Issue 2 / July 2015

Economics@Hull

A magazine about current economic issues and the research of Hull's economics group

In this issue

From the editor

Welcome to the second issue of Economics@Hull. The first article presents the results of research investigating local protectionism in the Chinese pharmaceutical industry. There are also articles summarizing research into the impact of credit ratings announcements on the stock market, and research on the convergence of the income of poorer countries and richer countries. The final article explores how market frictions can explain the observation that different stores often sell the same product at different prices.

If there are any economic issues that you would like to see examined in a future issue of Economics@Hull, or you have any general feedback or comments please email us at the address below.

Scott McCracken
Hull University Business School
University of Hull
Cottingham Road
Hull HU6 7RX
s.mccracken@hull.ac.uk

Revisiting the Razor's Edge: Local protectionism in China

Zheng Wang
University of Hull
z.wang@hull.ac.uk

Markus Eberhardt
University of Nottingham

Zhihong Yu
University of Nottingham

Over the past decades economists have followed the development of China with great interest, typically siding with one of two views on the sources and prospects of the Middle Kingdom's growth. One camp argues for the existence of a 'China model', variously defined but typically representing state control over key industries. A second camp views China on a trajectory from planned to socialist market to free market economy and emphasizes the progress made so far. Our research (Eberhardt et al., 2014) informs this debate, focusing narrowly on market integration, one aspect of a potential transition to a market-led system.

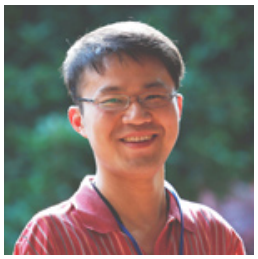

Dr Zheng Wang is a lecturer in economics at Hull University Business School and an external research fellow of GEP at the University of Nottingham. Prior to joining the University of Hull in 2013, he worked as an assistant professor at the University of Nottingham China Campus after getting his PhD from the same university (UK campus). His research concentrates on the areas of international trade, development economics, and political economy with a special focus on transition economies such as China. He teaches international trade modules at both the undergraduate and postgraduate levels.

Hull University Business School

A seminal study by Young (2000) compared data from China's Socialist period with periods well after the economic reforms began, and suggested that although China had opened up internationally it had become internally fragmented. The author argued that any economic system where the state played a significant role in determining prices, investment and output did not represent a viable strategy for long-run prosperity, and that China was in danger of falling off the 'razor's edge' that marks the narrow path towards establishing a successful market economy. Holz (2009) rejected these findings on grounds of weak evidence for a causal link between reforms and market fragmentation and argued that based on his own analysis, the 'razor's edge' may be more suitably described as a highway, along which China proceeded rather comfortably.

A common factor in these and other studies of inter-regional trade barriers in China is that they try to identify aggregate and implicit measures of protectionism: in the absence of significant regional protectionism one would for instance expect regional specialization in production and/or price convergence across provinces to be detectable in the data, which is the focus of analysis in the existing empirical literature. Our study is the first to use firm-level data to investigate market integration in China, employing information on the public disclosure of 'illegal' drug advertisements by provincial Food and Drug Administrations (FDAs). In present day China a large number of pharmaceutical enterprises are vying for a share of the 'over the counter' market in generic drugs. Since advertising is crucial for these firms and drug advertising standards, set by the State FDA in Beijing, are ambiguous, provincial FDAs can selectively inspect and 'disclose' nonlocal firms for 'illegal' advertising, which typically entails fines or more severe forms of punishment. The ambiguity of drug advertising regulations is indicated in a 2004 report by the state-run news agency Xinhua, which revealed that strict application of drug advertising standards would result in 62 per cent of all advertisements broadcast on television and 95 per cent of all newspaper advertisements being classified as 'illegal'. Provincial FDAs, whose funding and senior-level appointments are controlled by the local government, have an incentive to discriminate against nonlocal firms for at least two reasons. Firstly, fiscal decentralization in the 1990s redefined the split of tax revenue between central and provincial governments and was intended to provide incentives to push for local development and thus boost local tax revenue. Secondly, political promotions of provincial governors are closely linked to local economic performance, including gross provincial product and tax revenue.

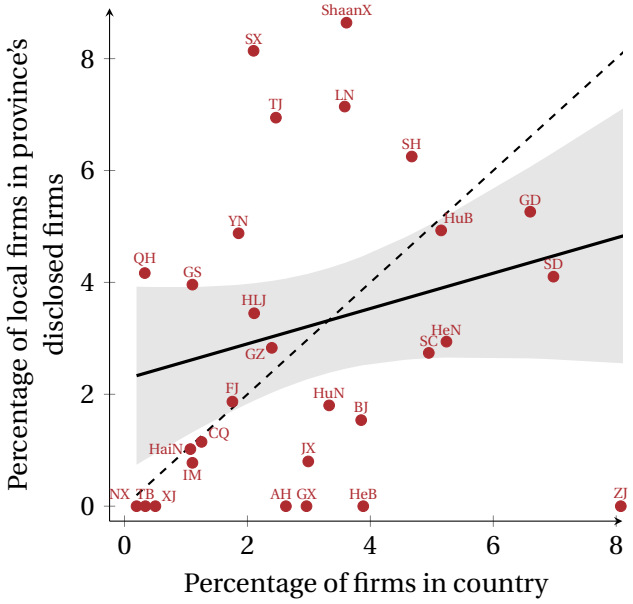

Figure 1: Each dot is a province, the solid line represents the line of best fit (from an OLS regression), the dashed line is the 45 degree diagonal, and the shaded area indicates the 90 per cent confidence interval.

Our descriptive analysis shows that nonlocal firms are disproportionally targeted for disclosure, constituting protectionism on behalf of the provincial FDAs. Figure 1 provides an illustration of the discrimination we detect in our sample of all state-owned and all medium and large non-state firms in China's pharmaceutical industry over the 2001–2005 period: if we assume that all firms sell their drugs in all 31 provinces then in the absence of protectionism the share of local firms being disclosed should be in line with the relative size of the local pharmaceuticals industry (the 45 degree line). As can be seen, local firms on average get disclosed disproportionally less than nonlocal ones. Since the assumption that all firms sell their products in all provinces is difficult to maintain, we carry out our regression analysis using additional information on compulsory advertising licences available for a subsample of provinces: Jiangsu, Zhejiang and Inner Mongolia. The first two host the largest number of pharmaceutical firms and are among the most developed regions in the country. In contrast, Inner Mongolia is a peripheral province characterized by mining and livestock breeding. Controlling for province-level idiosyncrasies and a host of other firm-level factors we find that the disclosure probability is 9–13 per cent higher for nonlocal firms. Investi-

gating whether this form of protectionism changed over time we find higher discrimination in the more recent sample years, in line with Young's findings at the aggregate level.

In further analysis we attempt to provide evidence for three specific motivations for discrimination against nonlocal firms. Using the same sample as above we first test the influence of government affiliation (lishu) on disclosure probability. The lishu system represents a uniquely Chinese institution where the 'iron fist' of the planned economy meets the 'invisible hand' of the market. A lishu relationship can be at various levels of government, is distinct from ownership status and entails government control as well as subsidies and support. Existing research on the economic implication of these relationships is somewhat mixed, finding improved access to credit for higher-level affiliated firms in one study and improved firm performance upon relinquishing of the affiliation in another. Our own analysis suggests that local firms with an affiliation to the provincial government are underrepresented among disclosed firms, whereas nonlocal firms of the same affiliation are targeted disproportionately. We interpret this finding as pointing to the strong incentives for provinces to compete with each other.

Provincial competition is also investigated more explicitly in its relationship with disclosure patterns: FDAs might discriminate more readily against firms from provinces with a large pharmaceutical industry, in order to hit the rival province's tax revenues. Using data for firms and disclosure in all 31 provinces we conclude that this channel indeed seems to be driving the patterns of discrimination.

A final motivation for discrimination against nonlocal firms relates to the importance of the pharmaceutical industry for the disclosing province: if there are only a small number of local drug firms which do not create substantial tax revenue then provincial FDAs may feel less inclined to protect them and thus safeguard their local fiscal base. Again using the full sample of firms we find evidence for selective discrimination along these lines. For both the provincial competition and fiscal base motives we conclude that the increase in nonlocal discrimination is economically significant along the channels described.

Since the mechanism of discrimination we detect points to relatively generic institutional roots of protectionism, we believe these findings have wider validity beyond the pharmaceutical sector. One policy implication is that reducing intra-national barriers perhaps poses a greater challenge than lowering international barriers since the former typically requires much deeper and wide-ranging domestic reforms, both economic and political.

References

Eberhardt, Markus, Zheng Wang, and Zhihong Yu (2014), "From One to Many Central Plans: Drug Advertising Inspections and Intra-National Protectionism in China." *CESifo Working Paper*.

Holz, Carsten A (2009), "No Razor's Edge: Reexamining Alwyn Young's Evidence for Increasing Interprovincial Trade Barriers in China." *The Review of Economics and Statistics*, 91, 599–616.

Young, Alwyn (2000), "The Razor's Edge: Distortions and Incremental Reform in the People's Republic of China." *The Quarterly Journal of Economics*, 115, pp. 1091–1135.

The information content of rating announcements on Frankfurt Stock Exchange

Amangeldi Kenjegaliev
University of Hull
a.kenjegaliev@hull.ac.uk

Djamila Mamedshakhova
Standard and Poor's

Investors often find making investment decisions challenging due to the time, effort and financial costs required to conduct a deep analysis of the riskiness of the investment alternatives available. To reduce costs, investors employ credit ratings to assess project riskiness. Rating announcements are treated as market signals which stem from informational asymmetry existing between debt issuers and investors. Credit ratings are meant to serve as yardsticks to investors, and the assertion that 'a low rating is better than no rating' is as relevant as ever. For institutional investors the distinction between investment and non-investment-grade ratings of particular securities is essential when considering their investment portfolios. Moreover, banks use credit ratings to set lending interest rates and to control the level of required capital. Hence, the rating announcements by credit agencies have a significant impact on the securities market.

Dr Amangeldi Kenjegaliev is a lecturer in economics at Hull University Business School. A graduate of Tashkent Financial Institute, he received his PhD from the Department of Economics, University of Leicester in 2012 and joined the University of Hull in 2013. Prior to the current post he was a Research Assistant at Fraser of Allander Institute, University of Strathclyde and associate tutor at the University of Leicester School of Management. His main research interests are international finance, exchange rates, time-varying coefficient regression, econometrics, and macroeconomics.

Not all academics support this argument. Altman and Saunders (2001), for example, argue that rating agencies are too slow and inflexible, and question the ability of ratings to predict default. Instead, according to Altman and Saunders, ratings merely follow the deterioration of credit quality; and usage of ratings in controlling for risk should be limited. Numerous empirical studies since the 1990s have investigated the actual capital market reaction to credit ratings changes. Ratings changes include rating upgrades and downgrades, as well as positive and negative rating reviews and outlooks. The most common objects of investigation are stock returns, bond yields and credit default swap spread changes. The studies have reached mixed results. While some authors like Jewell and Livingston (1999) point out a consensus in the academic literature that debt-security ratings provide useful information to the market, others like Partnoy (2002), Holthausen and Leftwich (1986) assert that credit ratings carry scant informational value, and much of the adjustment to security prices takes place long before the ratings announcement.

In Kenjegaliev and Mamedshakhova (2015) we extend the existing research on the price impact of rating actions on the stock market, and thus contribute to reducing the existing uncertainty regarding the information content of rating changes. We examine the stock price impact of the credit rating announcements made by three global rating agencies—Moody's Investors Service, Standard & Poor's and Fitch Ratings—using the data of the largest companies traded on the Frankfurt Stock Exchange between 2002 and 2007. The period following that is excluded from the research due to the 2008 financial crisis. The Frankfurt Stock Exchange was chosen due to the high liquidity of the market and intensity of trading of financial instruments, so that stock price volatility could be observed on a daily basis.

The analysis uses event study methodology and calculates abnormal stock returns of the HDAX (Frankfurt Stock Exchange index) constituents representing 110 major German companies as a reaction to rating announcements. An event study entails selection of a date when a particular activity has occurred (in our case a public announcement made by rating agencies) and examining the period surrounding this date. This

technique has been widely employed by researchers to investigate the impact of ratings announcements and to examine the impact of share splits, earnings and merger announcements. We set up an event window of four months, which starts 60 business days before a rating announcement and ends 20 business days after the announcement.

Table 1: Number of ratings announcements from January 1 2002 to September 1 2007

Agency	Upgrades	Downgrades	Total
Fitch	14	20	34
Moody's	17	17	34
S & P	29	34	63
Total	60	71	131

Table 1 presents the distribution of rating announcements by rating agencies. We assess the impact of positive and negative rating announcements on the stock market. In particular, if a rating change adds new information to the market, we expect that upgrades are associated with an immediate increase in share prices, while downgrades are associated with a decrease in share prices.

Fitch Ratings

Moody's

Standard & Poor's

A stock price reaction is a significant change in a stock return over the analysed period. To identify if these changes are caused by a general upward (or downward) trend in the market or credit ratings announcements we firstly calculated the abnormal return, which is the difference between the realized return and expected return on the security. Once the abnormal stock returns for all securities are calculated, the residuals are averaged across the firms to produce the average residual for each day (AAR). Such averaging helps to cancel out the 'noise' in the stock returns across the firms. Following that, abnormal returns are summed up over the analysed period to obtain the cumulative average abnormal return (CAAR).

Our results show that credit ratings carry scant informational value. This is evident from Figure 1 which shows the movement of the average abnormal returns observed around rating upgrades and downgrades. Figure 2 plots the the cumulative average abnormal returns. Share prices of the concerned companies do not react significantly to a rating change; instead, adjustments to stock prices are being made long before the rating announcement date. This market behaviour is noted for both upgrades and downgrades. The reported findings support the inference that rating agencies lack timeliness. The financial market immediately absorbs all the information related to the issuers, which becomes publicly available by the efforts of insiders (mostly in the case of negative information) or by means of the companies' management (mostly in the case of positive information). Rating agencies incorporate this information into rating grades later on.

The implication of this observation can be interpreted in terms of the position being held by rating agencies among the financial market participants. Based on the reported results, the predictive power of rating agencies is not proved. Thus, one can hardly claim that investors use the rating assessment as the major benchmark when pricing the companies. Apparently, investors rely on the more timely sources of news related to the issuers to which they have access. That being said, the importance of rating agencies should not be neglected. Investors may adjust share prices at the reported level precisely because they are confident in the rating upgrade/downgrade which will follow the released news in the near future, along with its consequences. Knowing this, the managers of the rated companies can envisage the possible implications of rating changes and adjust their strategy of communicating with investors by incorporating appropriate and timely decisions into it.

References

Altman, Edward I and Anthony Saunders (2001), "An Analysis and Critique of the BIS Proposal on Capital Adequacy and Ratings." *Journal of Banking & Finance*, 25, 25–46.

Holthausen, Robert W and Richard W Leftwich (1986), "The Effect of Bond Rating Changes on Common Stock Prices." *Journal of Financial Economics*, 17, 57–89.

Jewell, Jeff and Miles Livingston (1999), "A Comparison of Bond Ratings from Moody's S&P and Fitch IBCA." *Financial Markets, Institutions & Instruments*, 8, 1–45.

Kenjegaliev, Amangeldi and Djamila Mamedshakhova (2015), "Information content of rating announcements." Mimeo.

Partnoy, Frank (2002), "The Paradox of Credit Ratings." In *Ratings, Rating Agencies and the Global Financial System* (Richard M Levich, Giovanni Majnoni, and Carmen M Reinhart, eds.), 65–84, Springer US.

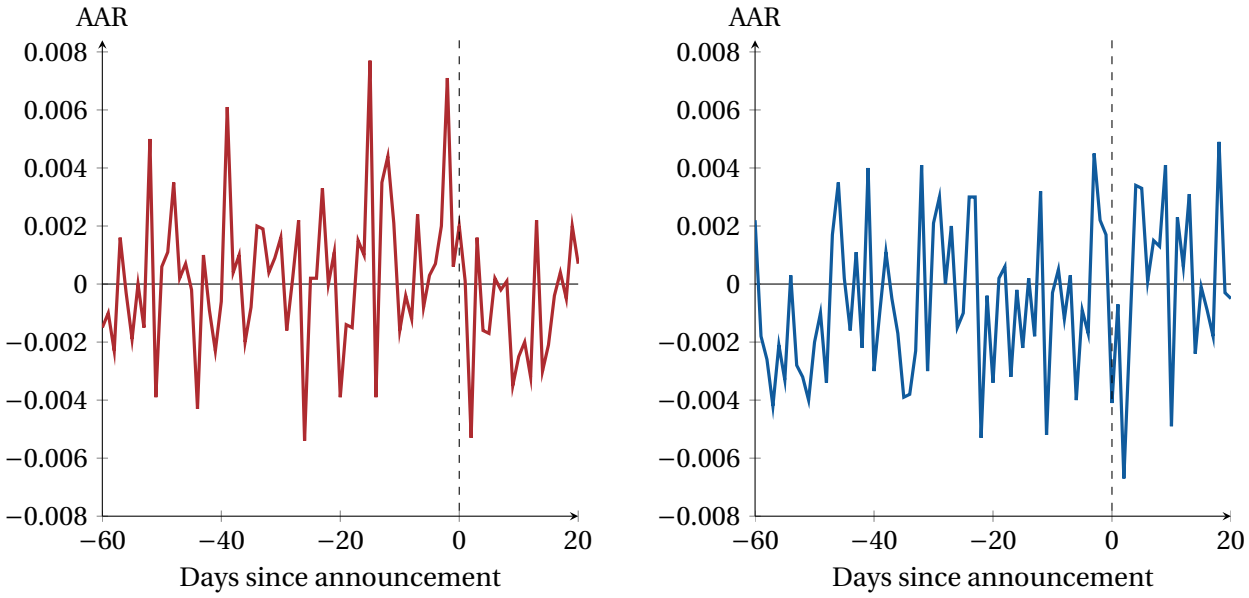

(a) Upgrades

(b) Downgrades

Figure 1: Average abnormal returns

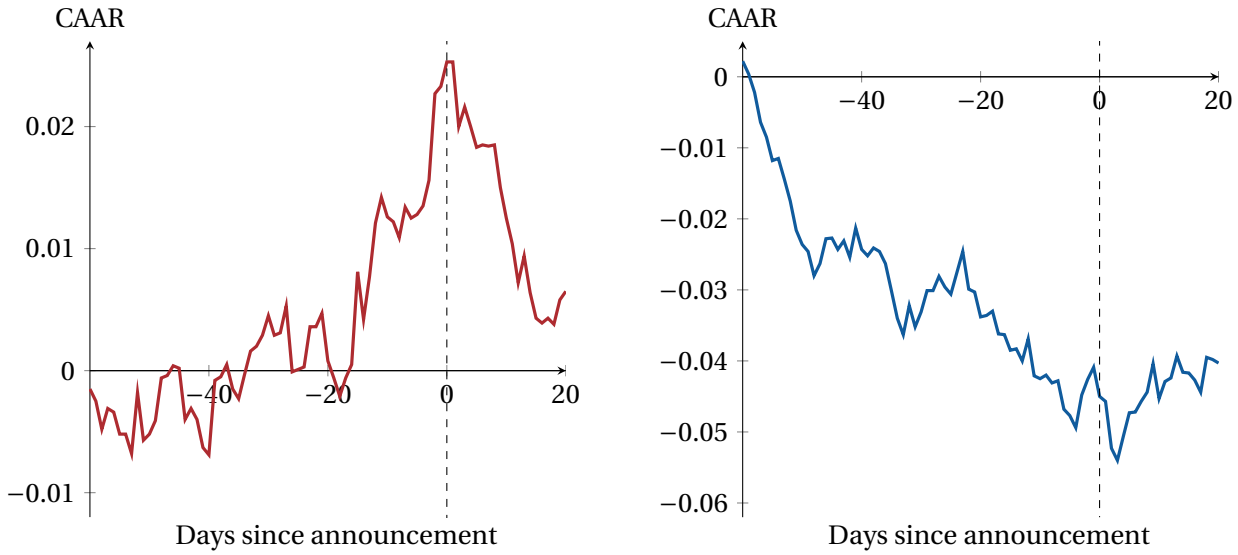

(a) Upgrades

(b) Downgrades

Figure 2: Cumulative average abnormal returns

Economic convergence: Is the catchup slowing down?

Gabriele Amorosi
University of Hull
g.amorosi@hull.ac.uk

Why should we be interested in economic convergence? If one spends some time travelling around the world, one will be very likely to see a tremendous degree of variation in terms of welfare and living standards. The per-capita incomes of some of the world's poorest countries are less than one-tenth the average levels in the wealthiest ones. The same also applies at the national level, as it is certainly not uncommon to observe certain regions of one country becoming richer and richer (at least on average) and others experiencing high levels of backwardness and industrial decline along with high unemployment.

Dr Gabriele Amorosi joined Hull University Business School as a lecturer in economics in November 2011. Prior to that, he was assistant lecturer at the University of Kent where he earned his PhD. Previously he had been a research affiliate of the Department of Economics at the University of Siena where he obtained undergraduate and master's degrees in economics and banking. He was also a Visiting Fellow at the Center for Global Change and Governance, Rutgers University, and research visitor at the Economic and Social Cohesion Laboratory, European Institute, London School of Economics. His main research interests are income distribution, financial development, economic convergence, and regional policy.

These sometimes very striking differences in income levels are, in turn, reflected in nearly every aspect of the quality of life—from the number of televisions and telephones per household to the infant mortality rate and, ultimately, life expectancy. Much research has been devoted to the question of whether economies converge over time to one another in terms of economic development. In particular, do economies that start off poor subsequently grow faster than economies that start off rich? If they do, will the world's poor economies then tend to catch up with the world's rich economies? This property of catching up is what economists call convergence; if convergence does not occur, then poor countries are likely to remain relatively poor.

The concept of covergence

The conventional theories on convergence are typically derived from the neoclassical model of economic growth proposed by Nobel Prize winner Robert Solow in 1956. According to this model, the rate of savings is the main factor promoting the growth of capital stocks in a particular economy and, thereby, determining the so-called steady-state level of growth in per-capita output and income income in the short run—indeed, at the aggregate level, the investment in capital stock within an economy is financed by saving. Assuming a constant labour force, a sustained increase in capital investment increases the amount of 'productive' capital per worker, thus leading to a higher steady-state-level of capital and, in turn, to a higher level of income per person. This concept can help explain the dramatic economic growth recorded in both Germany and Japan in the decade following World War II, during which the two countries experienced a sharp rise in investments—so as to restore their stocks of capital destroyed during the war.

All in all, the Solow model, in this simplified version, is not able to properly explain the phenomenon of persistent growth that is found in the majority of modern economies, especially in the second half of the past century; indeed, for economic growth to be achieved, it is important to implement ideas and new technology. Thus, the Solow growth model was extended by introducing into the model exogenous 'technological change', capable of augmenting factor productivity over time, thereby allowing the model to exhibit long-term economic growth.

The empirical debate on economic convergence is often regarded as being started by William Baumol, with his 1986 paper *Productivity Growth, Convergence and Welfare: What the Long-Run Data Show*. Since then, dozens of researchers have taken up his lead on this issue and have generated a huge strand of literature focusing on how cross-country and cross-

regional economic growth patterns differ. It was, however, not until the seminal work by Barro and Sala-i Martin (1992) that the fundamental concepts of σ (sigma) and β (beta) convergence were introduced.

The idea behind the notion of σ-convergence is quite intuitive; we have σ-convergence when there is a reduction in the dispersion of per-capita incomes over time. In other words, if we use the standard deviation as a measure of dispersion (which, in statistics, is denoted by the Greek letter σ), we have σ-convergence when the standard deviation of per-capita incomes at the end of the period of time considered is less than at the beginning. Beta convergence, which takes its name from the parameter β as set out below, is consistent with the Solow model and refers to the negative relationship between the rate of growth in per-capita income and the relative initial value in a group of economies. In other words, we have β-convergence when the economies with lower 'initial' levels of development (income) are growing faster than the countries with higher levels of income.

A more formal illustration of the concept of β-convergence is as follows. As explained in greater analytical detail in Islam (2003), if

$$\Delta y = \beta(y^* - y),$$

where Δy is the annual change in per-capita income, y^* is a the long-run income level, and y is the current income level of an economy, then β is a parameter that describes the 'speed' at which each economy converges towards the long-term income y^*. Thus, given a certain value for β, the larger the gap $(y^* - y)$, the higher the rate of growth in per-capita income.

Experts typically distinguish between two types of β-convergence: unconditional (or absolute) β-convergence and conditional β-convergence. Unconditional β-convergence occurs when the income gap between two or more countries decreases irrespective of these countries' 'characteristics' (e.g. institutions, policies, technology); this implies that all countries are supposed to have a common long-run per-capita income level y^*. Conditional β-convergence, which is usually regarded as a more plausible outcome, occurs when different countries have different long-run levels of income. This implies that a country is still growing faster the greater the gap between current income and long-term income, however, the full convergence of countries with different long-term income levels will never occur—unless some structural economic changes lead to the emergence of common long-term per-capita income levels. It can be useful at this point to underscore how the two concepts of σ-convergence and β-convergence are connected. In fact, using intuitive logic, it is easy to understand that if the income levels become increasingly similar over time this is most likely due to the fact that the weaker economies are growing at rates higher than those of the stronger economies.

The empirics of economic convergence

Here we provide an illustration of the way economic convergence has traditionally been tested in the literature. Let g_i denote the average annual growth rate of per-capita income in country i over the time period and let y_{i1} denote the initial per-capita income.[1] It is common to estimate the regression model specified below:

$$g_i = a - b\ln y_{i1} + u_i,$$

where ln denotes the natural logarithm. Ignoring the terms a and u_i in the equation above, notice that the annual rate of economic growth between year 1 and year T depends on the initial level of income, and b represents an estimate of how initial income affects growth. Finding the parameter b to be positive is consistent with unconditional (or absolute) β-convergence. Furthermore, it is possible to recover the proper coefficient of convergence β from the regression coefficient b.[2]

In order to test for conditional β-convergence, researchers account for the countries' structural characteristics by adding a number of variables to the equation above, each representing a distinguishing feature for each country.

Many cross-sectional studies applying β-convergence analysis to different sets of (mostly developed) countries have found estimated convergence rates (β) of about 2 per cent per year; this is regarded by the relevant literature as a sort of stylized fact. This result is found in a number of data sets, such as the OECD countries, the US states, the UK counties, the Japanese prefectures, the regions of Europe, and the Canadian provinces, among others; it is also found in data sets that range over time periods from the 1860s through to the late 1990s.

[1] Specifically, $g_i = \ln(y_{iT}/y_{i1})/T$ where T is the length of the sample period in years and y_{iT} is the per-capita income at the end of the sample period.

[2] They are related by the equation $b = (1 - e^{-\beta T})/T$ so that $\beta = -\ln(1 - bT)/T$.

Some new empirical evidence

The small recent empirical analysis reported below encompasses both a large sample of 122 countries (from advanced countries to less developed countries) and a subsample of 42 advanced and newly industrialized countries. Using data taken from the Penn World Table dataset, I have first computed the growth rate of per-capita income for each of the countries considered, and then I have used the model for unconditional convergence as specified in the previous section in order to test for unconditional β-convergence during the period 1980-2007. As outlined above, one should expect to obtain a negative β for the smaller sample of more advanced countries, while the outcome of the regression analysis relative to the large sample could plausibly be expected to be a relatively greater β. Looking at Figure 1, it becomes evident that, largely, economic convergence is not what the data show. Those countries with a higher income in 1980 actually tended to grow slightly *faster* over the period 1980–2007; incomes diverged at a rate of 0.09 per cent a year. Things change when the smaller sample is considered, as illustrated in Figure 2. In this case, there is evidence of economic convergence, taking place at an annual rate of 0.69 per cent. This is consistent with the traditional findings, even though the magnitude of such convergence is far below the 2 percent level claimed by most of the established empirical studies. The reason for such a slowdown is not simple to explain; Rodrik (2011) sees the need of further structural reforms in both mature and developing economies to remove impediments to growth, not to mention the low levels of productivity many Western countries are facing. These and other factors, not least the consequences of the recent financial crisis, will have to be taken into account when thinking of how the future of worldwide economic growth will unfold.

References

Barro, Robert J and Xavier X Sala-i Martin (1992), "Convergence." *Journal of Political Economy*, 100, 223–251.

Baumol, William J (1986), "Productivity Growth, Convergence, and Welfare: What the Long-Run Data Show." *The American Economic Review*, 76, 1072–1085.

Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer (2013), "Penn World Table." URL http://www.ggdc.net/pwt.

Islam, Nazrul (2003), "What have We Learnt from the Convergence Debate?" *Journal of Economic Surveys*, 17, 309–362.

Rodrik, Dani (2011), "The Future of Economic Convergence." Working Paper 17400, National Bureau of Economic Research.

Solow, Robert M (1956), "A Contribution to the Theory of Economic Growth." *The Quarterly Journal of Economics*, 70, 65–94.

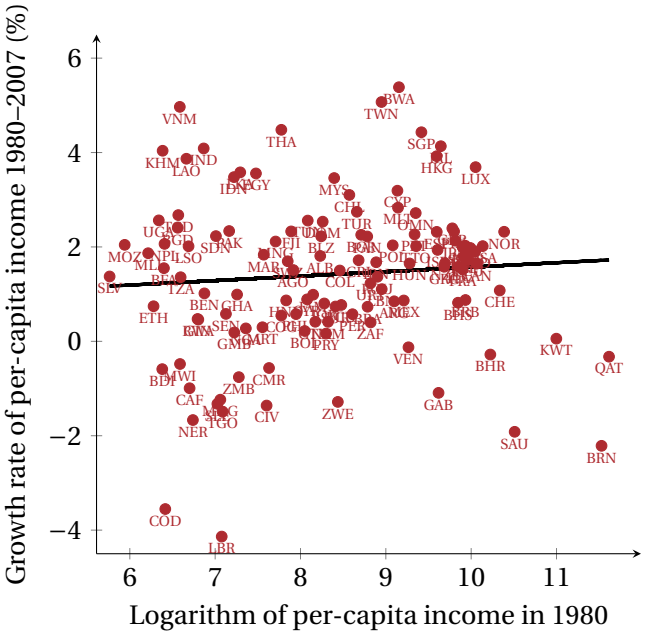

Figure 1: Large sample ($-\beta = 0.09\%$)

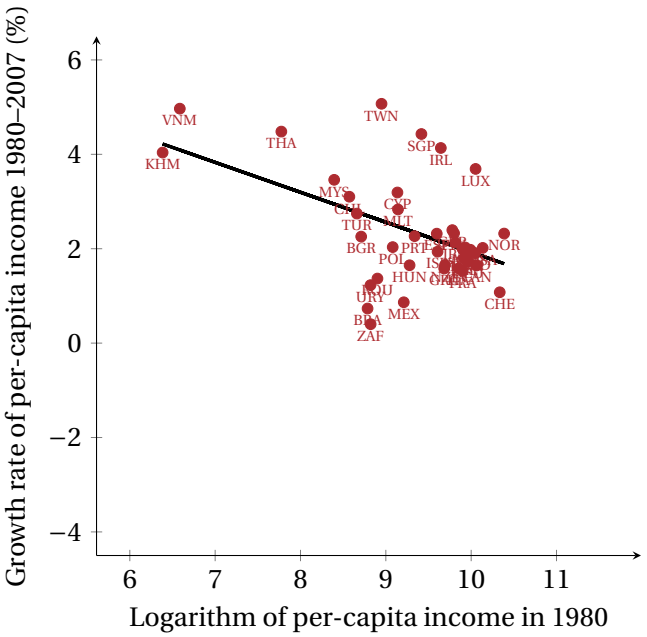

Figure 2: Subsample ($\beta = 0.69\%$)

Market frictions and their implications

Teng Ge
University of Hull
t.ge@hull.ac.uk

Our life is filled with imperfections and puzzles that are far beyond the stylized stories given in the textbooks. The idea that a well-functioning market under the 'invisible hand' clears itself is perhaps one of the most misunderstood arguments of economics by the public (and even by some professionals). The 'hand', i.e. the price of the market, is a strong tool for allocating resources, consummating trades, and distributing gains. Everyone is happy: all buyers and sellers willing to trade at the market-clearing price participate in the market and those who are not willing to trade do not; welfare is maximized.

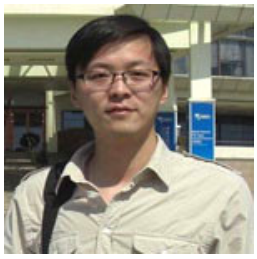

Dr Teng Ge joined Hull University Business School from Hertford College, University of Oxford, where he held the position of Lectureship in Economics from 2011 to 2013 after his PhD from University of Essex. Teng's research is primarily concerned with applied and theoretical microeconomics, with particular interest in the labour market. Current research focuses on how the specific mechanics of the job-matching process affect employment flows, and wage determination.

If this delightful tale were true (or even close to reality), we would not see long queues outside the job centre whilst job advertisements pop up in local newspapers, magazines and shop windows on the high street. If the tale were true, we would not see different prices for Diet Coke in Tesco and Sainsbury's on the same street, just a few hundred yards away from each other. If the tale were true, we would not see single men actively searching for a partner in a market for marriage, while in the same market there are single women looking for a partner. These observations from everyday life are opposed to our theory of a well-functioning market: if a worker is willing to sell their labour at a wage the market offers there must be some employers willing to buy their labour. Otherwise, the wage in the market would drop to such a level that employers are happy to employ all workers willing to work. So why are there unemployed people and unfilled jobs in the same market, at the same time? In the goods market example, if Tesco charges a penny more for Coke than its neighbour, is it not afraid of losing all customers to its competitor? Departing from the stylized textbook market, search theorists like Ken Burdett, Peter Diamond, Dale Mortensen, and Chris Pissarides (the last three won the Nobel Prize in 2010), told beautifully elegant stories about market frictions and how market frictions influence our everyday life.

A 'friction' is a restriction imposed by nature that may make it difficult to achieve good outcomes. Frictions can take a lot of different forms. For example, it takes time to meet a life partner. Consider a dating website—this is an online institution that helps you meet others. This institution helps mitigate the friction of finding a compatible mate by expediting your ability to meet individuals. In a market, there are many different kinds of frictions. In this essay, I focus particularly on the frictions in goods markets. In particular, instead of meeting at some centralized Walrasian market (where the price automatically adjusts to equate supply and demand), buyers need to engage in costly search in order to find sellers and discover their prices.

If you shop at different grocery stores you might notice that some similar goods—sometimes even identical ones—are priced differently in different stores. On the day I was writing this essay, I personally visited the local Tesco for a small can of Diet Coke. On the same street, just a few hundred yards away, there is a Sainsbury's. I found that Tesco asked for 42 pence for a Coke, whereas Sainsbury's charged 45 pence. Actually, if you visit a price comparison website, you will find most of the major dealers are charging differently for a can of Coke. Similar observations can be made in many local stores on the street where I pass by everyday: different prices for fruits and vegetables, different prices for fish and chips, etc. One can also try opening a web browser and visiting Amazon; by just clicking on the price tag you may find many examples of the same item being offered for different

prices by different sellers.

Search theory offers an elegant explanation to the follow question: why do sellers, even with identical costs, charge different prices for the same good, in the same market? The answer is, if the market is frictional then buyers must search for the sellers who charge a lower price. In response, if buyers search, sellers may price differently from each other and we shall see price dispersion in the market. To prove this idea, let us first look at how market frictions change buyers' behaviour. In a frictional world, buyers need to exert some effort to search for partners to trade. For example, I need to walk a little bit longer to visit other stores to find out the price from a different seller or I need to spend more time online to compare prices charged by different e-sellers. The market frictions prohibit buyers from obtaining complete information about the market: in the first case, I physically cannot visit two shops at the same time; in the second case, I can only sample a subset of all the e-sellers online. This small change in the assumptions, based on the fact of costly search by consumers, leads to some interesting implications.

The Diamond paradox

At first, search theorists found that imperfect information, or market friction, per se cannot generate price dispersion. This is the famous 'Diamond paradox', which says that if buyers are homogeneous and sequentially search for sellers in the market, all the sellers will charge the same price—the monopoly price—for that good. To prove this, let us suppose that a buyer wishing to purchase a can of Diet Coke knows only the distribution of prices of Coke, but does not know which seller charges any given price. To find out the price a particular seller charges, the buyer has to search sequentially one store at a time. Specifically, a buyer has a so-called 'reservation strategy': they sequentially visit the stores to find out who charges what, and stop searching and trade with the first seller who charges a price that is no higher than their 'reservation price', denoted by P^*. Intuitively, the reservation price is determined by a buyer's utility from purchasing Coke and the distribution of prices.

Specifically, the reservation price is the highest price at which a buyer is just indifferent between purchasing now and walking away and starting a new search. To see this intuitively, let V^S be the value that a consumer obtains from searching in the market and V^P be the value she obtains from purchasing a can of Coke from a seller charging (any) price P. The higher the price the seller charges, the lower the value a buyer obtains from purchasing. This implies a downward sloping V^P line in (Price, Value) space, as shown in Figure 1. On the other hand, the value from search is independent of the price that a seller charges, so that V^S is a horizontal line. Actually, the value from search depends only on the expected price (in other words the average price) of the market. For prices less than P^*, the value from purchasing is higher than the value of search; at these prices buyers are strictly willing to buy from a seller. For prices greater than P^*, the value from search is above the value of purchasing; buyers will walk away and keep looking for a better deal. There is only one price, i.e. P^*, at which the value from purchasing just equals the value from search. A price which is above this price by even a penny will be immediately rejected by buyers.

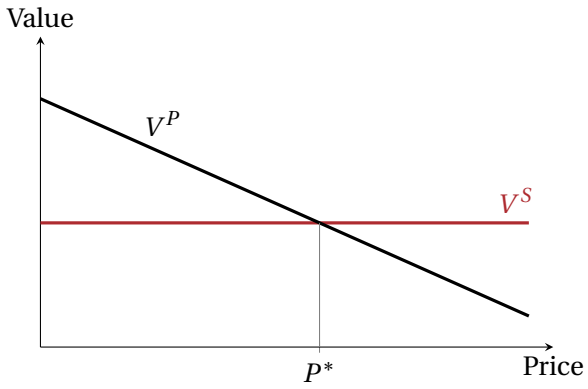

Figure 1: The reservation price (P^*) is the price such that the buyer's value from purchasing (V^P) is equal to the value from continuing to search (V^S).

Now, on the other side of the market, the sellers must respond to buyers' reservation strategies by adjusting their prices. Which price will a profit-maximizing seller charge? The answer is P^*, the reservation price of the buyers. Why? Well, the reason is that buyers will buy from a seller who charges a price less than or equal to their reservation price. Of course, by charging a penny lower than P^*, buyers will be happier but this will not change the fact that a buyer will purchase from this seller. However, the seller's profit is now a penny less per sale than from charging exactly P^*. On the other hand, if the seller charges a penny above P^*, buyers walking into this store will feel the price is unfairly expensive and so will walk away to other stores. The profit from setting a price a penny above P^* is zero. This simple logic proves that the best strategy for the sellers in response to the buyers' reservation strategy is to charge a price of exactly P^*.

Then the implication is quite straightforward: if all

the buyers value a can of Coke the same, then they must have the same reservation price P^*. All the Coke sellers in the market will charge the same price P^*, i.e. the monopoly price. That means, the sellers extract the entire surplus from trade and leave buyers with none. If there is a positive cost of search—*no matter how small*—buyers will stop searching in the market and simply buy from the first seller they encounter. Because all the sellers charge the same price, why bother searching? This is the famous 'Diamond Paradox': if buyers search sequentially, there is still a single price, and—even though there are many sellers in the market—it is the monopoly price!

Breaking the 'law of one price'

The key to breaking the seemingly strong 'law of one price'—although in the frictional world the one price is the monopoly price as opposed to marginal cost in the frictionless world—is to change the assumption of sequential search. Burdett and Judd (1983) played with the textbook models of monopoly and Bertrand competition in the following manner. Suppose that the consumers come to the market for a can of Coke. They know the distribution of prices in the market, but since the market is frictional, consumers have no idea who charges what until they actually meet with a seller. However, in the world of Burdett and Judd, some buyers are lucky enough to observe two prices set by two different sellers, but the others only observe one price from one of the sellers. For example, some people are too busy to visit two shops for a can of Coke, whilst others are happy to shop around.

What will happen in Burdett and Judd's 'noisy search' market? Well, it depends. Those busy buyers who only observe one price are willing to trade with a seller if the seller asked a price no greater than their reservation price. On the other hand, those buyers who observed two prices will obviously buy from the cheaper seller, if the lower price is not above their reservation price. If everyone only observes one price, then the market goes to the monopoly case, i.e. the Diamond Paradox; if everyone observes two prices and only buys from the cheapest sellers, we get Bertrand competition: all the sellers charge a price equal to their marginal cost. But if some buyers observe one price and some observe two, surprisingly, we shall have a market with price dispersion just as we see in the real world!

To prove this logic, we just need two steps: first, to demonstrate that the highest price and lowest price sellers will charge are different; second, to demonstrate that any price between these two will be fine, i.e. an equilibrium. What is the highest price a seller could charge? Obviously, the buyers' reservation price. Any price a seller charges above this level will be turned down. What is the lowest price a seller could charge? The answer is the seller's marginal cost. If a seller prices any lower it will incur a loss. Now, we have two candidate prices: the highest price sellers could charge, i.e. the reservation price, and the lowest price, i.e. the sellers' marginal cost. Which price(s) can survive economic reasoning? Suppose every seller in the market is charging the highest price for their Coke. If I were a seller, should I set my price as high as everyone else or a penny lower? Well, the answer is I should charge a penny lower. This is because although my profit per sale will be a penny less, there are some buyers who observe my price and that of one of the other sellers, and they will all come to me. Therefore, I will sell a lot more units than if I charged the same as everyone else. Hence, it is profitable to undercut! Consider the other extreme case, where every seller in the market charges the lowest price, marginal cost. If I were a seller, should I charge the lowest price or a penny more? Well, the answer is I could charge a penny more! The reason is that although I will lose those customers who observe two prices to the cheaper sellers, there are lazy/busy customers in the market who only visit one seller. If they visited me, they will buy! Therefore, my profit per sale is now a penny more, and hence I make a positive profit. Continuing with this logic, we see that the market price must be dispersed: there cannot exist two sellers charging the same price. In equilibrium, each firm charges some price below the reservation price and above marginal cost and all of these prices yield the same profit.

This elegant logic in Burdett and Judd (1983) answered the question of why the same goods are priced differently by different sellers. Burdett and Mortensen (1998) applied a similar idea to the labour market and predicted equally productive workers and firms will result in dispersed wage payments. Similar workers—meaning similar education, age, sex—could be paid substantially differently, just because of luck in job search. Empirically, economists (Christensen et al., 2005; Rosholm and Svarer, 2004) find that at least one third of income inequality is explained by market frictions. Although this finding implies some of the differences in wage earnings are outcomes of pure luck, it argues strongly for redistributive taxation. However, this is still the research frontier, as most of the literature about optimal taxation assumes a frictionless market.

References

Burdett, Kenneth and Kenneth L. Judd (1983), "Equilibrium Price Dispersion." *Econometrica*, 51, 955–969.

Burdett, Kenneth and Dale T. Mortensen (1998), "Wage Differentials, Employer Size, and Unemployment." *International Economic Review*, 39, 257–273.

Christensen, Bent Jesper, Rasmus Lentz, Dale T. Mortensen, George R. Neumann, and Axel Werwatz (2005), "On-the-Job Search and the Wage Distribution." *Journal of Labor Economics*, 23, 31–58.

Rosholm, Michael and Michael Svarer (2004), "Endogenous Wage dispersion in a Search-matching model." *Labour Economics*, 11, 623–645.