxy for financial liberalization, we will use the index of Chin and Ito (2007). This index captures financial openness related to the relaxation of government restriction measures. We think that this index corresponds better to the notion of financial integration as analyzed in the theoretical model. The Chin and Ito (2007) index comes in two forms. The first is a categorical variable taking 5 values from -2 to 2 with a step of 1. The second is a binary variable 0 or 1. In each of the definitions, from the smallest value to the highest one, we move from a state of financial autarchy to financial integration.

For validation purposes, we will also use a measure of financial integration based on capital movements. The database of Lane and Milesi-Ferretti (2015) will enable such a calculation. This database describes the annual capital movements of the different countries and covers the period from 1970 to 2015 (when available). To much our theoretical part, we will consider gross bonds flows measured as bond assets and liabilities as a share of GDP. This variable is continuous and measures the intensity of financial trade with the rest of the world.[18]

We complete the analysis with other control variables suggested by Umultu et al. (2009). The size of the domestic market is measured by the total market capitalization of the stock market in a country to the GDP. To account for the market liquidity effect, we also include the total value of shares traded in the market during the period, divided by the average market capitalization for the period. Depending on the specification, we also include other growth, monetary or government-related variables from the World Bank database such as GDP, inflation, trade openness, government consumption. Countries are also categorized as

[18]Alternatively, using lagged variables to minimize the endogeneity bias produces the same qualitative results.

advanced economies, emerging market and low-income countries following the IMF MAC DSA (Market access country- Debt sustainability analysis) and LIC DSA (Market access country- Debt sustainability analysis) decomposition.

Table 2 to 4 present some descriptive statistics for the sample period. On average, advanced economies tend to have lower idiosyncratic risk compared to emerging market and low income countries. The same goes for the aggregate return volatility. These countries are also more open, have a relatively larger market capitalization (1.5 time more), and have a higher fraction of world GDP (a typical advanced economy weighs twice as much in world GDP than a less advanced economy). Due to the availability of data and the choice of weighting, the values found for the aggregate return risk and the aggregate idiosyncratic risk are of an order of magnitude lower than those observed in the literature. However, compared to Pukthuanthong-Le and Visaltanachoti (2009), for example, the relative values (when comparing countries) are the same.

Table 2. Descriptives statistics: Advanced Economies

country	Aggregate Idiosyncratic volatility	Relative Aggregate Idiosyncratic volatility	Aggregate return volatility	Degree of openness (IC)	Relative Market capitalisation	GDP/GDP world	N
Australia	25.59	0.40	38.33	1.00	0.037	0.017	2420
Austria	6.16	-16.26	11.13	1.00	0.003	0.006	119
Belgium	8.56	-15.46	10.20	1.00	0.005	0.007	203
Canada	26.54	2.85	38.67	1.00	0.023	0.024	4562
Cyprus	20.09	-4.04	28.72	1.00	0.001	0.000	167
Denmark	11.93	-12.70	16.93	1.00	0.003	0.005	240
Finland	7.64	-16.49	9.35	1.00	0.006	0.004	167
France	11.78	-12.11	14.87	1.00	0.013	0.042	1150
Germany	11.91	-13.87	15.87	1.00	0.026	0.055	1262
Greece	14.28	-8.54	24.56	1.00	0.007	0.004	319
Hong Kong	17.57	-5.71	21.75	1.00	0.007	0.004	283
Ireland	14.41	-8.66	19.42	1.00	0.002	0.004	130
Israel	12.22	-11.51	19.93	1.00	0.010	0.003	696
Italy	11.51	-11.95	17.12	1.00	0.007	0.032	419
Japan	9.57	-14.36	12.55	1.00	0.106	0.081	4272
Korea, Republic of	20.77	-6.08	26.85	0.60	0.048	0.017	2181
Luxembourg	8.92	-15.21	15.30		0.001	0.001	101
Netherlands	8.51	-14.20	12.91	1.00	0.004	0.013	275
New Zealand	6.91	-16.92	9.14	1.00	0.006	0.002	215
Norway	14.66	-9.47	19.59	1.00	0.005	0.007	247
Portugal	10.46	-11.70	15.41	1.00	0.001	0.003	78
Singapore	14.43	-9.65	21.10	1.00	0.013	0.003	814
Slovakia	5.31	-14.87	6.71	1.00	0.000	0.001	92
Slovenia	4.31	-19.48	6.49	1.00	0.001	0.001	84
Spain	5.59	-17.39	9.05	1.00	0.008	0.022	2698
Sweden	10.20	-12.40	15.69	1.00	0.008	0.008	611
Switzerland	5.73	-16.81	9.88	1.00	0.007	0.009	386
United Kingdom	9.65	-11.13	12.44	1.00	0.137	0.043	2952
United States	22.43	1.00	25.22	1.00	0.175	0.244	12681
Total	12.33	-11.13	17.42	0.99	0.023	0.023	1373
Total	11.51	-12.11	15.69	1.00	0.007	0.007	319

The table shows some descriptives statistics for the advanced economies. The mean of the variable is taken over the sample of study (2002-2017).N denotes the number of firms in the database for a given country. The different volatilities are estimated in (‰).

Table 3. Descriptives statistics: Emerging Markets

country	Aggregate Idiosyncratic volatility	Relative Aggregate Idiosyncratic volatility	Aggregate return volatility	Degree of openness (IC)	Relative Market capitalisation	GDP/GDP world	N
Argentina	10.86	-11.86	20.13	0.07	0.002	0.006	77
Brazil	25.28	3.18	31.26	0.29	0.007	0.027	394
Bulgaria	29.80	5.67	35.92	1.00	0.001	0.001	244
Chile	8.26	-15.07	11.88	1.00	0.005	0.003	177
China	11.61	-10.61	25.26	0.00	0.115	0.089	2922
Croatia	10.52	-13.61	15.22	1.00	0.001	0.001	178
Egypt	20.62	-3.21	34.96	0.55	0.006	0.003	255
Hungary	7.67	-16.46	11.16	1.00	0.000	0.002	60
India	14.56	-6.97	23.37	0.00	0.039	0.022	4159
Indonesia	34.09	8.84	46.66	0.46	0.011	0.010	546
Iran	25.69	3.04	33.94	0.00	0.004	0.006	308
Jordan	25.50	5.58	30.58	1.00	0.002	0.000	242
Kuwait	12.31	-11.90	15.60	1.00	0.006	0.002	206
Malaysia	13.00	-13.34	23.11	0.07	0.016	0.004	1167
Mauritius	6.08	-18.05	7.95	1.00	0.001	0.000	102
Mexico	12.90	-10.25	16.09	1.00	0.004	0.016	127
Mongolia	64.51	50.37	75.22	1.00	0.000	0.000	187
Morocco	6.60	-16.75	9.29	0.00	0.006	0.001	85
Nigeria	27.10	5.00	38.15	0.00	0.003	0.006	196
Oman	19.63	-4.21	26.58	1.00	0.002	0.001	101
Pakistan	13.61	-10.14	17.89	0.00	0.009	0.003	556
Peru	12.53	-11.19	18.77	1.00	0.002	0.002	93
Philippines	26.26	0.80	35.91	0.00	0.005	0.003	263
Poland	33.07	17.43	40.76	0.21	0.006	0.007	959
Romania	19.97	-4.17	27.03	1.00	0.001	0.003	770
Russian Federation	30.16	6.03	40.53	0.70	0.004	0.025	280
Saudi Arabia	9.56	-14.57	17.57	1.00	0.008	0.009	171
South Africa	8.19	-15.57	11.60	0.00	0.009	0.005	437
Sri Lanka	15.67	-8.46	19.81	0.00	0.003	0.001	284
Thailand	11.54	-10.66	14.86	0.00	0.013	0.005	667
Tunisia	6.39	-18.09	8.45	0.00	0.002	0.001	76
Turkey	19.91	-1.12	27.07	0.00	0.007	0.011	456
Ukraine	61.15	37.02	79.37	0.00	0.001	0.002	163
United Arab Emirates	23.35	-0.78	26.41	1.00	0.004	0.005	106
Vietnam	11.56	-11.90	19.25	0.00	0.003	0.002	758
Total (mean)	19.70	-3.31	26.79	0.47	0.009	0.008	508
Total (median)	14.56	-8.46	23.37	0.29	0.004	0.003	244

The table shows some descriptives statistics for the emerging markets. The mean of the variable is taken over the sample of study (2002-2017).N denotes the number of firms in the database for a given country. The different volatilities are estimated in (‰).

Table 4. Descriptives statistics: Low Income Countries

country	Aggregate Idiosyncratic volatility	Relative Aggregate Idiosyncratic volatility	Aggregate return volatility	Degree of openness (IC)	Relative Market capitalisation	GDP/GDP world	N
Bangladesh	12.37	-10.86	18.56	0.00	0.004	0.002	301
Bermuda	27.36	6.74	34.35		0.023	0.000	805
Cayman Islands	24.83	5.66	31.45		0.025	0.000	1200
Kenya	24.38	2.17	29.97	1.00	0.002	0.001	62
Macedonia	5.82	-19.19	13.88		0.000	0.000	139
Nepal	397.49	412.20	904.61	0.00	0.002	0.000	165
Taiwan, Province of China	10.15	-14.30	14.80		0.040	0.007	1841
Virgin Islands, British	32.50	10.25	37.92		0.002	0.000	136
Zimbabwe	79.88	59.70	451.73	0.00	0.000	0.000	71
Total (mean)	68.31	50.26	170.81	0.25	0.011	0.001	524
Total (median)	24.83	5.66	31.45	0.00	0.002	0.000	165

The table shows some descriptives statistics for the Low Income Countries. The mean of the variable is taken over the sample of study (2002-2017).N denotes the number of firms in the database for a given country. The different volatilities are estimated in (‰).

Table 5. Effect of Idiosyncratic risk in presence of financial liberalization

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Aggregate return risk				Relative market capitalisation	GDP/GDP world	Investment/GDP
Idiosyncratic risk	50.282***	50.203***	53.233***	51.474***	-0.00387	0.028	-0.374**
	(4.449)	(4.451)	(3.595)	(3.686)	(0.059)	(0.026)	(0.168)
Gross debt	0.007	0.006					0.016
	(0.008)	(0.008)					(0.032)
Idiosyncratic risk* Gross debt	-0.191**	-0.213**					0.822***
	(0.077)	(0.082)					(0.235)
CI index			-0.046	-0.024	0.00107	-0.002***	
			(0.059)	(0.056)	(0.003)	(0.000)	
Idiosyncratic risk* CI index			-2.251*	-2.685*	-0.09552*	-0.086**	
			(1.179)	(1.345)	(0.054)	(0.032)	
N	520	520	596	531	397	375	412
R-sqr	0.879	0.880	0.880	0.891	0.958	0.980	0.811
Liquidity		Y		Y			
size		Y		Y			
GDP growth					Y	Y	Y
Inflation					Y	Y	Y
Gross domestic saving/GDP					Y	Y	Y
Credit to private sector/GDP					Y	Y	Y

The table shows the results of the estimation of equation 3.7.6 . Financial liberalization is measured by either Chin and Ito (2007) index or by the gross debt which is the sum of debt asset and liabilities. Robust standards errors are clustered at country level. All the specifications includes country and time fixed effects. Aggregate return risk is in estimated in logarithm. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

3.7.2. Empirical specification and results

In each estimate, we want to see how financial integration mitigates (or amplifies) the effects of idiosyncratic risk. Let y_{jt} be a dependant variable that could be either the aggregate return risk, the relative market value or GDP share, or the Investment to GDP ratio. Let $IdioRisk_{jt}$ the variable capturing the idiosyncratic risk, $Finlib_{jt}$ the one capturing financial liberalization. We estimate the following equation is estimated:

$$y_{jt} = \beta_0 + \beta_1 IdioRisk_{jt} + \beta_2 Finlib_{jt} + \alpha Z_{jt} + country_j + time_t + \nu_{jt} \tag{3.7.6}$$

Where $country_j$ is a set of country fixed effects, $time_t$ a set of time fixed effects, Z_{jt} are a set of controls variables that could include the size of the domestic market as measured by the market capitalisation over GDP, the liquidity of the domestic market, some macroeconomic variables (GDP growth, saving rate, domestic credit to private sector,inflation).

The results of the estimates are presented in Table 5. A variation of one unit of idiosyncratic risk (more than the average) increases the aggregate risk by 51% in an economy that is not in a phase of financial liberalization. On the other hand, this a one-unit increase in the idiosyncratic risk increase by 48% the aggregate risk in an economy which is in a phase of financial liberalization. The difference of around 3% is negative and statistically significant. This estimate is linked to the theoretical result which postulated that the raise of idiosyncratic risk in financial integration reduced the endogenous aggregate risk compared to a situation of financial autarchy. Further, the relative value of market capitalization (country market capitalization / global market capitalization), as well as the fraction of GDP are used proxies for the variable sa^t of the theoretical model. The table shows that in the presence of financial openness, this variable decreases much more in the event of a change in idiosyncratic risk. Finally, the effect of idiosyncratic risk on investment is negative and estimated at -0.374 in financial autarchy. However, this variation is greatly attenuated in the presence of financial openness. The cross coefficient is positive and significant. All these results confirm the main channel which governs the theoretical results of the previous section.

3.8. Conclusion and Policy Implications

This paper introduces uncertainty shock in a two-country general equilibrium model. It also investigates how taking this shock into account modifies the predictions on the potential effects of financial integration. It focuses on the uncertainty shock at the individual level which induces an ex-post heterogeneity between the agents. Compared to studies where the idiosyncratic risk is fixed, the variation of such risk causes a portfolio reallocation that allows agents to choose between risky and non-risky (or less risky) assets at each period. In such a scheme, the access or not to a less risky asset is greatly valuable.

In a two-country neoclassical growth model with aggregate uncertainty, the gains from risk sharing are offset by the losses of capital allocation. However, by introducing a time-varying idiosyncratic risk, first, this paper brings a distinction between the ability to invest (difficulty to invest even if the agent wants too) and the willingness to invest (fear of uncertainty that makes the agent reluctant to invest). In such a case, a lack of investment opportunities reduces the price of the risky asset because of fire sales and weakens the balance sheet of agents. Financial openness therefore allows agents facing a high risk to reallocate their portfolio in favour of the less risky asset. At the same time, the least risky country invests in capital, while the riskiest country recovers a financial strength through the less risky asset (wealth effect). Finally, a marginal increase in financial wealth gives an incentive to the riskiest country to invest much more in capital than it would have done in a situation of financial autarchy. Clearly, risk-sharing gains remain substantial despite the reallocation of capital. Three major effects enable to achieve such a result: a portfolio allocation effect, a price effect and a wealth effect.

In addition, the integrated economy exhibits less volatility than the autarchy one. Nevertheless, it becomes more vulnerable to external shocks. Welfare gains, however, are positive and dominated by gains in terms of risk sharing.

Therefore, the model has important implications for the problematic of financial integration. Business cycles costs of idiosyncratic capital income risk are high. Then, financial integration appears to be very useful in presence of this risk even at the expense of an exposure to the foreign shock. The riskiest country can build a stabilization policy using the trading of the financial assets with the safest country. Capital controls, thus, may be distortive with uncertainty shocks. As a policy, it seems better to reduce the long run value of idiosyncratic risk (helping investors to diversify away parts of their risk) instead of imposing capital control. Finally, this study calls for more cautiousness from policy makers in attempting to limit capital movement without considering heterogeneity at the individual production side level.

Bibliography

Abrams, D. S., U. Akcigit, and J. Popadak, 2013, Patent value and citations: Creative destruction or strategic disruption? Working Paper 19647, National Bureau of Economic Research.

Abrevaya, J., Hamermesh, D. S., 2012. Charity and favoritism in the field: Are female economists nicer (to each other)?. Review of Economics and Statistics, 94(1), 202-207.

Agénor, P-R., 2003. Benefits and costs of international financial markets. World Econ., 26 (2003), pp. 1089-1118.

Aghion, P., Bacchetta, P., Banerjee, A., 2004. Financial development and the instability of open economies. Journal of Monetary Economics 51 (2004) 1077–1106.

Albert, M., and S. Laberge, 2007, The legitimation and dissemination processes of the innovation system approach: the case of the Canadian and Quebec science and technology policy, Science, Technology and Human Values.

Alfaro, I., Bloom, N., Lin, X., 2016. The Finance-Uncertainty Multiplier. Working Paper.

Alfaro, L., Chari, A., Kanczuk, F., 2017. The real effects of capital controls: Firm-level evidence from a policy experiment. Journal of International Economics 108: 191–210.

Amara, N., and R. Landry, 2005, Sources of information as determinants of novelty of innovation in manufacturing firms: Evidence from the 1999 Statistics Canada innovation survey, Technovation.

Angeletos, M., 2007. Uninsured idiosyncratic investment risk and aggregate saving. Review of Economic Dynamics 10 (1): 1–30.

Angeletos,M. , Panousi, V., 2011. Financial Integration, Entrepreneurial Risk and Global Dynamics. Journal of Economic Theory, vol. 146:863-896.

Angrist, J., Azoulay, P., Ellison, G., Hill, R., Feng Lu, S., 2017. Economic Research Evolves: Fields and Styles. American Economic Review: Papers and Proceedings 107(5). 293-297.

Angrist, J., Azoulay, P., Ellison, G., Hill, R., Feng Lu, S., 2017. Inside job or deep impact? Using extramural citations to assess economic scholarship. National Bureau of Economic Research.

Antecol, H., Bedard, K., Stearns, J., 2018. Equal but inequitable: who benefits from gender-neutral tenure clock stopping policies?. American Economic Review, 108(9), 2420-41.

Antecol, H., Cobb-Clark, D. A., 2013. Do psychosocial traits help explain gender segregation in young people's occupations?. Labour Economics, 21, 59-73.

Arellano, C., Bai, Y., Kehoe, P., Forthcoming. Financial Frictions and Fluctuations in Volatility. Journal of Political Economy.

Audretsch, D., Keilbach, M., Lehmann, E., 2006. Entrepreneurship and Economic Growth. Oxford University Press, Oxford.

Auriol, E., Friebel, G., Wilhelm, S., 2019. Women in European Economics. Mimeo.

Bachmann, R., Bayer, C., 2013. "Wait-and-See" business cycles?. Journal of Monetary Economics 60: 704–719.

Bachmann, R., Moscarini, G., 2012. Business Cycles and Endogenous Uncertainty. Working paper.

Bai, Y., Zhang, J., 2017. Financial integration and international risk sharing. Journal of International Economics 86: 17–32.

Basu, S., Bundick, B., 2017. Uncertainty Shocks in a Model of Effective Demand. Econometrica, May 2017, 85 (3), 937–958.

Basu, S., Fernald, J., Shapiro, M., 2001. Productivity Growth in the 1990s: Technology, Utilization, or Adjustment?. NBER Working Paper No. 8359. Cambridge, MA: NBER.

Bayer, A., Rouse, C., 2016. Diversity in the Economics Profession: A New Attack on an Old Problem. Journal of Economic Perspectives, 30(4): 221-42.

Bengui, J., Mendoza, E., Quadrini, V., 2013, Capital mobility and international sharing of cyclical risk. Journal of Monetary Economics 60 (2013), pp. 42-62.

Benhima, K., 2013. A reappraisal of the allocation puzzle through the portfolio approach. Journal of International Economics 89(2): 331-346.

Blau, F. D., DeVaro, J., 2007. New evidence on gender differences in promotion rates: An empirical analysis of a sample of new hires. Industrial Relations: A Journal of Economy and Society, 46(3), 511-550.

Bloom, N., 2007. The impact of uncertainty shocks. Econometrica, 77(3): 623–685.

Bloom, N., 2014. Fluctuations in Uncertainty. Journal of Economic Perspectives 28, 153–176.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S.J., 2018. Really Uncertain Business Cycles. Econometrica 86, 1031-1065.

Boadway, R., and J.-F. Tremblay, 2017, Policy Forum: The uneasy case for a Canadian Patent Box, Canadian Tax Journal.

Bornmann, L., Butz, A., Wohlrabe, K., 2018. What are the top five journals in economics? A new meta-ranking. Applied Economics, 50(6), 659-675.

Broner, F. A., Ventura, J., 2016. Rethinking the effects of financial liberalization. The Quarterly Journal of Economics, 131(3): 1497-1542.

Brunnermeier, M., Pedersen, L., 2008. Market Liquidity and Funding Liquidity. Review of Financial Studies 22(6): 2201-2238.

Brunnermeier, M., Sannikov, Y., 2014. A Macroeconomic Model with a Financial Sector. American Economic Review 104 (2): 379–421.

Brunnermeier, M., Sannikov, Y., 2015. International Credit Flows and Pecuniary Externalities. American Economic Journal: Macroeconomics 7(1): 297–338.

Bubb R., 2013, The evolution of property rights: State law or informal norms?, The Journal of Law and Economics.

Buch, C.M., Döpke, J., Pierdzioch, C., 2005. Financial openness and business cycle volatility. Journal of International Money and Finance 24(5): 744–765.

Buera, F., Shin, Y., 2017. Productivity Growth and Capital Flows: The Dynamics of Reform. American Economic Journal: Macroeconomics, 9(3): 147–185.

Calabrese, T., J. A. C. Bau, and B. S. Silverman, 2000, Canadian biotechnology start-ups, 1991–1997: The role of incumbents' patents and strategic alliances in controlling competition, Social Science Research.

Canadian Intellectual Property Office, 2016 (a), Intellectual property and Canada's innovation agenda, 2016, Innovation, Science and Economic Development Canada.

Canadian Intellectual Property Office, IP Canada Report, 2016 (b).

Card, D., DellaVigna, S., 2013. Nine Facts About Top Journals in Economics. Journal of Economic Literature 51(1):144-161.

Card, D., DellaVigna, S., Funk, P., Iriberri, N., 2019. Are referees and editors in economics gender neutral? Quarterly Journal of Economics (forthcoming).

Ceci, S. J., Ginther, D. K., Kahn, S., Williams, W. M., 2014. Women in academic science: A changing landscape. Psychological Science in the Public Interest, 15(3), 75-141.

Chari, A., Goldsmith-Pinkham, P., 2017. Gender representation in economics across topics and time: Evidence from the NBER summer institute (No. w23953). National Bureau of Economic Research.

Chin, M.D., Ito, H., 2007. A new measure of financial openness. Journal of Comparative Policy Analysis: Research and Practice 10, 309–322.

Christiano, L. , Motto, R., Rostagno, M., 2014. Risk Shocks. American Economic Review 104(1): 27–65.

Coeurdacier, N., Rey, H., Winant, P., Forthcoming. Financial integration and growth in a risky world. Journal of Monetary Economics.

Cogneau, D., and A. Moraldi, 2014, Borders that divide: Education and religion in Ghana and Togo since colonial times, Journal of Economic History.

Colacito, R., Croce, M., 2011. The Short and Long Run Benefits of Financial Integration. The American Economic Review Papers and Proceedings, 100(2): 527-531.

Colussi, T., 2017. Social Ties in Academia: A Friend is a Treasure. Review of Economics and Statistics , forthcoming.

Corneli, F., 2017. Medium and long run implications of financial integration without financial development. Working Paper Bank of Italy.

Council of Canadian Academics, 2013, The state of industrial R&D in Canada, Expert panel on the state of industrial *R&D* in Canada.

Croce, M., 2013. Welfare Costs in the Long Run. UNC Kenan-Flagler Research Paper 2013-3.

D'Ippoliti, C. (2017). 'Many-Citedness': Citations Measure More Than Just Scientific Impact. Institute for New Economic Thinking Working Paper Series, (57).

Dahl, G., Kotsadam, A., Rooth, D. O., 2018. Does integration change gender attitudes? The effect of randomly assigning women to traditionally male teams (No. w24351). National Bureau of Economic Research.

Davis, J., Van Wincoop, E., 2018. Globalization and the increasing correlation between capital inflows and outflows. Journal of Monetary Economics 100: 83-100.

Decker, R., D'Erasmo, P., Boedo, M., 2016. Market Exposure and Endogenous Firm Volatility over the Business Cycle. American Economic Journal: Macroeconomics, 8(1):148-198.

Devereux, M. B., Yu, C., 2014. International financial integration and crisis contagion. NBER Working Paper.

Devereux, M., Sutherland, A., 2009. A portfolio model of capital flows to emerging markets. Journal of Development Economics 89: 181–193.

Di Tella, S., 2017. Uncertainty Shocks and Balance Sheet Recessions. Journal of Political Economy 125(6): 2038-2081.

Dietz, L., Bickel, S., Scheffer, T., 2007. Unsupervised prediction of citation influences. In Proceedings of the 24th International Conference on Machine Learning , pp. 233–240.

Dion, M., Sumner, J., Mitchell, S. M., 2018. Gendered citation patterns across political science and social science methodology fields. Political Analysis 26(3)312–327.

Doloreux, D., 2004, Regional innovation systems in Canada: A comparative study, Technology and Science.

Doloreux, D., and S. Parto, 2005, Regional innovation systems: current discourse and unresolved issues, Technology and Science.

Ductor, L., Goyal, S., Prummer, A., 2018. Gender and collaboration. Working Paper, School of Economics and Finance, Queen Mary University of London, No. 856.

Ellison, G., 2002. The slowdown of the economics publishing process. Journal of Political Economy 110, 947-993.

Ellison, G., 2002. Evolving standards for academic publishing: A q-r theory. Journal of Political Economy 110, 994-1034.

Ellison, G., 2011. Is Peer Review in Decline?. Economic Inquiry 49(3), 635-657.

Ellison, G., 2013. How Does the Market Use Citation Data? The Hirsch Index in Economics. American economic Journal: Applied Economics 5(3), 63-90.

Engemann, K., Wall, H., 2009. A journal ranking for the ambitious economist. Federal Reserve Bank of St. Louis Review, 91(3), 127-139.

Feng, F., 2017. Savings Gluts and Asset Price Booms in Emerging Economies. Working paper.

Ferber, M., 1986. Citations: Are They an Objective Measure of Scholarly Merit? Signs 11 (2), pp. 381–389.

Ferber, M., 1988. Citations and Networking. Gender and Society 2 (1), pp. 82–89.

Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramírez, J., Uribe, M., 2011. Risk Matters: The Real Effects of Volatility Shocks. The American Economic Review, 101(6): 2530-2561.

Fogli, A., Perri, F., 2015. Macroeconomic volatility and external imbalances. Journal of Monetary Economics 69: 1–15.

Fong, E.A., Wilhite, A.W., 2017. Authorship and citation manipulation in academic research. PLoS ONE 12(12): e01897394.

Fortin, M.-J., and M. Hadfield, 2016, Repenser l'innovation hors métropole : L'action publique dans le secteur bioalimentaire, Groupe de Recherche Interdisciplinaire sur le Développement Territorial, de l'Est du Quebec (GRIDEQ) et de le Centre de Recherche sur le Développement Territorial (CRDT).

Gibson, J., Anderson, D. L., Tressler, J., 2017. Citations or Journal Quality: Which Is Rewarded More in the Academic Labor Market?. Economic Inquiry 55 (4): 1945–65.

Gilchrist, S., Sim, J. W., Zakraje, E. 2014. Uncertainty, Financial Frictions, and Investment Dynamics. Working Paper 20038, National Bureau of Economic Research.

Ginther, D. K., Kahn, S., 2004. Women in economics: moving up or falling off the academic career ladder?. Journal of Economic perspectives, 18(3), 193-214.

Gourinchas, P-O., Jeanne, O., 2006. The Elusive Gains from International Financial Integration. The Review of Economic Studies 73, 715–741.

Gourinchas, P-O., Jeanne, O., 2013. Capital Flows to Developing Countries: The Allocation Puzzle. Review of Economic Studies 80: 1484–1515.

Gourio, F., Siemer, M., Verdelhan, A.,2013. International risk cycles. Journal of International Economics 89(2): 471-484.

Gourio, F., Siemer, M., Verdelhan, A., 2015. Uncertainty and international capital flows. Working Paper, MIT Sloan.

Greenspon, J., and E. Rodigues, 2017, Are trends in patenting reflective of innovative activity in Canada?, Center for the Study of Living Standards, Research Report.

Griliches, Z.,1990, Patent statistics as economic indicators: A survey. Journal of Economic Literature.

Hall, L.A. and S. Bagchi-Sen, 2001, A study of R&D, innovation, and business performance in the Canadian biotechnology industry, Technovation.

Hall, B. and R. Ziedonis, 2001, The patent paradox revisited: An empirical study of patenting in the US semiconductor industry, 1979-1995. The RAND Journal of Economics.

Hamermesh, D., 2018. Citations In Economics: Measurement, Uses, and Impacts. Journal of Economic Literature 56, 115-156.

Hamermesh, D., Pfann A., 2012. Reputation and Earnings: The Roles of Quality and Quantity in Academe. Economic Inquiry 50 (1):1–16.

Heckman, J., Moktan, S., 2018. Publishing and Promotion in Economics: The Tyranny of the Top Five.Journal of Economic Literature (forthcoming).

Hengel, E., 2017. Publishing while Female. Are women held to higher standards? Evidence from peer review.

Hengel, E., Moon, E., 2019. Gender and quality at top economics journals.

Hilmer, M., Ransom, M., Hilmer, C., 2015. Fame and the Fortune of Academic Economists: How the Market Rewards Influential Research in Economics. Southern Economic Journal 82 (2): 430–52.

Holness, P., 2016, Linking Canadian holders of US patents to Statistics Canada's business register: 2000 to 2011, Statistics Canada.

Hospido, L., Sanz, C., 2019. Gender gaps in the evaluation of research: evidence from submissions to economics conferences.

Hoxhaa, I., Kalemli-Ozcan, S., Vollrath, D., 2013. How big are the gains from international financial integration?. Journal of Development Economics 103: 90-98.

Hunt, J., 2016. Why do women leave science and engineering?. ILR Review, 69(1), 199-226.

Impact Center, 2017, Canada's patent puzzle: Do we have a problem with securing patents or with the commercialization of them?, An impact brief, University of Toronto.

Jensen, P., Rouquier, J-B., Croissant,Y., 2009. Testing Bibliometric Indicators by Their Prediction of Scientists Promotions. Scientometrics 78 (3): 467–79.

Johnson, D. K. N., 2002, The OECD technology concordance: Patents by industry of manufacture and sector of use, OECD Publications.

Kalaitzidakis, P., Mamuneas, T., Stengos, T., 2003. Rankings of Academic Journals and Institutions in Economics. Journal of the European Economic Association 1(6): 1346-1366.

Kalaitzidakis, P., Mamuneas, T., Stengos, T., 2011. An updated ranking of academic journals in economics. Canadian Journal of Economics, 44(4), 1525-1538.

Kaminsky, L., Schmukler, S., 2008. Short-Run Pain, Long-Run Gain: Financial Liberalization and Stock Market Cycles. Review of Finance 12: 253–292.

Kehrig, M., 2015. The Cyclical Nature of the Productivity Distribution. Working paper.

Kelly, B., Papanikolaou, D., Seru, A., Taddy, M., 2018. Measuring technological innovation over the long run (No. w25266). National Bureau of Economic Research.

King, R., Levine, R., 1993. Finance, entrepreneurship, and growth. Journal of Monetary Economics,32: 513-542.

Kiyotaki, N., Moore, J., 1997. Credit Cycles. Journal of Political Economy, 105: 211-248.

Kodrzycki, Y., Yu, P., 2006. New approaches to ranking economics journals. The BE Journal of Economic Analysis and Policy, 5(1).

Koffi, M., Panousi, V., 2019. Patents, Innovation and Growth in Canadian Pharmaceuticals.

Kogan, L., Papanikolaou, D., Stoffman, N. (2020). Left behind: Creative destruction, inequality, and the stock market. Journal of Political Economy, 128(3), 855-906.

Kosnik, L. R., 2015. What have economists been doing for the last 50 years? A text analysis of published academic research from 1960–2010. working paper.

Kraay, A., Ventura, J., 2000. Current Accounts in Debtor and Creditor Countries. The Quarterly Journal of Economics, 115(4): 1137-1166.

Laband, D. N., 2013. On the Use and Abuse of Economics Journal Rankings. Economic Journal 123(570): F223–54.

Laband, D. N., Piette, M. J., 1994. The relative impacts of economics journals: 1970-1990. Journal of economic Literature, 32(2), 640-666.

Lampe, R., 2012. Strategic citation. Review of Economics and Statistics, 94(1), 320–333.

Lewis, K. , 2000. Why do stocks and consumption imply such different gains from international risk sharing? J. Int. Econ. 52, 1–35 .

Lewis, K. , Liu, E.X. , 2015. Evaluating international consumption risk sharing gains: an asset return view. J. Monet. Econ. 71, 94–98.

Lexchin, J., 1993, Pharmaceuticals, patents, and politics: Canada and Bill C-22, International Journal of Health Services.

Ludvigson, S., Ma, S., Ng, S., 2016. Uncertainty and business cycles: Exogenous impulse or endogenous response? Working Paper.

Lundberg, S. J., Stearns, J, 2019. Women in Economics: Stalled Progress. Journal of Economic Perspectives 33 (1), pp. 3–22.

Mendoza, E., Quadrini, V., and Rìos-Rull, J-V., 2007. On the Welfare Implications of Financial Globalization Without Financial Development. NBER working paper.

Michalopoulos, S. and E. Papaioannou, 2013, Pre-colonial ethnic institutions and contemporary African development, Econometrica.

Michalopoulos, S. and E. Papaioannou, 2014, National institutions and subnational development in Africa, Quarterly Journal of Economics.

Miguel, Satyanath and Sergenti, 2004, Economic shocks and civil conflict: An instrumental variables approach, Journal of Political Economy.

Moss-Racusin, C., Dovidio, J., Brescoll, V., Graham, M., Handelsman, J., 2012. Science faculty's subtle gender biases favor male students. Proceedings of the National Academy of Sciences, 109(41), 16474-16479.

Moyse, P.-E., 2015, Towards increased innovation: exploring the effectiveness of sovereign patent funds in Canada, Research Seminar in Intellectual Property Policy in the Making.

Obstfeld, M., 1994. Risk-Taking, Global Diversification, and Growth. The American Economic Review, 84(5): 1310-1329.

Panousi, V., Papanikolaou D., 2012. Investment, Idiosyncratic Risk and Ownership. Journal of Finance 67: 1113–1148.

Panousi, V., 2012. Capital taxation with entrepreneurial risk.

Pinkovskiy, M.L., 2013, Economic discontinuities at Borders: Evidence from satellite data on lights at night.

Pinkovskiy, M.L., 2017, Growth discontinuities at borders, Journal of Economic Growth.

Preston, A. E., 1994. Why have all the women gone? A study of exit of women from the science and engineering professions. The American Economic Review, 84(5), 1446-1462.

Romer, P. M. (1986). Increasing returns and long-run growth. Journal of political economy, 94(5), 1002-1037.

Sandri, D., 2014. Growth and Capital Flows with Risky Entrepreneurship. American Economic Journal: Macroeconomics 6(3): 102-123.

Sarsons, H., 2019. Gender differences in recognition for group work. working paper.

Schumpeter, J., 1911, Theorie der wirtschaftlichen Entwicklung (Dunker and Humboldt, Leipzig); translated by Redvers Opie, 1934, The theory of economic development (Harvard University Press, Cambridge, MA).

Schwanen, D., 2017, Innovation policy in Canada: A holistic approach, Institute C.D. Howe, Commentary.

Storesletten, K., Telmer, C. I. and Yaron A., 2004. Cyclical dynamics in idiosyncratic labor market risk. Journal of Political Economy, 112(3): 695-717.

Tallarini, T., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45: 507-532.

Teele, D. L., Thelen, K., 2017. Gender in the journals: Publication patterns in political science. PS: Political Science and Politics, 50(2), 433-447.

Umutlu, M., Akdeniz, L., Altay-Salih, A., 2010. The degree of financial liberalization and aggregated stock-return volatility in emerging markets. Journal of Banking and Finance 34: 509–521.

Wei Dou, W., 2016. Embrace or Fear Uncertainty: Growth Options, Limited Risk Sharing, and Asset Prices. Working Paper.

Wilhite, A. W., Fong, E., 2012. Coercive citation in academic publishing. Science 335:542–543.

Wu, A., 2018. Gendered Language on the Economics Job Market Rumors Forum. AEA Papers and Proceedings 108: 175–79.

Zafar, B., 2013. College major choice and the gender gap. Journal of Human Resources, 48(3), 545-595.

Zhu, X., Turney, P., Lemire, D., Vellino, A., 2015. Measuring academic influence: Not all citations are equal. Journal of the Association for Information Science and Technology, 66 (2) (2015), pp. 408-427.

Appendix A

Appendix to Chapter 1

Main variables

A description of the main variables is given below.

Patent application number. This is a unique number that can be used to identify each patent. Patents are numbered consecutively based on filing data, starting from patent number 1 in 1868 growing to about 3,000,000 for the most recently available patents in 2017. Patent numbers higher than 2,000,000 indicate patent applications filed after October 1, 1989, i.e. under the most recent patent act (New Act or NA).

International Patent Classification (IPC) codes. These codes indicate the industry of the firm owning each patent, according to a unified international classification system, same as for the US. These codes are A for human necessities; B for performing operations and transporting; C for chemistry and metallurgy; D for textiles and paper; E for fixed constructions; F for mechanical engineering, lighting, heating, weapons, and blasting; G for physics; and H for electricity.

Language. About 5 percent of patent applications are in French and 95 percent are in English. In the analysis that follows, we will focus on applications filed in English.

Patent status and dates. The processing and maintenance of a patent application proceeds through a number of states from filing to expiration. States include application, examination, granting, transfer, expiration or abandonment of a patent.

License for sale indicator. This element indicates if the owner is willing to sell or license the rights to the patent. This variable takes on the values true or false.

Ownership, assignees, or applicants. All current and previous owners. Name and address of a current or previous owner, original or preserved name of the owner, owner address. Owner enabling date, on which the owner received all or part of the ownership. Owner disabling date, on which the owner stopped having ownership. Owner grantee indicator indicates if the owner is the grantee of the patent Y = yes, N = no. Nationality and residence of each owner.

Inventors. Name and address of the inventors. Original preserved name of the inventor or the name of inventor. Inventor address contains the name and address of the inventor(s), including detailed street, zip-code, province or state, and country information.

Patent agent. This element contains the address the name and address of the current registered patent agent of record representing the owner.

Filing date. The date on which the patent application was filed.

Grant date. The date on which the patent application was granted.

Abstract. Short, non-technical, textual description of the patent and proposed innovation.

Short description. Two-to-five page textual description of each patent. They vary in terms of technical language.

Patent claims. Textual documents, usually dozens of pages long, with technical descriptions of the patent. They outline the extent or scope of the protection conferred by a patent, or the protection sought in a patent application. Their purpose is to define which subject-matter is protected by the patent or sought to be protected by the patent application. This is termed as the "notice function" of a patent claim, to warn others of what they must not do if they are to avoid infringement liability. The claims are of the utmost importance during prosecution and litigation.

Long description. Extended textual documents, averaging hundreds of pages, with the complete description of each patent. These documents contain complicated, technical and scientific language. They also include chemical formulas of molecular structures or drawings of the proposed innovations.

World Intellectual Property Organization (WIPO) number. Number assigned by the WIPO to a patent application at the time it is filed (if it is filed) in the Patent Cooperation Treaty (PCT) system. WIPO publication number is the number assigned by the WIPO to an application filed through the PCT when the application becomes open to public inspection. Publication date is the date assigned by the WIPO to an application filed through the PCT when the application becomes open to public inspection. WIPO is the global forum for intellectual property services, policy, information and cooperation. It is a self-funding agency of the United Nations, with 191 member states. Their mission is to lead the development of a balanced and effective international intellectual property system that enables innovation and creativity for the benefit of all. Their mandate, governing bodies and procedures are set out in the WIPO Convention, which established WIPO in 1967.

Patent Cooperation Treaty (PCT) information. The PCT is an international patent law treaty, concluded in 1970. It provides a unified procedure for filing patent applications to protect inventions in each of its contracting states. A patent application filed under the PCT is called an international application, or PCT application. The Washington Diplomatic Conference on the Patent Cooperation Treaty was held in Washington from 25 May to 19 June 1970. The Patent Cooperation Treaty was signed on the last day of the conference on 19 June 1970. The Treaty entered into force on 24 January 1978, initially with 18 contracting states. The first international applications were filed on 1 June 1978. The Treaty was subsequently amended in 1979, and modified in 1984 and 2001. A single filing of a PCT application is made with a Receiving Office (RO) in one language. It then results in a search performed by an International Searching Authority (ISA), accompanied by a written opinion regarding the patentability of the invention, which is the subject of the application. It is optionally followed by a preliminary examination, performed by an International Preliminary Examining Authority (IPEA). Finally, the relevant national or regional authorities administer matters related to the examination of application (if provided by national law) and issuance of patent. A PCT application (also called "international patent application") has two phases. The first phase is the international phase in which patent protection is pending under a single patent application filed with the patent office of a contracting state of the PCT. The second phase is the national and regional phase which follows the international phase in which rights are continued by filing necessary documents with the patent offices of separate contracting states of the PCT. A PCT application, as such, is not an actual request that a patent be granted, and it is not converted into one unless and until it enters the "national phase". Finally, at 30 months from the filing date of the PCT application or from the earliest priority date of the application if a priority is claimed, the international phase ends and the PCT application enters in national and regional phase. However, any national law may fix time limits which expire later than 30 months. For instance, it is possible to enter the European regional phase at 31 months from the earliest priority date. National and regional phases can also be started earlier on the express request of the applicant, even before publication of the international application. If the entry into national or regional phase is not performed within the prescribed time limit, the PCT application generally ceases to have the effect of a national or regional application

Example. Canadian Alexander Graham Bell invented the telephone, a critical milestone in global communication systems. There are three patents related to the telephone. First, Patent number: 7,789, Year issued: 1877, Title: Improvements on electric telephony. Second, Patent number: 13809 Year issued: 1881, Title: Improvements in electric telephony. Third, Patent number: 13810, Year issued: 1881, Title: Improvements in electric telephony.

Patent dates

Each patent document includes the different date of the process of treatment: the priority claim (the application date), the PCT filling date, the examination date, the publishing date, the date of PCT publishing and the grant date.

The data on application date has started around 1944, but this is really near 1977 that this part of the database is quite filled. An interesting fact clearly revealed by this figure is the truncation problem. In fact, very old granted patents are not reported with their application date and very new patents are reported with their application date but are less likely to be granted. This is why we observe this increase before 1980 and the decrease after 2005. After 2005, we observed application that have been granted relatively fast. This contrasts a lot with figure (**??**) where we have an upward trend after 2005. One conclusion emerging from that fact is the time of administrative procedure in the CIPO. The gap between the date where the inventor filled for the patent (application date) and the date when the patent was granted gives an idea of the duration of the examination process of the patent office. The average time between the date of application and the granted date is around 6 years with a standard deviation of three years. Looking at table 1, this average number varies a lot depending on the sub period considered.

Between 1975 and 1985, half of the patents were granted during the four years after the application. Between 1985 and 2009, there was a little slowdown in the application process and only 10% of the patents were granted during the four years after the filling date. Although we observe an acceleration in the examination process those recent year, one must be cautious in interpreting the apparent reduction of those recent years. Indeed, the drop starting around 2005 in Figure seems to indicate that a lot of application are still in the process of grant.

Patent classification

International Patent Classification codes. Eight categories and many sub-categories. The database contains patents classified by product or by process with the International Patent Classification (IPC). The different categories are: Human Necessities, Transporting, Chemistry, textiles, construction, Engineering, Physics and Electricity. Whereas in the early 1990s most of the patents were related to chemistry and transportation, the recent years are characterized by more patents for human necessities, electricity, physics and mechanical engineering. Transportation experiences the biggest decline of the last ten years: roughly a 30% drop. The human necessities categories is dominated by the medical science, where a mayor subgroup is composed by pharmaceutical activities. Roughly 16% of the patents issued in Canada belong to the medical science in 2016.

Primary sector (agriculture, fishing, hunting) and personal and domestic articles (shoes, jewels, furniture) received each around 3% of the total patents throughout the period. 6% of the patents are used in measuring. Measuring includes length, weight, temperature, force, stress, work, mechanics, balance. It is a broader group that encompasses all domains related to the measurement of an object. The most stringent fact is the surge of artificial intelligence starting around 2008. Indeed, this increases in artificial intelligence drives the recent higher rates of physics related inventions in patent.From 2% in 2008, artificial intelligence reaches today more than 6% of the patents issued in Canada. Although medical sciences and pharmaceutical in particular, tends to give a higher value to patent (because of the high cost of invention and the relatively low cost of imitation) we surprisingly found that medical sciences are among the top lapsed rate with 13% of patents issued abandoned in the first five years after the patents is issued. Moreover, the application duration by field shows a relatively long time for Human necessities (around 92 months) and Chemistry categories (87 months). Indeed in the overall sample, human necessities took 15% more times to be issued compared to patent in physics and 30% more times to be issued compared to patent in engineering.

Inventors, assignees, applicants

The dataset contains detailed information about the inventors: full names, addresses, city, country. Most patents have multiple inventors, the average being over 2 inventors per patents. The maximum number of inventors in the database is 94, while 50% of the patents have only one inventor.

There are different type of assignees: Canadian corporations, non-Canadian corporations, Canadian individual, non-Canadian individuals, Canadian government and non-Canadian government. Company are registered in the same way as individual. The current patent database does not make any distinction between individual and companies in the section* "Assignee". The section* applicant gives the name of the firm or a person, but is only available for a small number of patents (794 705 over more than 2 millions of patents). To distinguish the name of the companies from the name of individual, we used a matching process by defining companies as assignee that do not correspond to the inventors. Further treatments help to refine the database and have only the companies involved in the patenting activities.

On the firms side, the top five patenting firms in the database includes: General Electric Company, E.I. Du Pont and Co (specialized in agriculture, materials science and speciality products), Westinghouse Electric Company, International Business MAC, Procter and Gamble (specialized in consumption goods). However, the last ten years are characterized by an increase in the number of patents of technological companies such as: Blackberry limited, Qualcomm incorporated, Schlumberger Canada limited, Haliburton energy services. Those firms describecapture the changes observed in the dynamics of patents, with the high tech industry at the top and the quite stable tendency of chemistry as Canada has a comparative advantage in natural resources. Among all those companies, only 8% are Canadian, 50% are US companies patenting in Canada, and 22% are companies from other G7 countries.

Looking at the nationality of the inventors, the patents issued in Canada are dominated by US inventors: 50% of the patents issud are from US citizens and this share is quite stable over time. The top 10 countries are mostly form the G7 with Switzerland, Netherlands and Sweden. The evolution of the share of Canadian inventors shows an upward trend whereas the number of Japanese inventors is declining. Canadian inventors represent in 2016 13% of the total number of inventors. This share was 10% in 1997. It is also interesting to notice the increase in the share of emerging markets like China and South Korea, even if those shares are quite small compared to the one of the G7 countries. There is a discrepancy between the share of patents owned by Canadian inventors (around 13%) and the share of patents owned by Canadian firms (around 9%).

Patents and provinces

There are 163, 433 patents granted with at least one Canadian inventor in the whole database. Those patents are mostly focus on Human necessities and Transporting in the 90s. But, recently Canadians tend to more specialised in patent in field such as Electricity and Physics with more than 20% of the patents in 2015.

In fact, the province of Ontario has almost the half of the whole patents with at least one Canadian inventor, followed by Quebec, Alberta and British Columbia. The share of Quebeco patents increased around 2005, accompanied by a decline in the share of Toronto patents. But the share of Quebec is steadily decreasing, by contrast to Alberta where the share is gradually going up over the years. Nevertheless, Ontario drives almost perfectly the evolution of the number of patents by Canadian residents whether the composition (dominant field) or the level.

Patent expiration

For patent applications filed on or after October 1, 1989 (patent numbers 2,000,000 and above), the maximum term of the patent is 20 years from the date the application was filed, after which time it expires. For patent applications filed before October 1, 1989 (patent numbers below 1,262,016), the maximum term of the patent is normally 17 years from the date of issue, after which time it is considered expired. In the special case of patent applications filed before October 1, 1989, (patent numbers between 1,262,017 and 1,999,999) that had not expired by July 12, 2001, the maximum term is 17 years from the date on which the patent had been issued or 20 years from the filing date, whichever occurs later. Canada does not currently grant extension of patent terms beyond the expired date.

Gender of innovators

We use machine-learning in the programming language Python, employing probabilistic Bayesian dictionaries that identify a person's gender by their first name. In the tables, "andy" indicates androgynous, "mostly" indicates probability higher than 85%, and "unassigned" means that the gender cannot be determined from the name. The overwhelming majority of innovators are males. The fraction of women has increased over time, from 1 percent in the older period to 10% in the most recent period. However, CIPO data (not shown) indicate that in Quebec the number of women inventors has been steadily declined since the mid-2000s.

Tables and Figures

Table 1: Cosine Similarity

	Mean	Standard deviations	Percentiles					Number Patents
			P50	P75	P90	P95	P99	
Cosine similarity score	0.09	0.04	0.07	0.10	0.15	0.19	0.27	15919

Distribution of cosine similarity

Table 2: Distribution (log) quality measure

Log Quality measure

	Mean	Standard deviations	Percentiles					Number Patents
			P50	P75	P90	P95	P99	
Quality (FS/BS): 0-1 year	0.33	0.15	0.33	0.40	0.52	0.60	0.81	15504
Quality (FS/BS): 0-3 years	0.37	0.24	0.34	0.5	0.71	0.83	1.08	15564
Quality (FS/BS): 0-5 years	0.41	0.32	0.34	0.59	0.88	1.03	1.36	15565
Quality (FS/BS): 0-7 years	0.43	0.38	0.31	0.65	0.99	1.20	1.60	15567

Table 3: **Relationship with patent citations.**

Log(citations)		
Log(quality)	**0.22*** **	**0.31*** **
Rsquare	0.42	0.63
Observations	19150	19150
Grant Year FE	Y	
Assignee FE	Y	
GY*A FE		Y

Regression that relates the log of (one plus) the number of patent citations to our measures of patent quality. We controlled for the granted year, the assignee, the interaction between the granted year and the assignee. There is no need to control for the technology class as we are in the drug industry. The sample covers the period 1992-2016. As patents can be assigned to multiple assignees, observations are at the patent-assignee level. Last, we cluster the standard errors by the patent grant year. The quality index and the citation index are measured on an horizon of 5 years.

Figure 1: Cosine Measure

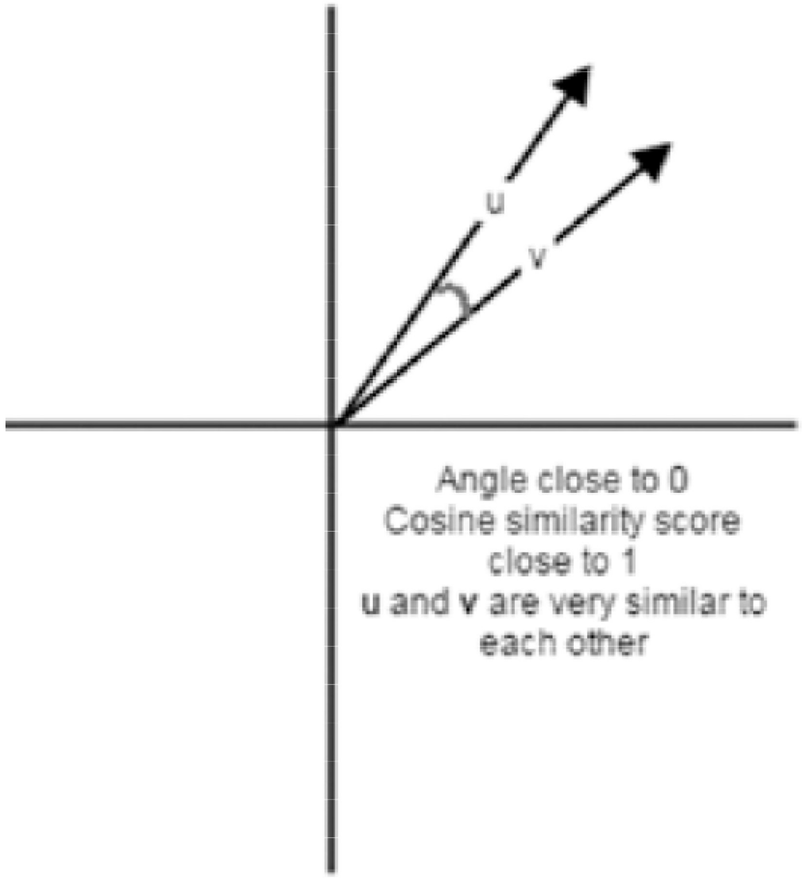

Figure 2: Cumulative density function of patent-similarity or cosine measure

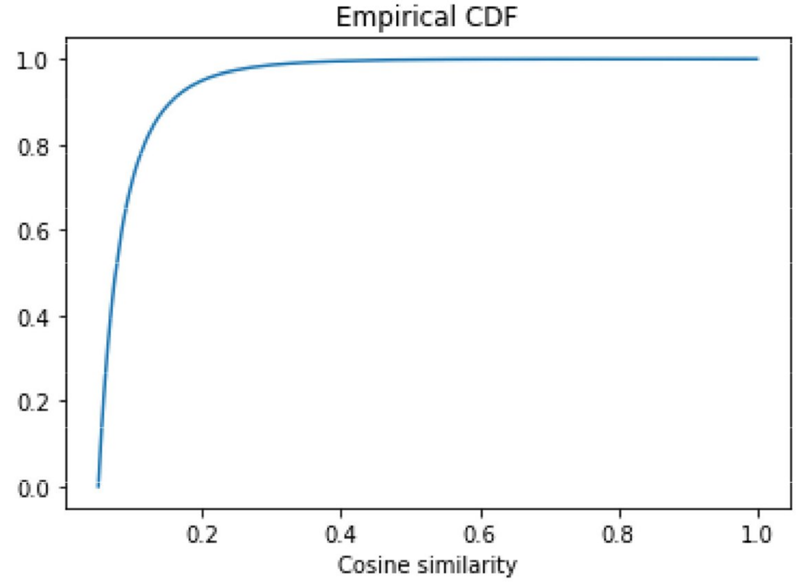

Figure 3: Cumulative density function of patent-quality measure

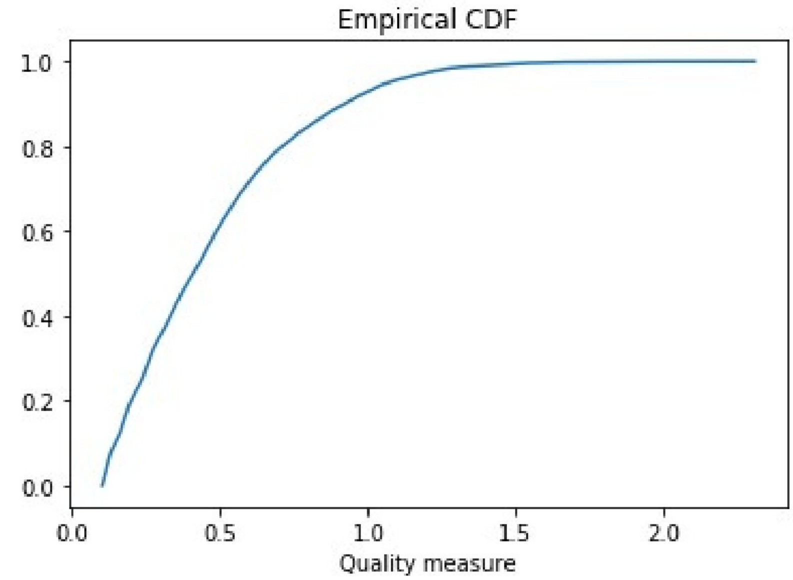

Figure 4: Probability density function of patent-quality measure

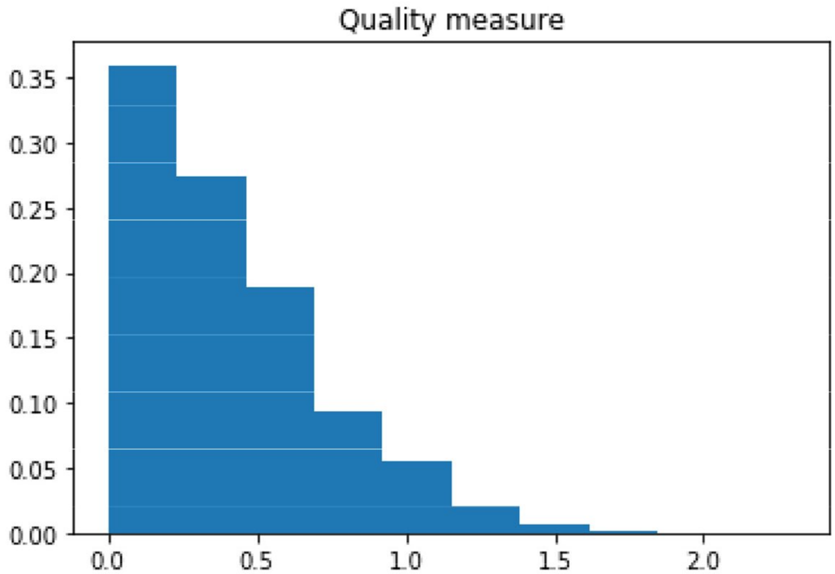

Figure 5: Distribution of patent-quality measure over time

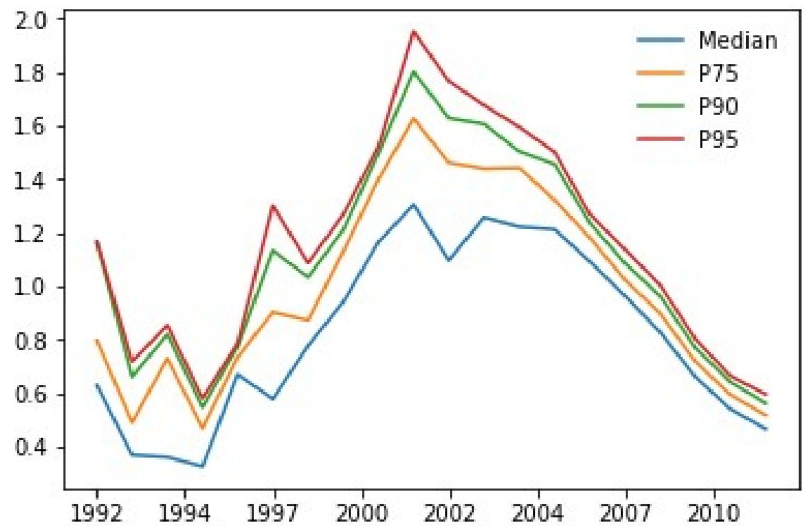

Figure 6: Relationship with patent citations

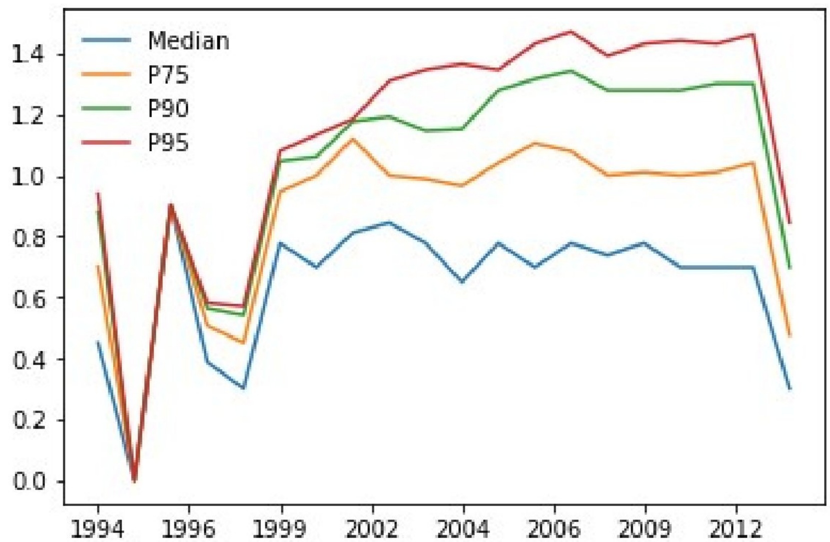

We use the Canadian patents that have a similar application in the US. We use a web scrapping to extract information about patent citations and construct the forward citation measure. To be in line with the fact that our similarity measure is conducted using the year of issuance, we will also define the forward citation as the citation after the patent has been issued (the control for this method will be done when we will use the filling year instead of the issuance year).

Appendix B

Appendix to Chapter 2

Tables

Table 1. Summary Statistics Journal Level

	Full sample	Male	Female	Mixed	Unknown
	(1)	(2)	(3)	(4)	(5)
			Overall		
Total Number	24033	17,785	1,091	4,255	902
			Journal		
American Economic Review	1931	1,398	108	378	47
Econometrica	1485	1,244	39	186	16
Journal of Econometrics	3089	2,230	96	532	231
Journal of economic growth	260	193	10	47	10
Journal of economic literature	340	266	26	41	7
Journal of economic perspectives	1026	809	60	144	13
Journal of Economic Theory	2651	2,138	70	312	131
Journal of Finance	2001	1,509	60	395	37
Journal of Financial Economics	2247	1,533	87	517	110
Journal of International Economics	1775	1,146	167	371	91
Journal of Labour Economics	836	513	87	193	43
Journal of Monetary Economics	1193	881	74	198	40
Journal of Political Economics	1122	889	49	163	21
Quaterly Journal of Economics	1105	830	49	213	13
Review of Economics Studies	1201	936	56	183	26
Review of Financial Economics	1771	1,270	53	382	66

This table presents the journals in the database and the total number of papers per journal over the sample period 1991-2019. The selected papers exclude proceedings, comments, articles of less than three pages, books reviews, bibliographical items, articles without references and without abstracts, editorial material, letters and corrections.

Table 2. Summary Statistics Publication Level

	Full sample	Male	Female	Mixed	Unknown
	(1)	(2)	(3)	(4)	(5)
Total	24033	17,785	1,091	4,255	902
			Authors		
Single authored	6,949	5,950	836	0	163
Coauthored	17,084	11,835	255	4,255	739
			Field		
Mathematical	4,259	3,217	128	664	250
Microeconomics	4,051	3,243	123	536	149
Macroeconomics	1,938	1,514	79	305	40
International Economics	2,235	1,544	191	423	77
Finance	6,184	4,467	219	1,284	214
Labour & Education	2,514	1,645	211	577	81
IO	236	182	16	28	10
Other	2,616	1,973	124	438	81
			Institutions		
Very Top tier	7,203	5,406	267	1,387	143
Middle tier	10,283	7,392	419	2,067	405
Low tier	4,486	3,242	252	709	283
Undefined	2,061	1,745	153	92	71
			References		
Average Number	38.04	37.31	38.87	41.29	36.17
Average Number (database)	9.25	8.83	8.72	11.18	9.08

The table describes the papers published per journal in the database over the period 1991-2019. The field selection is based on the Journal of Economic Literature (JEL) codes. The category *Other* includes public economics, agricultural economics, general economics, urban economics, law and economics, business administration, economic history, and economics systems.

Table 3. Distribution of the cosine

Mean	0.33
Standard deviation	0.14
1th percentile	0.13
5th percentile	0.16
10th percentile	0.19
25th percentile	0.23
Median	0.3
75th percentile	0.39
90th percentile	0.51
95th percentile	0.61
99th percentile	0.86

The table shows the distribution of the relative cosine. For paper P and P_{max} such as: $P_{max} = argmax_{P' \in C} cos(P,P')$, for any given article P' in the database, the relative cosine of paper P and paper P' is defined as: $\tilde{\lambda}_{p,p'} = \frac{\lambda_{p,p'}}{\lambda_{p,p_{max}}}$.

Table 4. Number of Citations of the chosen paper and the omitted papers

	Mean	Standard deviations	Median	75th	90th	95th
Omitted among the relevant prior literature	0.85	0.84	0.77	0.98	1.14	1.23
Cited among the relevant prior literature	1.20	1.29	1.03	1.29	1.51	1.66
Overall paper cited in the database	1.30	1.26	1.19	1.41	1.60	1.72

The table shows the distribution of average number of citations of the papers omitted among the relevant prior literature (second row), the papers cited among the relevant prior literature (third row) and the papers from the database that an article chooses to cite (last row). Citations are adjusted by a monotonic transformation (logarithm of one plus citation). The citations are constructed on various horizons using the database in the paper. Basically, the citations considered here are the average citations by the time a given paper is published. Most recent citation update was in January 2019.

Table 5. Innovativeness (quality) index distribution

Quality	Quality, 0-5 years
Mean	0.82
Standard deviation	0.26
1th percentile	0.37
5th percentile	0.49
10th percentile	0.55
25th percentile	0.65
Median	0.77
75th percentile	0.93
90th percentile	1.12
95th percentile	1.29
99th percentile	1.73
p90/p50	1.44
p50/p10	1.41

This table shows the distribution of the Innovativeness index defined in the main text. Three horizons are considered: 0 to 1 year, 0 to 3 years and 0 to 5 years.

Table 6. Quality index, citations and gender: 5-year horizon

	Outcome Variable: Forward Citations					
	(1)	(2)	(3)	(4)	(5)	(6)
log(1 + q)	0.586***	0.585***	0.586***	0.578***		
	(0.036)	(0.036)	(0.037)	(0.039)		
male		0.303***				
		(0.030)				
female		0.242***				
		(0.038)				
mixed		0.272***				
		(0.033)				
$A1f$			-0.038**		-0.037**	
			(0.016)		(0.015)	
Only f				-0.057**		-0.047*
				(0.027)		(0.027)
quality					0.266***	0.271***
					(0.018)	(0.018)
$A1f\cdot$ quality					0.008	
					(0.028)	
(Only f) $\cdot$ quality						-0.081*
						(0.047)
N	17,183	17,173	16,623	13,903	16,623	13,903
R-sqr	0.31	0.33	0.33	0.32	0.33	0.32

This table shows the relationship between the quality index, q, the number of citations, and gender. The controls include dummies for journal, field, affiliation, year of publication, number of authors, and NBER membership. See equation 2.7.1 in text for the construction of the quality index q. The variable *quality* is the standardized q. Standard errors are clustered by journal of publication and year and are reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

Table 7. Quality, citations and gender: Journal and affiliation

	Outcome Variable: Forward Citation 0-5 years						
	Journal				Institutions		
	Top 5		non Top 5		Top tier	Mid tier	Low tier
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$A1f$	-0.033		-0.038**		-0.015	-0.050**	-0.009
	(0.032)		(0.018)		(0.027)	(0.021)	(0.033)
quality'	0.239***	0.243***	0.287***	0.283***	0.336***	0.266***	0.202***
	(0.030)	(0.030)	(0.022)	(0.023)	(0.032)	(0.027)	(0.028)
$A1f$·quality'	-0.060		0.002		-0.058	0.018	0.073
	(0.052)		(0.033)		(0.054)	(0.041)	(0.065)
f		-0.146***		-0.016			
		(0.055)		(0.030)			
f·quality'		-0.079		-0.078			
		(0.095)		(0.054)			
N	4748	4060	11875	9843	5410	7569	3119
R-sqr	0.213	0.207	0.346	0.345	0.279	0.255	0.206

This table shows the relationship between the quality index (q-index), the number of citations and the gender of the paper by type of journals and by type of affiliations. The controls include dummies for journals, field, institutions, year of publications, number of authors, NBER member. The q-index is built following equation 2.7.1. *quality'* is the standardized quality index. Standard errors are clustered by journal of publication and years reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

Figures

Fig. 1. Life Cycle of Citations

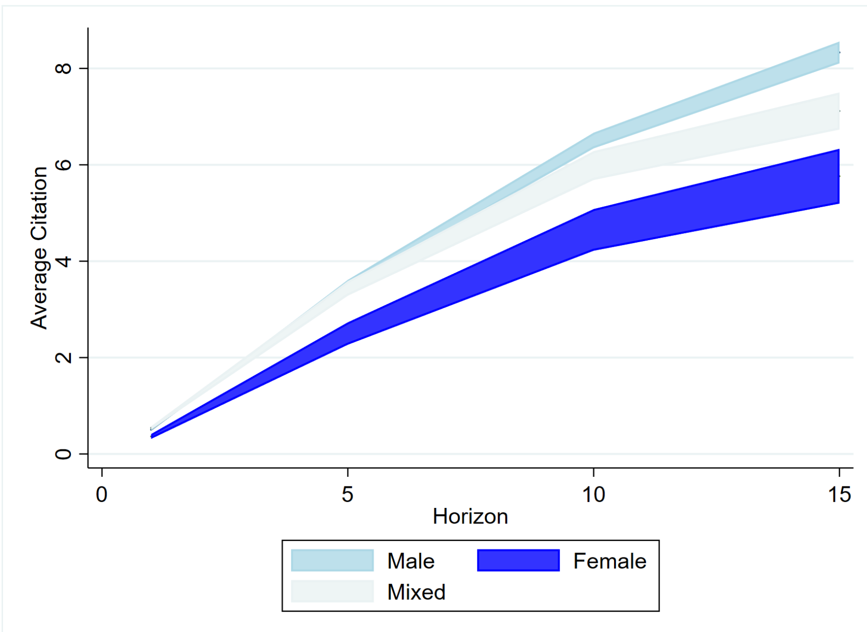

The figure shows the evolution of the average number of citation for different years after the publication of the paper. The plots distinguish papers written by men only, from papers written by women only, from papers written by mixed gender teams. Each band width represents the 5% confidence intervals.

Fig. 2. Field and references

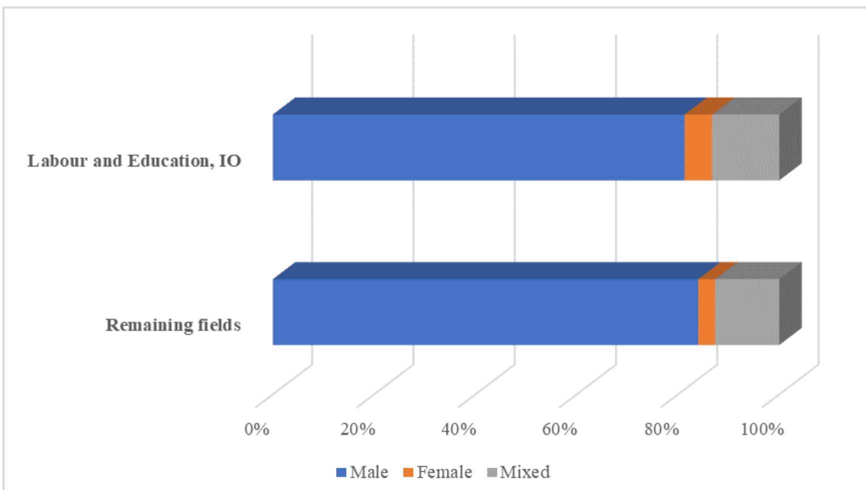

This figures presents the share of references that are attributed to a given gender by field. Basically, the plot answers to the question: what is the average fraction of citations that refers to papers written by men, women or mixed gender? *Male* designed paper written by only men; *female* designed paper written by only women; *Mixed* designed paper written by a team of females and males.

Fig. 3. Aggregate overview of the omission propensity

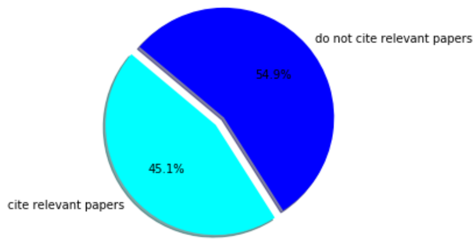

(a) $\mathcal{P}_p$ with $n = 5$

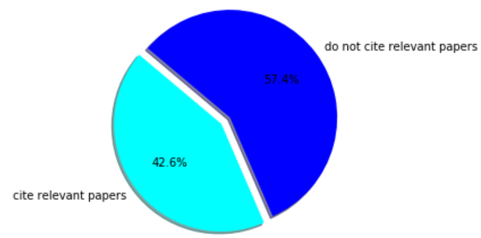

(b) $\mathcal{P}_p$ with $n = 5$ constrained

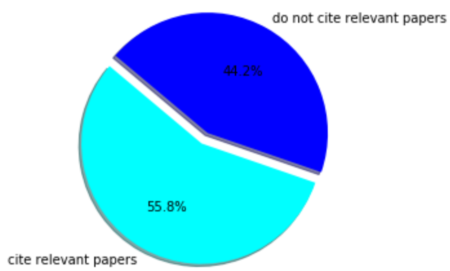

(c) $\mathcal{P}_p$ with $n = 10$

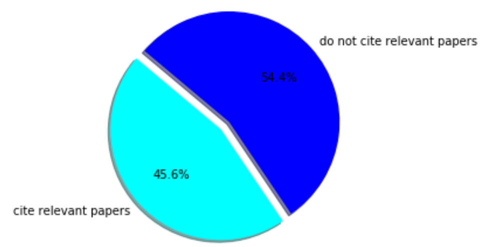

(d) $\mathcal{P}_p$ with $n = 10$ constrained

The figure gives an aggregate overview of the propensity to omit relevant prior literature $\mathcal{P}_p$. A paper will be said to omit its relevant prior literature if it does not cite at least one paper from its relevant prior literature. At the opposite, a paper will be said to cite its relevant prior literature if it cites at least one paper from its relevant prior literature. Pie chart (a) shows the propensity of omission when $\mathcal{P}_p$ has 5 elements. Pie chart (b) shows the propensity of omission when $\mathcal{P}_p$ has 5 elements. For paper P and P_{max} such as: $P_{max} = argmax_{P' \in C} cos(P,P')$, for any given article P' in the database, the relative cosine of paper P and paper P' is defined as: $\tilde{\lambda}_{p,p'} = \frac{\lambda_{p,p'}}{\lambda_{p,p_{max}}}$. In the constrained specification, the relative cosine should be greater than 0.5. Similarly, Pie chart (c) shows the propensity of omission when $\mathcal{P}_p$ has 10 elements. Pie chart (d) shows the propensity of omission when $\mathcal{P}_p$ has 5 elements, but the relative cosine should be greater than 0.5.

Fig. 4. Odds of omission by field

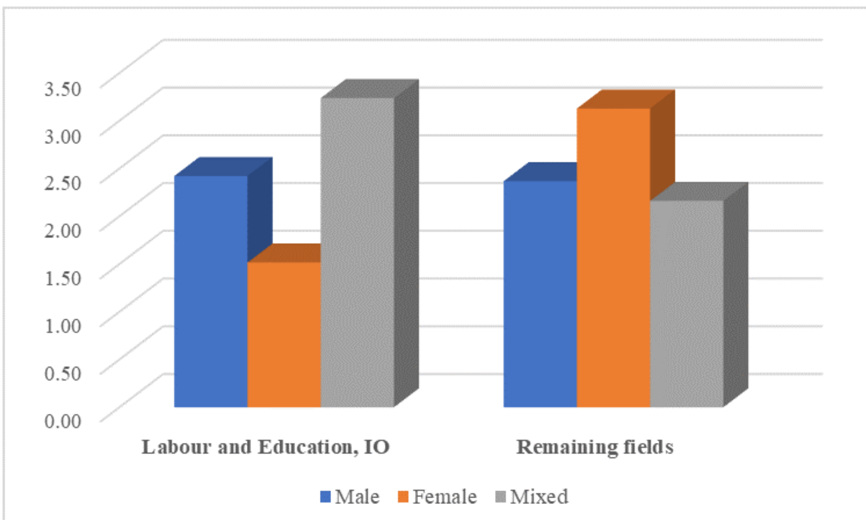

Panel (a) decomposes the likelihood of being omitted compared to the one of being cited with respect to the gender and the field. *Male* designed paper written by only men; *female* designed paper written by only women; *Mixed* designed paper written by a team of females and males.

Fig. 5. Experiment 2: Actual distribution-Target distribution (1/2)

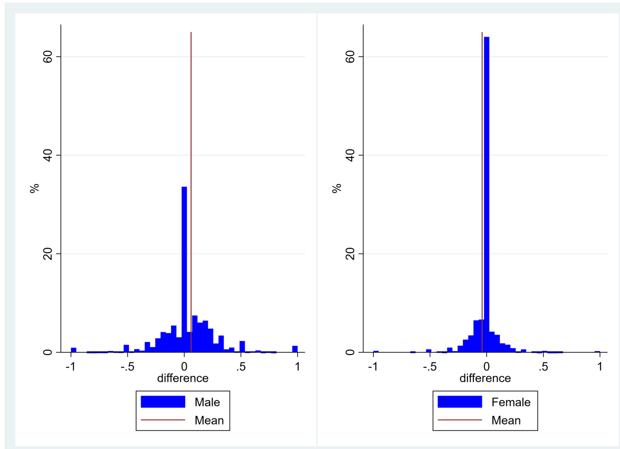

(a) Males versus Females

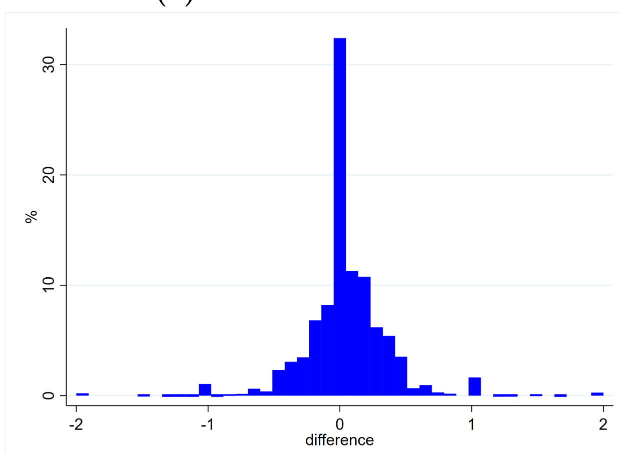

(b) Difference Males-Difference Females

This figure plots the difference between the actual distribution and the target distribution of gender type g. For each paper p, the actual distribution of gender is the share of papers in its references belonging to each category of gender (only males, only females, mixed gender). The target distribution of gender is the share of the closest papers (in the sense of the relative cosine) in its prior literature belonging to each category of gender (only males, only females, mixed gender). For each paper, the difference between the actual distribution of gender g and the target distribution is taken. Panel (a) plots the distribution of this difference for males versus females. A positive difference means that the actual distribution of a certain gender type is higher than the target distribution of this gender type. A negative difference means that the actual distribution of a certain gender type is lower than the target distribution of this gender type. Finally, a null difference means that the actual distribution of a certain gender type is the same as the target distribution of this gender type. Panel (b) takes the difference of the difference for males and for females. A positive double difference means that males are more *over-cited* or less *under-cited* than females.

Fig. 6. Distribution of innovativeness (quality) index

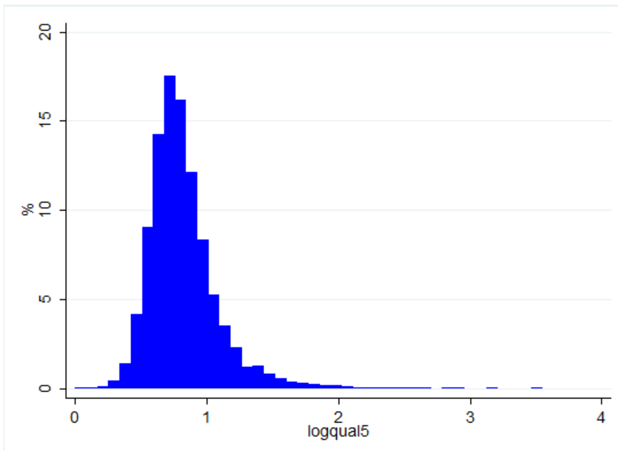

(a) Overall distribution

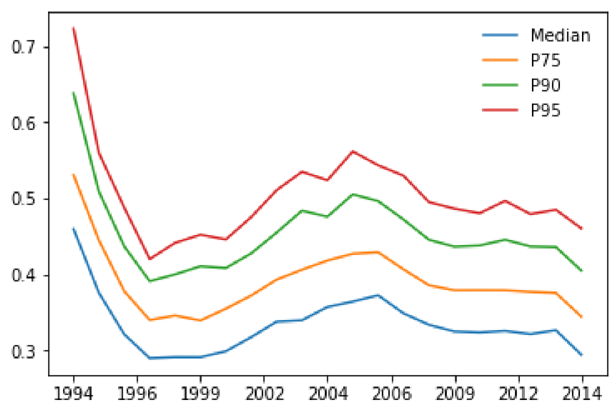

(b) Distribution over time

Panel (a) shows the overall distribution of the innovativeness index (q-index). Panel (b) shows the distribution of this index over time. The index is built following equation 2.7.1.

Fig. 7. Quality and citations

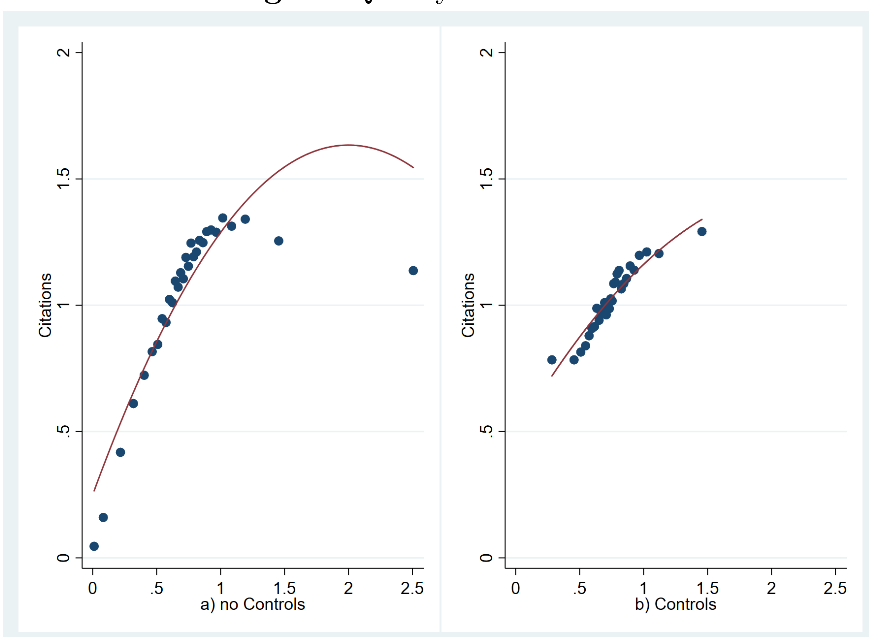

The figure plots the link between the number of forward citations after and the quality index over 0-5 years. The binned scatter plot b) controls for journals, field, institutions, year of publications, number of authors, NBER member. The corresponding estimates are in table 9.

Fig. 8. Quality, Citation and Gender

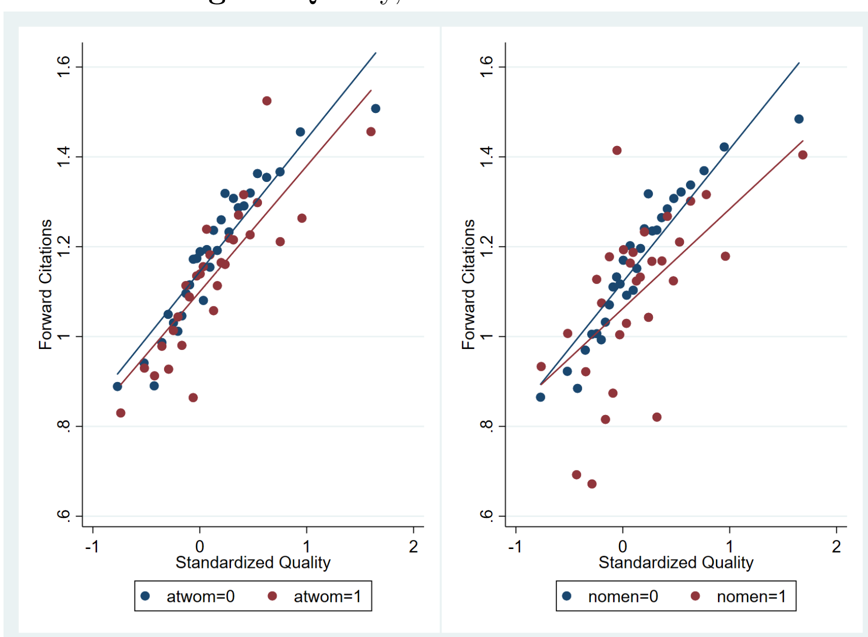

The figure plots the link between the number of forward citations and the q-index for papers written by males and papers written by at least one female (left) and for males and females (right). The variable *atwom* indicates a dummy variable that takes one if the paper is written by at least one female and 0 otherwise. The variable *nomen* indicates a dummy variable that takes one if the paper is written by only females and 0 otherwise. The binned scatter plot controls for journals, field, institutions, year of publications, number of authors, NBER member. The corresponding estimates are in table 6.

Data and Variables

Data

Data cleaning. Several steps were used to process the textual information: remove the stopwords, drop punctuations, lower letter, lemmatize, tokenize. Stopwords are words commonly used such as “a",“the"... A list of stopwords is built using information available online and internal to the software.[1] Lemmatazing or stemmatizing means put the words at their root. For example, “taxing" becomes “taxe". Finally, by tokenizing, the paragraphs are split into sentences and the sentences are split into words.

Institution Ranking. Information about author institutions were retrieved using the web of science database. When the data is available, each author could have one or many affiliations. All the affiliations available were extracted. Thereafter, the number of papers in the database that refer to a particular affiliation can be found. To avoid changes in university rankings, a more conservative method as in Engel (2019) was used. The institutions are ranked according to the number of papers they have in the database.

Primary Field. Articles are classified based on the Journal of Economic Literature Code (JEL) codes. As mentioned above the database of web of science does not include jel code. The one of Ideas Repec include the JEL codes of certain articles. Moreover, the articles usually have several JEL codes. The classification is done on the articles of the 16 newspapers selected over the period 1991-2019. Areas for the classification include: Microeconomics, macroeconomics, public finance, labor, industrial organization, development, urban economics, environmental, econometrics, finance, international, experimental (lab), economic history, economic economy, productivity, law and economics, and other. To assign each article a primary, the following machine learning algorithm is considered. The dataset is split in two: a training dataset and a using dataset. With the training dataset, the algorithm can recognize the characteristics of the different categories. This dataset is composed of articles that have a single primary JEL code. For articles in newspapers whose the field is widely admitted (ex Journal of Labor economics), the field of the journal is assigned (see Angrist et al. (2019)). To avoid a high level of success due to an over-representation of a class, the sample is chosen to ensure proportional share of each field. The idea is to predict the field using the characteristics of the article. The dependent variables will therefore be the titles and the keywords.[2]

The training database was used to train a Support Vector Classifier (SVC) using the titles and keywords to predict the JEL code. The package used is the “Scikit-learn" package

[1] See for example $https://www.nltk.org/nltk_data/$.

[2] One can also use the abstracts. The results remain the same.

(Pedregosa et al., 2011).[3] Titles and keywords were subject to a cleaning procedure (remove stopwords, consider ngrams, drop punctuations, lemmatize ...) and transform into digital data by the TF-IDF. We used a grid search to get the optimal hyper-parameter values for the classification. Thus, in a subset with 90-10 training-test sample, the accuracy of the algorithm is estimated to approximately 90%.

Alternative Omission score definition

Another way to construct the omission score is to used the value of the similarity. In fact, the metric used to assess whether an article should be quoted by another is the similarity between the two articles. Indeed, the semantic similarity makes it possible to know if two articles tackle the same questions, the same subject or the same problematic. The intensity of the similarity is then evaluated by the relative distribution of words from one article to another. To establish the omission score, we look at if an article mentions the most similar items attached to it. One way to measure this is to use the backward similarities and evaluate the score by weighting the citation of an article by the value of the similarity between the two articles.

More specifically, we define the backward similarity index as the sum of pairwise cosine similarities of paper p, published in t, with papers published in $t-T$:

$$BS^0_{-T}(p) = \Sigma_{p'} cos(p,p')$$

Similarly, the forward similarity is the sum of pairwise cosine similarities of paper p, published in t, with papers published in $t+T$:

$$FS^T_0(p) = \Sigma_{p'} cos(p,p')$$

Moreover, let BNc: dummy vector that takes 1 if paper P does not cite backward paper P' and $Vcos_p$ the vector of cosine similarities between P and the other papers. The omission score will be:

$$Omis_p = \frac{BNc.Vcos_p}{BS^0_{-\infty}}$$

The numerator is the dot product of the vector represented the paper not cited by paper P and the one represented the cosine between P and the other papers. The aim is to capture the sum of the cosine similarity of the papers similar to P but that P does not cited. The denominator capture the sum of the cosine similarity of paper prior to P. We could rewrite the omission score formula as followed:

$$Omis_p = \frac{\Sigma_{p' \in \mathscr{B}_{nc}} cos(p,p')}{\Sigma_{p' \in \mathscr{B}_{nc}} cos(p,p') + \Sigma_{p' \in \mathscr{B}_{c}} cos(p,p')}$$

[3] The SVC classifier was the best classifier among a set of tests made with other classifiers like Random Forest, decision trees ...

$\mathscr{B}_{nc}$ is the set of papers prior to P that have not being cited by P whereas $\mathscr{B}_c$ is the set of papers prior to P that have being cited by P. In other words, the omission score is the relative share of papers prior to P that have a non zero cosine with P and are not cited by P over the whole set of papers prior to P with non zero cosine (cited or not by P). A high omission score means that paper P does not cite relevant papers on which it is based and a low omission score means that paper P mostly cite relevant prior literature to him. To avoid noisy or misleading index, we will set some thresholds for the cosine varying from 0.10 to 0.2 to have a reasonable and meaningful set of papers in the $\mathscr{B}_{nc}$ and $\mathscr{B}_c$.

Additional tables and graphs

Table 8. Paper Citations

	Full sample	Male	Female	Mixed	Unknown
	(1)	(2)	(3)	(4)	(5)
	Citation over horizons				
Citations, 0-1 years	0.27	0.28	0.21	0.29	0.11
Citations, 0-2 years	0.53	0.54	0.43	0.56	0.25
Citations, 0-3 years	0.74	0.76	0.62	0.77	0.37
Citations, 0-4 years	0.90	0.92	0.75	0.93	0.46
Citations, 0-5 years	1.02	1.05	0.86	1.04	0.53
Citations, 0-6 years	1.12	1.15	0.96	1.13	0.58
Citations, 0-7 years	1.19	1.23	1.04	1.20	0.62
Citations, 0-8 years	1.26	1.30	1.10	1.26	0.65
Citations, 0-9 years	1.30	1.35	1.15	1.30	0.68
Citations, 0-10 years	1.34	1.39	1.18	1.33	0.70
Citations, 0-11 years	1.38	1.43	1.21	1.35	0.73
Citations, 0-12 years	1.40	1.46	1.24	1.37	0.74
Citations, 0-13 years	1.43	1.48	1.26	1.39	0.75
Citations, 0-14 years	1.44	1.50	1.28	1.40	0.76
Citations, 0-15 years	1.46	1.52	1.29	1.41	0.77
Citations, 0-16 years	1.47	1.53	1.30	1.42	0.78
Citations, 0-17 years	1.48	1.54	1.31	1.42	0.79
Citations, 0-18 years	1.49	1.55	1.32	1.43	0.80
Citations, 0-19 years	1.50	1.56	1.32	1.43	0.80
Citations, 0-20 years	1.50	1.56	1.33	1.43	0.80
	Overall database				
Mean	1.56	1.62	1.39	1.49	0.84
Standard deviation	1.18	1.19	1.09	1.14	0.91
Median	1.38	1.61	1.38	1.38	0.69
75th percentile	2.39	2.48	2.20	2.30	1.38
90th percentile	3.17	3.26	2.94	3.04	2.19
95th percentile	3.61	3.69	3.33	3.49	2.63
	WoS database				
Mean	3.26	3.32	3.14	3.18	2.53
Standard deviation	1.57	1.56	1.54	1.59	1.47
Median	3.33	3.40	3.29	3.29	2.60
75th percentile	4.35	4.41	4.30	4.31	3.58
90th percentile	5.25	5.30	5.08	5.18	4.34
95th percentile	5.73	5.78	5.44	5.63	4.76
	Correlation				
correlation: database-wos	0.81	0.81	0.77	0.81	0.71
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

The table shows the mean value of the number of citations (logarithm of one plus number of citation). The citations are constructed on various horizons using the database in the paper. The last line provides the correlation between the overall citation in the database used in this paper and the citation obtained from Web of Science. The p-value are in parenthesis. Citations were downloaded on January 2019.

Table 9. Relationship between omission and gender

	Outcome variable: Omission									
	Baseline gender					Gender with historiy of top 5				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
female	0.193***	0.291***	0.194***	0.193***	0.172***	0.068*	0.196***	0.129***	0.132***	0.112***
	(0.043)	(0.045)	(0.045)	(0.046)	(0.046)	(0.038)	(0.039)	(0.040)	(0.041)	(0.041)
mixed	-0.215***	-0.072***	-0.010	-0.003	0.008	0.079**	0.189***	0.161***	0.143***	0.125***
	(0.024)	(0.026)	(0.027)	(0.027)	(0.028)	(0.033)	(0.034)	(0.035)	(0.035)	(0.036)
unknown	0.785***	0.706***	0.576***	0.532***	0.514***	0.816***	0.725***	0.586***	0.541***	0.521***
	(0.058)	(0.059)	(0.059)	(0.060)	(0.061)	(0.058)	(0.059)	(0.059)	(0.060)	(0.061)
top5 j		-0.500***	-0.449***	-0.556***	-0.654***		-0.494***	-0.446***	-0.553***	-0.652***
		(0.018)	(0.019)	(0.019)	(0.020)		(0.018)	(0.019)	(0.019)	(0.020)
same primary field		-0.972***	-0.965***	-0.922***	-0.859***		-0.971***	-0.964***	-0.921***	-0.859***
		(0.019)	(0.019)	(0.019)	(0.020)		(0.019)	(0.019)	(0.019)	(0.020)
Difference of publication years		0.028***	0.022***	0.042***	0.057***		0.028***	0.023***	0.042***	0.057***
		(0.001)	(0.001)	(0.002)	(0.002)		(0.001)	(0.001)	(0.002)	(0.002)
Gender structure relevant prior literature		-0.606***	-0.477***	-0.229***	-0.209***		-0.733***	-0.544***	-0.284***	-0.249***
		(0.058)	(0.058)	(0.059)	(0.059)		(0.056)	(0.056)	(0.057)	(0.057)
Number of references				-0.048***	-0.042***				-0.048***	-0.042***
				(0.001)	(0.002)				(0.001)	(0.002)
Number Authors					-0.013					-0.014
					(0.011)					(0.011)
Institution of j FE			Y	Y	Y			Y	Y	Y
Institution of i FE					Y					Y
Journal of i FE					Y					Y
Year of publication of i FE					Y					Y
Field FE					Y					Y
N	117694	117694	117694	117694	117661	117694	117694	117694	117694	117661
R-sqr	0.004	0.046	0.065	0.089	0.102	0.003	0.046	0.065	0.089	0.102

This table shows the relationship between the omission and the gender of the omitted paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. *female* represents papers written by only women. *mixed* represents papers written by a team of females and males. *unknown* represents papers for which we cannot recover the gender of at least one of the author. The reference variable is *male*, which represents papers written by only men. *Top 5 j* is binary and indicates if paper j is published in a top 5 journal or not. *Same primary field* is binary and indicates if paper i and paper j have the same primary field. *Difference of publication years* is the difference between the publication year of paper i and the publication year of paper j. *Gender structure relevant prior literature* is the share of paper written by at least one female author in the relevant prior literature. *Number of references* is the number of references recovered from the database. *Number Authors* is the number of authors writing the paper. The equations are estimated using a logit model, distinguishing the case where the gender is attributed as in the baseline and the case where the gender for mixed teams is the gender of the authors with the highest number of top 5. *Mixed* in that case are teams with a tie (same number of top 5 publications for both gender). The odd ratio for a variable is the exponential of its given coefficient. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$* * * = p < 0.01$

Table 10. Relationship between omission and gender: effect of the gender of the citing

	Outcome variable: Omission										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
female j	0.017***					0.016***	0.036***			0.029***	0.036***
	(0.005)					(0.005)	(0.005)			(0.005)	(0.005)
At least one female in j		0.0055*		0.005*	0.021***			0.017***	0.021***		
		(0.002)		(0.003)	(0.003)			(0.003)	(0.004)		
Same gender			-0.017***								
			(0.002)								
female i								-0.018***	-0.004	-0.016**	-0.006
								(0.006)	(0.006)	(0.006)	(0.006)
At least one female in i				-0.007**	0.005*	0.000	0.005				
				(0.003)	(0.003)	(0.003)	(0.003)				
(female j)*(at least one female i)							-0.073***				
							(0.012)				
(at least one female j)*(at least one female i)					-0.063***						
					(0.007)						
(female j)*(female i)											-0.106***
											(0.022)
(at least one female j)*(female i)									-0.058***		
									(0.013)		
N	98417	113247	117661	109055	109055	94760	94760	88725	88725	77832	77832
R-sqr	0.08	0.10	0.10	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the gender of the citing paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. $female_x$ represents paper x written by only women. *at least one* $female_x$ represents paper x with at least one female author. All the specifications include controls for paper j published in a top 5 journal; paper i and paper j having the same primary field; difference between the publication year of paper i and the publication year of paper j; the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. The equations are estimated using a logit model. The table displayed the marginal probabilities. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$* * * = p < 0.01$

Table 11. Relationship between omission and gender: heterogeneous effect and robustness

	Outcome variable: Omission										
	All						Top 5 j		Institution j		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
							Top5	non Top5	Top tier	Mid tier	Low tier
at least one female j	0.019***	0.019***				0.006					
	(0.003)	(0.003)				(0.008)					
cosine	-0.767***	-.758***	-0.758***	-0.-.757***	-0.762***	-0.78***					
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)					
at least one female i	0.007**		0.006**								
	(0.003)		(0.003)								
(at least one female j)* (at least one female i)	-0.060***										
	(0.007)										
female i		-0.002		-0.004			-0.009	0.001	-0.030**	0.002	0.007
		(0.006)		(0.006)			(0.01)	(0.006)	(0.013)	(0.01)	(0.011)
(at least one female j)* (female i)		-0.052***									
		(0.013)									
female j			0.036***	0.035***	-0.018		0.050***	0.032***	0.038**	0.061***	0.019
			(0.005)	(0.005)	(0.013)		(0.010)	(0.006)	(0.016)	(0.009)	(0.012)
(female j)* (at least one female i)			-0.069***								
			(0.012)								
(female j)* (female i)				-0.096***			-0.093**	-0.098***	-0.102*	-0.171***	-0.061
				(0.022)			(0.039)	(0.025)	(0.056)	(0.043)	(0.05)
cosine * female j					0.134**						
					(0.06)						
cosine * (at least one female in j)						-0.04					
						(0.03)					
N	109055	88725	94760	77832	113280	88725	27735	50097	18959	25461	11497
R-sqr	0.108	0.107	0.105	0.107	0.11	0.11	0.13	0.10	0.12	0.09	0.09

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the gender of the citing paper. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. $female_x$ represents paper x written by only women. *at least one* $female_x$ represents paper x with at least one female author. *Cosine* is the value of the cosine between i and j. All the specifications include controls for paper j published in a top 5 journal; paper i and paper j having the same primary field; difference between the publication year of paper i and the publication year of paper j; the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. The equations are estimated using a logit model. The table displayed the marginal probabilities. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

Table 12. Relationship between omission and gender: Peer effects

	Outcome variable: Omission			
	(1)	(2)	(3)	(4)
at least one female j	0.002	0.003		
	(0.003)	(0.003)		
same affiliation	-0.101***	-0.100***	-0.099***	-0.101***
	(0.004)	(0.005)	(0.005)	(0.005)
at least one female i	-0.005*	-0.005*	0.002	0.002
	(0.003)	(0.003)	(0.003)	(0.003)
(at least one female j) * (same affiliation)		-0.006		
		(0.011)		
female j			-0.003	-0.006
			(0.005)	(0.005)
(female j) * (same affiliation)				0.031**
				(0.015)
N	109055	109055	94760	94760
R-sqr	0.167	0.167	0.165	0.165

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the effect of being in the same affiliation. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. The control variables include the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

Table 13. Relationship between omission and gender: Existing recognition

	Outcome variable: Omission					
	(1)	(2)	(3)	(4)	(5)	(6)
at least female j	0.019***	0.014***	0.012***			
	(0.003)	(0.003)	(0.003)			
at least female i	0.007**	0.008***		0.006**	0.008**	
	(0.003)	(0.003)		(0.003)	(0.003)	
(at least female j)*(at least female i)	-0.060***	-0.054***				
	(0.007)	(0.006)				
citation j		-0.090***	-0.088***		-0.089***	-0.088***
		(0.001)	(0.001)		(0.001)	(0.001)
(at least female j)*(citation j)			-0.008***			
			(0.002)			
female j				0.036***	0.017***	0.008*
				(0.005)	(0.005)	(0.005)
(female j)*(at least female i)				-0.069***	-0.060***	
				(0.012)	(0.011)	
(female j)*(citation j)						-0.006
						(0.004)
N	109055	109055	113280	94760	94760	98417
R-sqr	0.108	0.182	0.182	0.105	0.180	0.180

	(1)	(2)	(3)	(4)	(5)	(6)
at least female j	0.019***	0.018***	0.009***			
	(0.003)	(0.003)	(0.003)			
at least female i	0.007**	0.007**		0.006**	0.006**	
	(0.003)	(0.003)		(0.003)	(0.003)	
(at least female j)*(at least female i)	-0.060***	-0.060***				
	(0.007)	(0.007)				
Max T5 j		-0.021***	-0.019***		-0.019***	-0.019***
		(0.002)	(0.002)		(0.002)	(0.002)
(at least female j)*(Max T5 j)			-0.013***			
			(0.004)			
female j				0.036***	0.030***	0.017***
				(0.005)	(0.005)	(0.005)
(female j)*(at least female i)				-0.069***	-0.070***	
				(0.012)	(0.012)	
(female j)*(Max T5 j)						-0.030**
						(0.013)
N	109055	109055	113280	94760	94760	98417
R-sqr	0.108	0.109	0.109	0.105	0.106	0.107

This table shows the relationship between the omission and the gender of the omitted paper emphasizing the existing number of citation or the number of top 5 of the most prolific author of paper j. The dependent variable, omission, is binary and indicates whether a paper i cites a paper j in the database given that j is in the relevant prior literature of i. The relevant prior literature is defined by equation 2.4.5. The control variables include the share of paper written by at least one female author in the relevant prior literature; the number of references recovered from the database; the number of authors writing the paper; field fixed effect, journal fixed effects, year fixed effect, institutions fixed effect. Standard errors are clustered by papers and reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$

Table 14. Relationship between omission index and gender

	Outcome variable: Omission					
	Raw Count		Intensity		Compensation	
	(1)	(2)	(3)	(4)	(5)	(6)
female	0.054**		0.180*		0.638**	
	(0.03)		(0.11)		(0.26)	
mixed	0.026		0.104**		0.446***	
	(0.02)		(0.05)		(0.17)	
unknown	0.046		0.248*		1.682***	
	(0.03)		(0.13)		(0.25)	
N	11162		11162		11162	
R-sqr	0.072		0.061		0.172	
	Journal					
	Top 5					
at least one female	0.069***		0.174***		0.534*	
	(0.02)		(0.07)		(0.32)	
female		0.124**		0.299*		1.103**
		(0.05)		(0.18)		(0.51)
N	3600	3095	3600	3095	3600	3095
R-sqr	0.151	0.158	0.130	0.134	0.134	0.201
	Non Top 5					
at least one female	0.016		0.082		0.443***	
	(0.02)		(0.06)		(0.14)	
female		0.022		0.089		0.387
		(0.03)		(0.12)		(0.24)
N	8592	7136	8592	7136	8592	7136
R-sqr	0.086	0.090	0.069	0.069	0.226	0.223
	Institution					
	Top Tier					
at least one female	0.028		0.060		0.840***	
	(0.02)		(0.07)		(0.31)	
female		0.035		-0.126		0.786
		(0.05)		(0.16)		(0.55)
N	3661	3035	3661	3035	3661	3035
R-sqr	0.084	0.085	0.066	0.066	0.168	0.160
	Mid Tier					
at least one female	0.028		0.158**		0.348**	
	(0.02)		(0.07)		(0.17)	
female		0.045		0.302*		0.320
		(0.04)		(0.18)		(0.36)
N	5125	4174	5125	4174	5125	4174
R-sqr	0.069	0.068	0.062	0.057	0.184	0.175
	Low Tier					
at least one female	0.043		0.120		-0.017	
	(0.03)		(0.11)		(0.23)	
female		0.088		0.223		0.360
		(0.05)		(0.20)		(0.33)
N	2035	1721	2035	1721	2035	1721
R-sqr	0.109	0.122	0.105	0.115	0.135	0.131

This table shows the relationship between the omission index and the gender a given paper. The dependent variable, omission index, indicates: the raw count of omission; the intensity of omission (sum omitted over sum cited and in prior literature) and the compensation (sum omitted minus sum citation). *female* represents papers written by only women. *at least one female* represents paper with at least one female author. The specifications include controls variables for the journal of publication, the institution, the number of authors, the year of publication, the field. Standard errors are clustered by journal of publication and years reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

Table 15. Relationship between Q-index and citation

log(1 + Forward citations, 0-1 year)	(1)	(2)	(3)	(4)	(5)	(6)
log(1 + q-index, 0-1 years)	0.105***	0.103***	0.105***	0.104***	0.102***	0.103***
	(0.011)	(0.010)	(0.010)	(0.009)	(0.009)	(0.009)
nbr_author		0.048***	0.039***	0.017***	0.008**	0.008**
		(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
dummynber		0.200***	0.195***	0.164***	0.112***	0.106***
		(0.013)	(0.013)	(0.011)	(0.012)	(0.012)
N	21286	21286	21286	21286	21286	21286
R-sqr	0.007	0.041	0.049	0.121	0.142	0.146
log(1 + Forward citations, 0-3 years)	(1)	(2)	(3)	(4)	(5)	(6)
log(1 + q-index, 0-3 years)	0.319***	0.309***	0.314***	0.299***	0.287***	0.292***
	(0.045)	(0.042)	(0.037)	(0.031)	(0.030)	(0.029)
nbr_author		0.114***	0.099***	0.042***	0.023***	0.022***
		(0.009)	(0.008)	(0.006)	(0.006)	(0.006)
dummynber		0.403***	0.392***	0.316***	0.199***	0.188***
		(0.021)	(0.022)	(0.016)	(0.017)	(0.017)
N	19471	19471	19471	19471	19471	19471
R-sqr	0.006	0.058	0.065	0.227	0.262	0.270
log(1 + Forward citations, 0-5 years)	(1)	(2)	(3)	(4)	(5)	(6)
log(1 + q-index, 0-5 years)	0.517***	0.513***	0.697***	0.609***	0.585***	0.587***
	(0.068)	(0.061)	(0.047)	(0.038)	(0.037)	(0.036)
nbr_author		0.159***	0.135***	0.061***	0.038***	0.038***
		(0.012)	(0.010)	(0.008)	(0.008)	(0.008)
dummynber		0.485***	0.459***	0.383***	0.243***	0.227***
		(0.026)	(0.027)	(0.018)	(0.020)	(0.020)
N	17173	17173	17173	17173	17173	17173
R-sqr	0.022	0.081	0.097	0.282	0.318	0.328
Publication Year FE			Y	Y	Y	Y
Journal FE				Y	Y	Y
Institution FE					Y	Y
Field FE						Y

This table shows the relationship between the quality index (q-index) and the number of citations. The controls include dummies for journals, field, institutions, year of publications, number of authors, NBER member. The q-index is built following equation 2.7.1. Standard errors are clustered by journal of publication and years reported in parentheses. ($* = p < 0.10$, $** = p < 0.05$,$*** = p < 0.01$)

Fig. 9. Number of authors per article

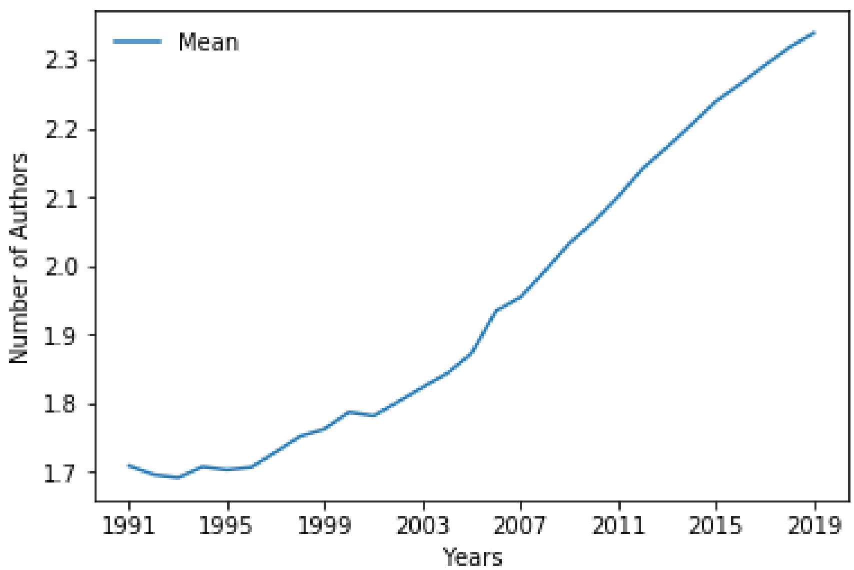

The figure shows the evolution of the average number of authors per article over time since 1991 to 2018.

Fig. 10. Experiment 2: Actual distribution-Target distribution

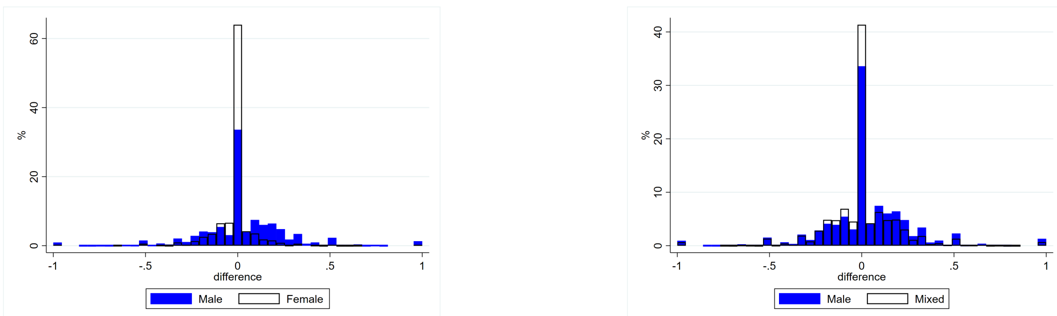

(a) Males versus Females **(b)** Males versus Mixed

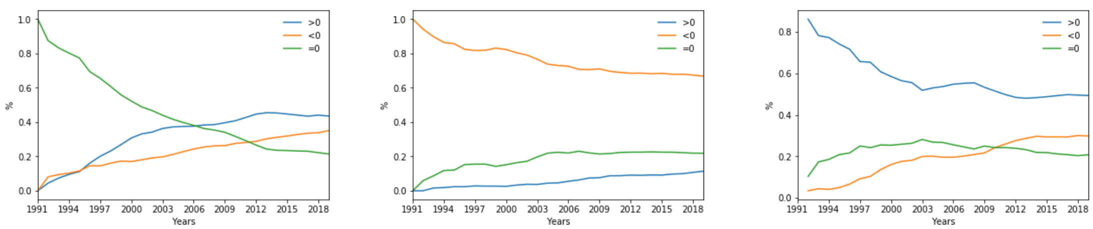

(c) Males **(d)** Females **(e)** Mixed

This figure plots the difference between the actual distribution and the target distribution of gender type g. For each paper p, the actual distribution of gender is the share of papers in its references belonging to each category of gender (only males, only females, mixed gender). The target distribution of gender is the share of the closest papers (in the sense of the relative cosine) in its prior literature belonging to each category of gender (only males, only females, mixed gender). For each paper, the difference between the actual distribution of gender g and the target distribution is taken. Panel (a) plots the distribution of this difference for males versus females. Panel (b) plots the distribution of this difference for males versus mixed gender teams. A positive difference means that the actual distribution of a certain gender type is higher than the target distribution of this gender type. A negative difference means that the actual distribution of a certain gender type is lower than the target distribution of this gender type. Finally, a null difference means that the actual distribution of a certain gender type is the same as the target distribution of this gender type. Panel (c), (d) and (e) plot the evolution of the share of papers with a positive, negative and null difference with respect to gender type male, female and mixed gender teams.

Exemple de citation : Consultez le LaTeX companien de Mittelbach *et al.* **?**.

Appendix C

Appendix to Chapter 3

Mathematical appendix

Some notions on Wiener Process:

Definition: There exists a probability distribution over the set of continuous functions $B : R \to R$ satisfying the following conditions:

(1) B(0) = 0.

(2) (stationary) for all 0 = s < t, the distribution of B(t) - B(s) is the normal distribution with mean 0 and variance t - s, and

(3) (independent increment) the random variables $B(t_i) - B(s_i)$ are mutually independent if the intervals $[s_i, t_i]$ are nonoverlapping.

Thus, dA_t is a Brownian motion means that $dA_t \sim N(0,1)$. Similarly is defined dz_{it}. But, whereas dA_t is drawn at each period, one time for all agents, dz_{it} is drawn at each period and for each agent. An example, assume there are two investors:

Period	1	2	3	4	5	6	7
dA							
Inv 1	-1	0	1	1	0	-1	0
Inv 2	-1	0	1	1	0	-1	0
dz							
Inv 1	1	-1	0	1	-1	-1	1
Inv 2	0	0	0	-1	1	1	-1
Cumulation							
Inv 1	0	-1	1	2	-1	-2	0
Inv 2	-1	0	1	0	1	0	0

Capital return:

$$dR_t^k = \frac{dP_t k_{i,t}^j + (a - \Gamma(\iota_{i,t}^j))k_{i,t}^j}{P_t k_{i,t}^j} \tag{C.0.1}$$

Where $dP_t k_{i,t}^j$ is obtained using generalized Ito product. Note for intuition that Ito product with a single brownian is written as followed. Consider two variables F and G such that:

$$\frac{dF}{F} = \mu_F dt + \sigma_F dZ_t$$

and

$$\frac{dG}{G} = \mu_F dt + \sigma_F dZ_t$$

Then

$$\frac{dFG}{FG} = (\mu_F + \mu_G + \sigma_F \sigma_G)dt + (\sigma_F + \sigma_G)dZ_t$$

Applying this formula to $P_t k_{i,t}^j$ gives:

$$\frac{dP_t k_{i,t}^j}{P_t k_{i,t}^j} = (\mu_{P,t} + \iota_t^j + \sigma^A \sigma_{P,t}^A)dt + (\sigma^A \sigma_{P,t}^A)dZ_t^A + (\sigma^B \sigma_{P,t}^B)dZ_t^B + v_t^A dW_{i,t}$$

The return from capital follows with equation (**??**).

Proof Proposition 1:

We focus on country A. The results are analogous for country B. The HJB of household, where we simplify the notation by abstracting from time- and household- indexes is:

$$0 = \max_{m^A, \iota^A, \hat{k}^A} \Bigg[- \tfrac{\rho}{1-\theta} + \tfrac{(m^A)^{1-\theta}}{1-\theta} \rho (\chi^A)^{\theta - 1} + \mu_{n^A}^A - m^A + \mu_{\chi^A}^A + \tfrac{\gamma}{2} \Big((\sigma_{n^A}^A)^2 + (\sigma_{\chi^A}^A)^2 \tag{C.0.2}$$

$$2 \tfrac{1-\gamma}{\gamma} \sigma_{n^A}^A \sigma_{\xi^A}^A + (\sigma_{n^A}^B)^2 + (\sigma_{\chi^A}^B)^2 - 2 \tfrac{1-\gamma}{\gamma} \sigma_{n^A}^B \sigma_{\xi^A}^B + (\tilde{\sigma}_{n^A}^A)^2 \Big) \Bigg] \tag{C.0.3}$$

Take the derivatives with respect to m^A, ι^A and $\hat{k}^A$, this give proposition 1.

$$(\sigma^A + \sigma_{p,t}^A)(\gamma \sigma_{n^A,t}^A - (1-\gamma)\sigma_{\chi^A,t}^A) + \sigma_{p,t}^B(\gamma \sigma_{n^A,t}^B - (1-\gamma)\sigma_{\chi^A,t}^B) + \gamma P_t \hat{k}_t^A (v_t^A)^2 - E(dR_t^A) + r_t$$

$$(\sigma^A + \sigma_{p,t}^A)(\gamma P_t \hat{k}_t^A (\sigma^A + \sigma_{p,t}^A) - (1-\gamma)\sigma_{\chi^A,t}^A) + \sigma_{p,t}^B(\gamma P_t \hat{k}_t^A \sigma_{p,t}^B) - (1-\gamma)\sigma_{\chi^A,t}^B) + \gamma P_t \hat{k}_t^A (v_t^A)^2 - E(dR_t^A) + r_t$$

The risk premium is not a function of individual because the expectation of the return from capital is the same for all individual. Looking at the fraction of net worth investing in capital:

$$P\hat{k}^A = \frac{E(dR) - r + (1-\gamma)\sigma_{\chi^A,t}^A(\sigma^A + \sigma_{p,t}^A) + (1-\gamma)\sigma_{\chi^A,t}^B \sigma_{p,t}^B}{\gamma((v^A)^2 + (\sigma^A + \sigma_{p,t}^A)^2 + (\sigma_{p,t}^B)^2)}$$

None of those variables are investor specific so each individual invests the same share of net worth in capital. Moreover, all the volatilities that enter in the denominator will tend to decrease this share.

Proof Proposition 2:

$$\begin{aligned} dK_t &= dK_t^A + dK_t^B \\ &= g_t K_t^A dt + \sigma^A K_t^A dZ_t^A + g_t K_t^B dt + \sigma^B K_t^B dZ_t^B \\ \frac{dK_t}{K_t} &= g_t dt + s^A \sigma^A dZ_t^A + s^B \sigma^B dZ_t^B \end{aligned} \tag{C.0.4}$$

Using Ito's lemma (for products) we get:

$$\frac{dP_tK_t}{P_tK_t} = (\iota_t + \mu_{p,t} + \sigma^A s_t^A \sigma_{p,t}^A + \sigma^B s_t^B \sigma_{p,t}^B)dt + (s^A\sigma^A + \sigma_{p,t}^A)dZ_t^A + (s^B\sigma^B + \sigma_{p,t}^B)dZ_t^B \tag{C.0.5}$$

Then:

$$\frac{d(P_tK_t)^{-1}}{(P_tK_t)^{-1}} = \left[-\iota_t - \mu_{p,t} - \sigma^A s_t^A \sigma_{p,t}^A - \sigma^B s_t^B \sigma_{p,t}^B + (s^A\sigma^A + \sigma_{p,t}^A)^2 + (s^B\sigma^B + \sigma_{p,t}^B)^2\right] dt - (s^A\sigma^A + \sigma_{p,t}^A)dZ_t^A - (s \tag{C.0.6}$$

Recall:

$$\frac{dN_t^A}{N_t^A} = (\mu_{n^A,t} - m_t^A)dt + P_t\hat{k}_t^A(\sigma^A + \sigma_{p,t}^A)dZ_t^A + P_t\hat{k}_t^A\sigma_{p,t}^B dZ_t^B \tag{C.0.7}$$

Hence, the process for x_t is:

$$\begin{aligned} dx_t &= d\frac{N_t}{P_tK_t} \\ dx_t &= d(N_t(P_tK_t)^{-1}) \end{aligned}$$

Hence, the process for x_t is:

$$\begin{aligned} \mu_{x,t} = x_t[&\mu_{n^A,t} - m_t^A - g_t - \mu_{p,t} - \sigma_{p,t}^A\sigma^A s_t^A - \sigma_{p,t}^B \sigma^B s_t^B + (\sigma_{p,t}^A + s^A\sigma^A)^2 + \\ &(\sigma_{p,t}^B + s^B\sigma^B)^2 - \sigma_{n^A,t}^A(\sigma_{p,t}^A + s^A\sigma^A) - \sigma_{n^A,t}^B(\sigma_{p,t}^B + s^B\sigma^B)] \end{aligned} \tag{C.0.8}$$

$$\sigma_{x,t}^A = x_t[\sigma_{n^A,t}^A - s^A\sigma^A - \sigma_{p,t}^A] \tag{C.0.9}$$

$$\sigma_{x,t}^B = x_t[\sigma_{n^A,t}^B - s^B\sigma^B - \sigma_{p,t}^B] \tag{C.0.10}$$

$$
\begin{aligned}
\sigma_{x,t}^A &= x_t[\sigma_{n^A,t}^A - s^A\sigma^A - \sigma_{p,t}^A] \\
&= x_t[P_t\hat{k}_t^A(\sigma^A + \sigma_{p,t}^A) - s^A\sigma^A - \sigma_{p,t}^A] \\
&= x_t[(P_t\hat{k}_t^A - 1)\sigma_{p,t}^A + (P_t\hat{k}_t^A - s^A)\sigma^A] \\
&= x_t[(\frac{s^A}{x_t} - 1)\sigma_{p,t}^A + (\frac{s^A}{x_t} - s^A)\sigma^A] \\
&= [(s_t^A - x_t)\sigma_{p,t}^A + (1 - x_t)s_t^A\sigma^A]
\end{aligned}
$$

Because $\sigma_{p,t}^A = \frac{1}{P}\left[\frac{\partial P}{\partial v^A}\sigma(v^A) + \frac{\partial P}{\partial x}\sigma_x^A\right]$, we then have:

$$
\sigma_{x,t}^A = \frac{(1 - x_t)s_t^A\sigma^A + \frac{(s_t^A - x_t)\frac{\partial P_t}{\partial v_t^A}}{P_t}\sigma(v^A)}{1 - \frac{(s_t^A - x_t)\frac{\partial P_t}{\partial x_t}}{P_t}} \tag{C.0.11}
$$

Similarly,

$$
\sigma_{x,t}^B = \frac{-x_t s_t^B\sigma^B + \frac{(s_t^A - x_t)\frac{\partial P_t}{\partial v_t^B}}{P_t}\sigma(v^B)}{1 - \frac{(s_t^A - x_t)\frac{\partial P_t}{\partial x_t}}{P_t}} \tag{C.0.12}
$$

For the special case of a closed economy, those equations are reduced to:

$$
\begin{aligned}
dx_t &= ds_t^A \\
&= d\frac{K_t^A}{K_t} \\
&= x_t(1 - x_t)[-x_t(\sigma^A)^2 + (1 - x_t)(\sigma^B)^2]dt + x_t(1 - x_t)\sigma^A dZ_t^A - x_t(1 - x_t)\sigma^B dZ_t^B
\end{aligned}
$$

The dynamic of x_t and s_t^A are the same.

Relation between share of capital and share of aggregate wealth :

Let s^j be the share of capital in country j as a fraction of total capital in the world. It then follows that:

$$
\begin{aligned}
P_t\hat{k}_t^A &= P_t\frac{K_t^A}{N_t^A} \\
&= P_t\frac{s_t^A K_t}{x_t P_t K_t} \\
&= \frac{s_t^A}{x_t}
\end{aligned}
$$

The relationship between s_t^j and x_t is given by:

$$\begin{aligned} x_t &= \frac{N_t^A}{P_t K_t} \\ &= \frac{P_t K_t^A + B_t^A}{P_t K_t} \\ &= s_t^A + \frac{\frac{B_t^A}{K_t}}{P_t} \end{aligned}$$

Thus, when the economies are closed, then $x_t = s_t^A$

Proof of Proposition 3

To show proposition 3, it is sufficient to prove that $\frac{P_x}{P} < 1$. As P is increasing in x and concave, suppose that there is an $\tilde{x}$ such that $\frac{P_x(x)}{P(x)} > 1$ for $x \succeq \tilde{x}$. We can write:$\tilde{x} = \bar{x} + \epsilon$ with $\frac{P_x(\bar{x})}{P(\bar{x})} < 1$.

$$\begin{aligned} \frac{P_x(\tilde{x})}{P(\tilde{x})} > 1 \Rightarrow \frac{P_x(\bar{x}+\epsilon)}{P(\bar{x}+\epsilon)} &> 1 \\ \Rightarrow \frac{P'_x(\bar{x}) + \epsilon P''_x(\bar{x}) + o(\epsilon)}{P(\bar{x}) + \epsilon P'_x(\bar{x}) + o(\epsilon)} &> 1 \\ \Rightarrow P'_x(\bar{x}) + \epsilon P''_x(\bar{x}) + o(\epsilon) > P(\bar{x}) + \epsilon P'_x(\bar{x}) + o(\epsilon) \\ \Rightarrow P'_x(\bar{x}) + \epsilon P''_x(\bar{x}) + o(\epsilon) > P(\bar{x}) + \epsilon P'_x(\bar{x}) + o(\epsilon) \\ \Rightarrow (1-\epsilon)P'_x(\bar{x}) + \epsilon P''_x(\bar{x}) > P(\bar{x}) \\ \Rightarrow (1-\epsilon)P'_x(\bar{x}) + \epsilon P''_x(\bar{x}) > P(\bar{x}) > P'_x(\bar{x}) \text{ By definition of } \bar{x} \\ \Rightarrow P''_x(\bar{x}) > P'_x(\bar{x}) \end{aligned}$$

But $P'_x > 0$ and $P''_x < 0$, this brings a contradiction.

Sign of B_t

Let us consider the following equality:

$$B_t^A = \tilde{B}_t^A(v^A, v^B, x) P_t K_t$$

Then,

$$\frac{\sigma^A_{B_t^A}}{\tilde{B}_t^A} = -(\sigma^A + \sigma_p^A)$$

For $(\sigma^A + \sigma_p^A) > 0$, $\frac{\sigma^A_{B_t^A}}{\tilde{B}_t^A} < 0$ so that an increase in v^A tends to be associated with an increase in $\tilde{B}_t^A(v^A, v^B, x)$: $\tilde{B}_t^A(v^A, v^B, x)$ is an increasing function of v^A. Similarly, $\tilde{B}_t^B(v^A, v^B, x)$

is an increasing function of v^B. Therefore, for $v^A < v^B$, $\tilde{B}_t^B(v^A, v^A, x) < \tilde{B}_t^B(v^A, v^B, x)$ and using bond market clearing condition $\tilde{B}_t^A(v^A, v^B, x) = -\tilde{B}_t^B(v^A, v^B, x)$, we find that: $\tilde{B}_t^A(v^A, v^B, x) < \underbrace{\tilde{B}_t^A(v^A, v^B, x)}_{0} < \tilde{B}_t^B(v^A, v^B, x)$.

This implies that $x_t - s_t^A > 0$ for the riskiest country and $x_t - s_t^A < 0$ for the other country.

Proof of proposition 5

$$\frac{ds_t^A}{s_t^A} = \frac{d\frac{K_t^A}{K_t}}{\frac{K_t^A}{K_t}} \tag{C.0.13}$$

$$= (1 - s_t^A)[(1 - s_t^A)(\sigma^B)^2 - s_t^A(\sigma^A)^2]dt + (1 - s_t^A)\sigma^A dZ_t^A - (1 - s_t^A)\sigma^B dZ_t^B \tag{C.0.14}$$

Let $\sigma^A_{s^A,t} = (1 - s_t^A)\sigma^A$ and $\mu_{s^A,t} = (1 - s_t^A)[(1 - s_t^A)(\sigma^B)^2 - s_t^A(\sigma^A)^2]$.
$\mu_{s^A,t} > 0$. $\sigma^A_{s^A,t} > 0$ then $\frac{\partial s^A}{\partial v^A} < 0$. A bad shock is associated with a decrease in s_t^A. Moreover,

$$\frac{\partial \sigma^A_{s^A,t}}{\partial v_t^A} = -\sigma^A \frac{\partial s_t^A}{\partial v_t^A} \tag{C.0.15}$$

As $\frac{\partial s_t^A}{\partial v_t^A} < 0$, then $\frac{\partial \sigma^A_{s^A,t}}{\partial v_t^A} > 0$.
The sign of $\frac{\partial \tilde{B}_t^A(v^A,v^B,x)}{\partial v^A}$ could also be viewed by the direct effect of v^A given x. In fact: $\frac{\partial \tilde{B}_t^A(v^A,v^B,x)}{\partial v^A} = -\frac{\partial s^A}{v^A} > 0$.

Investors' balance sheet in autarchy versus openness

In autarchy:

$$\frac{(\sigma_x^A)_{aut}}{x} = (1 - x)\sigma^A$$

In Integration:

$$\frac{(\sigma^A_{x,t})_{int}}{x_t} = \frac{(1 - x_t)\frac{s_t^A}{x_t}}{1 - (s_t^A - x_t)\frac{P_{xt}}{P_t}}\sigma^A + \frac{(\frac{s_t^A}{x_t} - 1)\frac{P_{vt}}{P_t}}{1 - (s_t^A - x_t)\frac{P_{xt}}{P_t}}\sigma(v^A)$$

For $0 < 1 - (s_t^A - x_t)\frac{P_{xt}}{P_t} < 1$, and $\frac{s_t^A}{x_t} > 1$, $\frac{(1-x_t)\frac{s_t^A}{x_t}}{1-(s_t^A-x_t)\frac{P_{xt}}{P_t}}\sigma^A > (1 - x_t)\sigma^A$ and $\frac{(\frac{s_t^A}{x_t}-1)\frac{P_{vt}}{P_t}}{1-(s_t^A-x_t)\frac{P_{xt}}{P_t}}\sigma(v^A) > 0$. Then, $(\sigma^A_{x,t})_{aut} < (\sigma^A_{x,t})_{int}$ for $x \in (0,1)$.

In Autarchy:

$$(\sigma^B_{1-x})_{aut} = x_t(1 - x_t)\sigma^B$$

In Integration:

$$(\sigma^B_{1-x})_{int} = \frac{(1-s^A_t)x_t}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}\sigma^B + \frac{(s^A_t - x_t)\frac{P_{vt}}{P_t}}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}\sigma(v^B)$$

Let $Diff^B = (\sigma^B_{1-x})_{int} - (\sigma^B_{1-x})_{aut}$ For $v^B \to 0$, $Diff^B \to \frac{(1-s^A)x}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}\sigma^B - x(1-x)\sigma^B > 0$ because of $0 < 1-(s^A_t - x_t)\frac{P_{xt}}{P_t} < 1$ and $1-s^A > 1-x$ for relatively low values of v^B.

For $v^B \to 1$, $Diff^B \to \frac{-(1-x)\frac{P_{vt}}{P_t}}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}\sigma(v^B) - x(1-x)\sigma^B < 0$ because of $0 < 1-(s^A_t - x_t)\frac{P_{xt}}{P_t} < 1$ and $P_{vt} < 0$ for relatively high values of v^B. Following Bolzano's theorem, $Diff^B$ must be 0 at some point. Therefore, for high values of v^B, $(\sigma^B_{1-x})_{aut} > (\sigma^B_{1-x})_{int}$ and for low values of v^B, $(\sigma^B_{1-x})_{aut} < (\sigma^B_{1-x})_{int}$.

Exposure to foreign shock

Consider $\sigma^B_{x,t}$ the exposure of the wealth share of country A to foreign shock and $\sigma^A_{1-x,t}$ the exposure of wealth share of country B to foreign shock:

$$\frac{\sigma^B_{x,t}}{x_t} = \underbrace{\frac{-s^B_t}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}}_{\text{Aggregate Amplification}}\sigma^B + \underbrace{\frac{(\frac{s^A_t}{x_t}-1)\frac{P_{v^B t}}{P_t}}{1-(s^A_t - x_t)\frac{P_{xt}}{P_t}}}_{\text{Idiosyncratic Amplification}}\sigma(v^A)$$

$$\frac{\sigma^B_{1-x,t}}{1-x_t} = \underbrace{\frac{-s^A_t}{1+(s^B_t-(1-x_t))\frac{P_{xt}}{P_t}}}_{\text{Aggregate Amplification}}\sigma^A + \underbrace{\frac{-(s^A_t - x_t)\frac{P_{vt}}{P_t}}{1+(s^B_t-(1-x_t))\frac{P_{xt}}{P_t}}}_{\text{Idiosyncratic Amplification}}\sigma(v^A)$$

NB: error to correct in this equation

Let T^j_3 and T^j_4 be respectively the expression that gather the aggregate and the idiosyncratic amplification, $j \in A,B$. Assume that A is the riskiest country and $1+(s^B_t-(1-x_t))\frac{P_{xt}}{P_t} > 0$. Case 1: Economy with aggregate shock and time-varying idiosyncratic risk.

- $T^A_3 < 0$ and $T^A_3 < 0$
- $T^B_3 < 0$ and $T^B_3 > 0$

Case 2: Economy with aggregate shock and no time-varying idiosyncratic risk.

- $T^A_3 < 0$ and no T^A_3
- $T^B_3 < 0$ and no T^B_3

Case 3: autarchy Economy

$$\sigma^B_{x,t} = -x_t(1-x_t)\sigma^B$$
$$\sigma^A_{1-x,t} = -x_t(1-x_t)\sigma^A$$

- $\sigma^B_{x,t} < 0$

- $\sigma^A_{x,t} < 0$

Again, the mitigation is only present when there is a time-varying risk.

Balance sheet effect

As v^A_t increases, $\sigma^A_{x,t}$ tends to decrease, same for $\sigma^A_{p,t}$.

$$\sigma^A_{n^A,t} = P_t \hat{k}^A_t(\sigma^A + \sigma^A_{p,t}) \tag{C.0.16}$$

$$= \tfrac{s^A_t}{x_t}(\sigma^A + \sigma^A_{p,t}) \tag{C.0.17}$$

Then:

- $s^A_t > x_t \Rightarrow \sigma^A_{n^A,t} > \sigma^A + \sigma^A_{p,t}$: The effect on balance sheet is amplified when the country hold more capital.
- $s^A_t < x_t \Rightarrow \sigma^A_{n^A,t} < \sigma^A + \sigma^A_{p,t}$: the effect on balance sheet is reduced when the country hold less capital.
- $s^A_t = x_t \Rightarrow \sigma^A_{n^A,t} = \sigma^A + \sigma^A_{p,t}$: the effect of s^A/x disappears.

As v^A_t increases, $s^A_t < x_t$ so that $\sigma^A_{n^A,t}$ will also tend to be lower in financial openness compared to the autarchy case.

Dynamic of P_t, χ^A_t and χ^B_t

Using Ito's lemma, all variables can be written as a function of the state-variable vector. Consider a variable H will be a function $H = H(x_t,\nu^A_t,\nu^B_t)$ and therefore:

$$\begin{aligned} dH(x_t,v^A_t,v^B_t) = \frac{\partial H}{\partial v^A}dv^A_t + \frac{\partial H}{\partial v^B}dv^B_t + \frac{\partial H}{\partial x}dx_t + \frac{1}{2}[\frac{\partial^2 H}{\partial (v^A)^2}dv^A_t.dv^A_t + \frac{\partial^2 H}{\partial (v^B)^2}dv^B_t.dv^B_t + \\ \frac{\partial^2 H}{\partial x^2}dx_t.dx_t + 2\frac{\partial^2 H}{\partial v^A \partial v^B}dv^A_t.dv^B_t + 2\frac{\partial^2 H}{\partial v^A \partial x}dv^A_t.dx_t + 2\frac{\partial^2 H}{\partial v^B \partial x}dv^B_t.dx_t] \end{aligned} \tag{C.0.18}$$

The expressions for the derivatives are therefore:

$$H\sigma^A_H = \frac{\partial H}{\partial v^A}\sigma(v^A) + \frac{\partial H}{\partial x}\sigma^A_x \tag{C.0.19}$$

$$H\sigma^B_H = \frac{\partial H}{\partial v^B}\sigma(v^B) + \frac{\partial H}{\partial x}\sigma^B_x \tag{C.0.20}$$

$$\begin{aligned} H\mu_H = \frac{\partial H}{\partial v^A}\mu_{v^A} + \frac{\partial H}{\partial v^B}\mu_{v^B} + \frac{\partial H}{\partial x}\mu_x + \frac{1}{2}[\frac{\partial^2 H}{\partial (v^A)^2}(\sigma(v^A))^2 + \\ \frac{\partial^2 H}{\partial (v^B)^2}(\sigma(v^B))^2 + \frac{\partial^2 H}{\partial x^2}((\sigma^A_x)^2 + (\sigma^B_x)^2) + 2\frac{\partial^2 H}{\partial v^A \partial x}\sigma(v^A)\sigma^A_x + 2\frac{\partial^2 H}{\partial v^B \partial x}\sigma(v^B)\sigma^B_x] \end{aligned} \tag{C.0.21}$$

Some Empirical Facts

In this part, we present some empirical facts based on the database “External Wealth of Nations Dataset" Lane and Milesi-Ferretti(2007)

Fig. 1. World Gross financial flows over World Total GDP

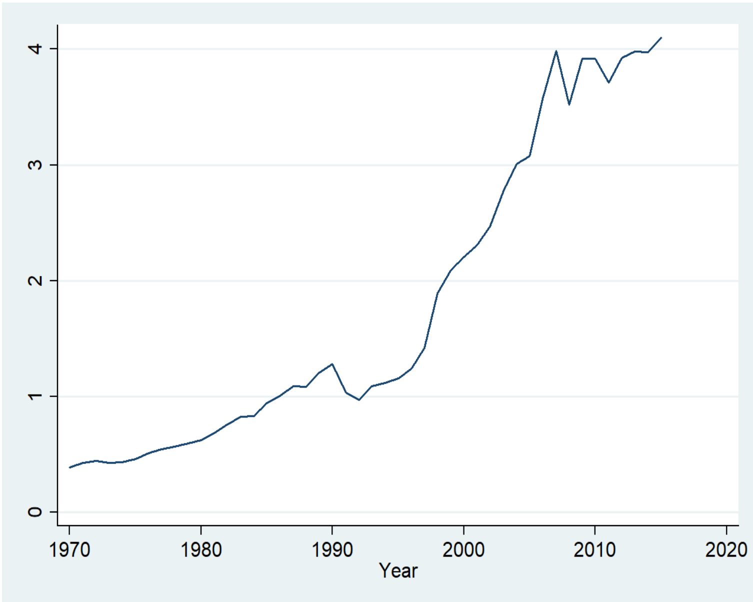

Source: Updated database “External Wealth of Nations Dataset" Lane and Milesi-Ferretti(2007); Gross flows included equities, bonds, FDI assets and liabilities. There is a slowdown in the evolution of Gross financial flows.

Fig. 2. Growth rate of the World Gross financial flows over World Total GDP and the growth rate of the world GDP

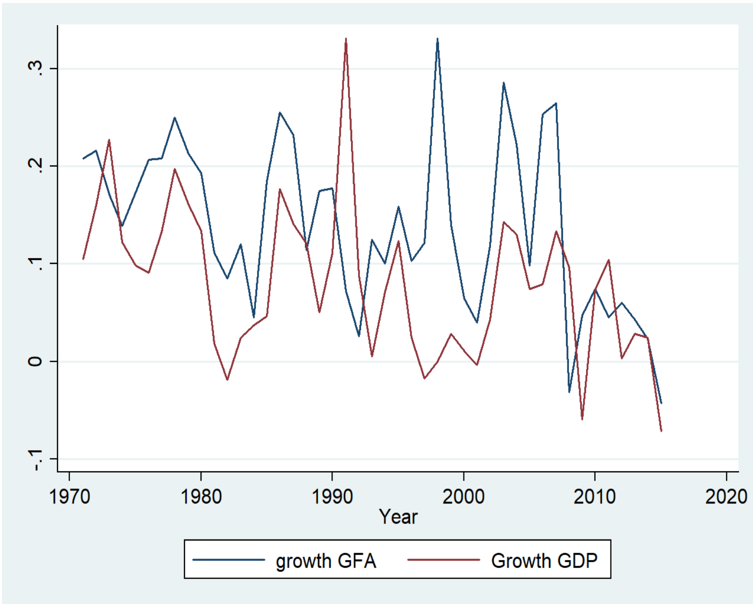

This figure presents the growth rate of the World Gross financial flows over World Total GDP and the growth rate of the world GDP; Source: Updated database “External Wealth of Nations Dataset" Lane and Milesi-Ferretti(2007); Gross flows included equities, bonds, FDI assets and liabilities. An increase in the GDP growth rate is associated with an increase in the world gross financial flows.

Fig. 3. NFA/GDP by country category

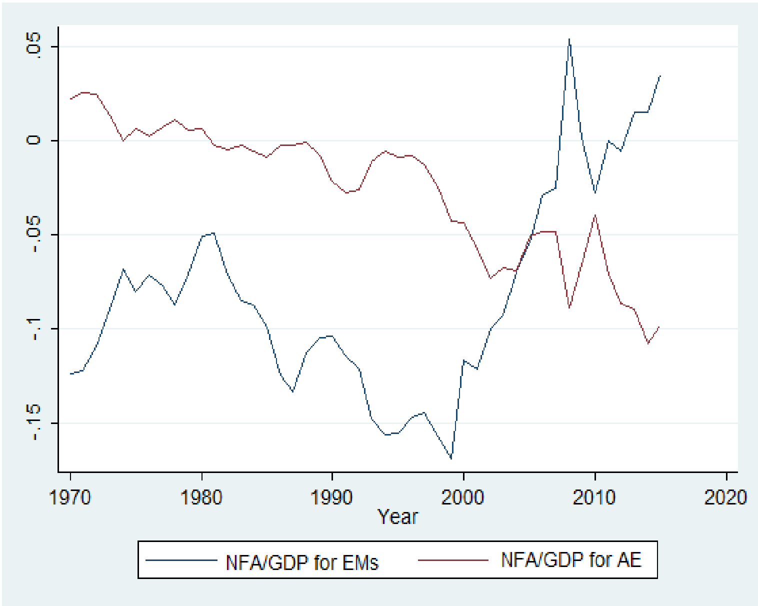

Emerging markets net foreign assets over Emerging market GDP (blue curve) and advanced economies net foreign assets over advanced economy GDP(red curve); Source: Updated database "External Wealth of Nations Dataset" Lane and Milesi-Ferretti(2007); Financial flows included equities, bonds, FDI assets and liabilities. Emerging markets tend to experiment higher net foreign assets compared to advanced economies in the recent years.

Fig. 4. Relative idiosyncratic risk and net foreign asset

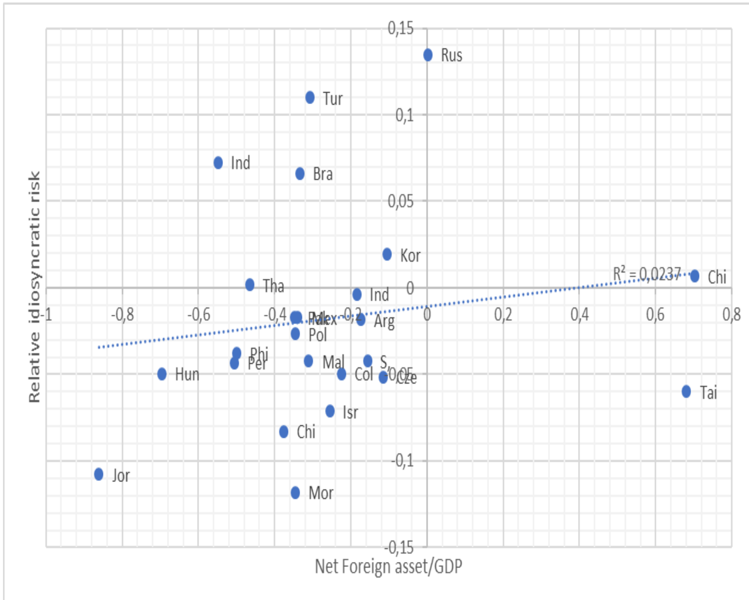

Relative idiosyncratic risk and net foreign asset for a set of 21 emerging markets. The data for the idiosyncratic risk are form Umultu et al. (2009) and cover the period 1980-2005. We compute the relative idiosyncratic risk by substracting from the level of the idiosyncratic risk for a given country the average value of the idiosyncratic risk from the other countries. The data for the net foreign asset (NFA) are from the updated database “External Wealth of Nations Dataset" Lane and Milesi-Ferretti(2007); Financial flows included equities, bonds, FDI assets and liabilities. An increase in the relative idiosyncratic risk is associated with an increase in the NFA positions.