BANK OF CANADA
BANQUE DU CANADA

Staff Discussion Paper/Document d’analyse du personnel 2019-3

Global Commodity Markets and Rebalancing in China: The Case of Copper

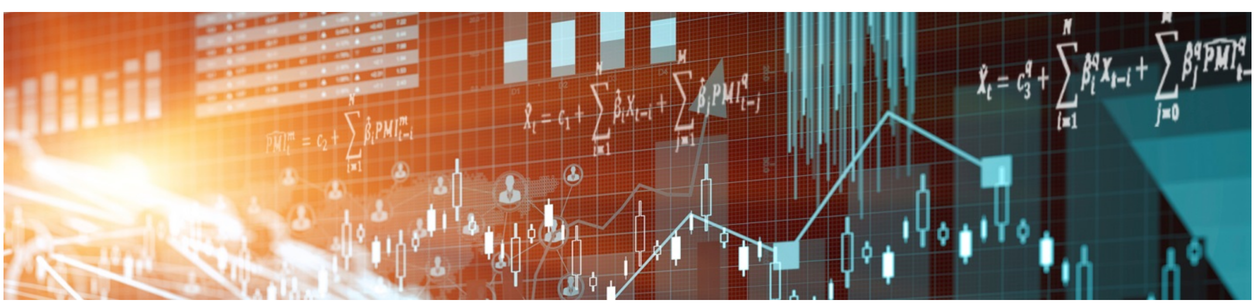

by Jeannine Bailliu, Doga Bilgin, Kun Mo, Kurt Niquidet and Benjamin Sawatzky

www.bank-banque-canada.ca

Bank of Canada Staff Discussion Paper 2019-3

April 2019

Global Commodity Markets and Rebalancing in China: The Case of Copper

by

Jeannine Bailliu, Doga Bilgin, Kun Mo, Kurt Niquidet and Benjamin Sawatzky

International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
jbailliu@bankofcanada.ca
kmo@bankofcanada.ca
kniquidet@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

The authors would like to thank Rose Cunningham, Christopher Hajzler, Mark Kruger and Konrad Zmitrowicz for helpful comments, as well as participants at seminars at the Bank of Canada. The views expressed in this paper are solely those of the authors.

Abstract

Given that China accounts for about half of global copper consumption, it is reasonable to expect that any significant change in Chinese copper consumption will have an impact on the global market. This paper examines the likely impact of the rebalancing of the Chinese economy on its copper consumption over the next decade, focusing on the relationship between the copper intensity of GDP and the share of investment in GDP. We use a panel smooth transition regression model to account for potential non-linearities in this relationship at different levels of urbanization and income. Our findings suggest that there is indeed a significant relationship between a country's copper intensity of GDP and its investment share. Our baseline rebalancing scenario for China implies that copper intensity in China has already peaked and is expected to decline steadily through the next decade. This anticipated reduction in Chinese copper intensity is the result of the dampening impact of rebalancing and higher per capita income on copper intensity, which more than offsets the upward pressure stemming from the ongoing process of urbanization. An exploration of alternative rebalancing scenarios suggests that China's rebalancing path could have a significant impact on global copper consumption.

Bank topics: International topics; Econometric and statistical methods
JEL codes: O13, O14, Q02

Résumé

Étant donné que la Chine compte pour environ la moitié de la consommation mondiale de cuivre, il est raisonnable de penser qu'une variation importante de la consommation chinoise aura un impact sur le marché mondial de ce métal. Nous examinons les répercussions probables du rééquilibrage de l'économie chinoise sur la consommation intérieure de cuivre dans les dix prochaines années et nous intéressons à cette fin au lien entre l'intensité-cuivre du PIB et la part de l'investissement dans le PIB. Pour tenir compte d'éventuelles non-linéarités dans cette relation pour différents degrés d'urbanisation et niveaux de revenu, nous recourons à un modèle de panel à seuil à transition lisse. Nos résultats tendent à montrer qu'il existe bel et bien une relation significative entre l'intensité-cuivre et la composante investissement du PIB d'un pays. Il ressort de notre scénario de rééquilibrage de référence pour la Chine que l'intensité-cuivre y a déjà culminé et devrait être en diminution constante dans les dix années à venir. Cette baisse prévue en Chine est attribuable à l'incidence modératrice du rééquilibrage et à la hausse du revenu par habitant, qui fait plus que compenser la pression haussière exercée par le processus d'urbanisation en cours. Il se dégage des autres scénarios envisagés pour la Chine que la trajectoire de rééquilibrage pourrait avoir un effet marqué sur la consommation mondiale de cuivre.

Sujets : Questions internationales; Méthodes économétriques et statistiques

Codes JEL : O13, O14, Q02

1 Introduction

China is currently the most important player in the global base metals market, accounting for roughly 50 per cent of base metal consumption in 2014, growing from about 20 per cent in 2001 (Figure 1). China's GDP is very metal intensive, reflecting the fact that its output is weighted toward metal-intensive sectors, such as housing and infrastructure. However, Chinese consumption for some commodities may wane as China's economy transitions to a more balanced and slower growth path characterized by a rotation away from investment and heavy industry and toward private consumption and services. The term "rebalancing" is often used to describe this transition that is underway in the Chinese economy—one of its main features is a reduction in the investment share of GDP from its currently high level.

In this paper, we examine how the rebalancing of China's economy will likely affect its copper consumption over the next decade. We focus on China given that it accounts for roughly half of global copper consumption. Thus, it is reasonable to expect that any significant change in Chinese copper consumption will have an impact on the global market. An existing literature based on what is known as the Environmental Kuznets Curve (EKC) suggests that the relationship between metal intensity and output can break down at higher levels of income, due to some combination of changing consumer preferences and economic structure. To capture these changes and quantify the impact of Chinese rebalancing, we focus on the relationship between a country's investment share of GDP and its ratio of copper consumption to GDP (henceforth, copper intensity), as investment tends to be a major source of copper demand. However, there are reasons to believe that this relationship may not be stable over time due to the changing composition of investment.

To address potential non-linearities between copper intensity and the investment share of GDP, we apply a panel smooth transition regression (PSTR) model on a panel of countries using data over the period from 1950 to 2015. We use the urbanization rate and level of real GDP per capita as transition variables to determine regimes in the PSTR. These transition variables are intended to account for the changing composition of investment. Based on these results, we then analyze the impact of different rebalancing outcomes on Chinese copper consumption. We construct our baseline scenario drawing on the early experience of the Japanese economy, which underwent a similar period of rebalancing characterized by a decline in its high investment-to-GDP ratio.

Our paper distinguishes itself from existing empirical work on the impact of the Chinese economy on commodity markets in that we emphasize the link between the rebalancing of the Chinese economy and copper consumption. Studies such as Roberts et al. (2016), Roache (2012), McKay et al. (2010) and Ghoshray & Pundit (2016) examine the impact of a slowdown in Chinese economic activity on world commodity markets but do not examine the implications from a rebalancing of the Chinese economy. In addition to developing a baseline rebalancing scenario, we also consider alternative scenarios under which rebalancing is either faster or slower relative to the baseline.

Our paper also distinguishes itself from existing empirical work examining potential non-linearities in the relationship between commodity consumption and output by more deeply assessing the non-linearities in the relationship. While other papers have used the level of GDP per capita as a proxy for factors, including the changing structure of an economy, our use of the investment share of GDP as an explanatory variable more directly accounts for the effect of economic composition on copper consumption. Our use of real GDP per capita and the urbanization rate as transition variables can then account for further non-linearities that are not captured by the aggregate investment ratio.

Finally, we also contribute to the literature through a more rigorous econometric analysis of the relationship between a country's copper intensity and its key determinants. Many earlier papers (e.g., Döhrn & Krätschell (2013); Roberts et al. (2016); Warell & Olsson (2009)) use the levels of commodity consumption and GDP per capita as dependent and independent variables in either quadratic panels or PSTR specifications. In either situation, both variables are non-stationary, which raises concerns that the results may be spurious. Our use of stationary variables like copper intensity and the investment share in the PSTR framework allows us to provide a more robust assessment.

The rest of the paper is structured as follows. Section 2 provides an overview of the drivers of global copper consumption. In Section 3, we discuss our methodology for estimating the relationship between copper intensity and its determinants. In Section 4, we present our estimation results. Section 5 reviews the different rebalancing outcomes that we consider and the implications of each for Chinese copper consumption over the next decade. We extend this analysis in Section 6 by exploring potential structural changes in the world copper market that may offset somewhat the moderation in Chinese copper consumption. We offer some concluding remarks in Section 7.

2 Drivers of copper consumption

Copper is one of the world's most traded and widely used commodities. In 2016, refined copper consumption was 23.5 million tonnes with approximately 50 per cent of consumption attributed to China. Copper is traded on several exchanges worldwide, with the London Metals Exchange historically being most important. However, as measured by open interest on copper futures contracts, the Shanghai Futures Exchange has seen more use in recent years, reflecting China's increasingly dominant role in the global market. Over the past 15 years, copper prices have fluctuated between a range of US$1 to $5 (in constant 2018 US dollars) per pound (Figure 2). It is generally known and accepted that economic growth in China was a key driver of increasing metal consumption, including copper, and responsible for a significant run-up in prices between 2004 and 2011 (Kruger et al. (2016)).

However, copper intensity is not expected to be stable through a country's economic development process. The typical pattern involves copper intensity rising at the early stage of economic development, when an agrarian economy industrializes and the manufacturing and construction sectors become the main drivers of often rapid economic growth. This situation characterizes China over the last couple of decades. However, as the economy continues to develop, the relative importance of investment and manufacturing shrinks at the expense of consumption and services, and thus copper intensity typically falls.

Previous studies have generally modelled a non-linear relationship between metal consumption and the level of real GDP per capita, focusing on the income level at which metal consumption reaches the inflection point (e.g., Warell & Olsson (2009); McKay et al. (2010);

Döhrn & Krätschell (2013)). This is justified in part by the idea that the structure of the economy changes at higher income levels and consumer preferences shift as households become wealthier. In addition to potential spurious results from using the non-stationary level of GDP as a regressor, this approach is also inefficient, because the true driver of the non-linearity is left unspecified. We instead approach the empirical analysis by examining the relationship between copper intensity and the underlying economic structure directly. Specifically, from the perspective of the expenditure approach to GDP accounting, we expect to observe a positive relationship between copper intensity and the share of GDP devoted to gross capital formation, as investment in infrastructure, housing and capital goods are copper-intensive GDP expenditure components.

Even after accounting for changes in the high-level economic structure, there may still be further important non-linearities in the data arising from the changing nature of investment. As an economy matures, the importance of research and development in investment can be expected to increase at the expense of investment in durable and capital goods. These secondary transformations within the sectors can be expected to cause non-linearity in the relationships between copper intensity and the share of investment in GDP.

Another potential source of non-linearity arises from the fact that the composition of investment could be different for different levels of urbanization. In general, urban centres require greater investment in transportation and housing infrastructure compared with rural areas. For example, several studies have linked urbanization to increased electricity consumption (Jones (1991); Zhao & Zhang (2018)). These types of infrastructure investment typically require large amount of base metals, including copper (Kruger et al. (2016)). Hence, for a given investment share of GDP, we hypothesize that copper intensity would be different depending on the urbanization rate.

3 Methodology

3.1 The PSTR model

To assess the impact of Chinese rebalancing on copper consumption, we use a framework that accounts for potential non-linearities in the relationship between economic development and commodity consumption. The EKC literature suggests the presence of such a non-linearity, with findings that the relationship between metal consumption and output breaks down at higher levels of income. Traditionally, the EKC literature as applied to commodity consumption has used polynomial models, which consist of estimating a quadratic equation that includes GDP per capita and GDP per capita squared using a standard fixed-effects model (e.g., Canas et al. (2003); Warell & Olsson (2009); McKay et al. (2010); Cole & McCoskey (2013); Döhrn & Krätschell (2013)).

However, other studies have used PSTR models to capture non-linear relationships across regimes (Destais et al. (2007); Aslanidis & Xepapadeas (2008); Duarte et al. (2013); Roberts & Rush (2010)), which is an extension of the panel threshold regression model (Hansen (1999)) that was first applied by González et al. (2005). These models have the advantage of allowing for a smooth transition between regimes, and also of calculating the threshold within the model rather than giving it a priori. As such, we estimate the relationship between copper consumption and its key determinants using a PSTR model on a panel of 25 countries, both advanced and emerging economies, over the period 1970–2015. Appendix B provides a list of sample countries.

The basic PSTR model with two regimes is defined as follows:

$$y_{it} = \mu_i + \beta_0' x_{it} + \beta_1' x_{it} g(q_{it}; \gamma, c) + \beta_3' trend + \upsilon_{it}, \tag{1}$$

for $i = 1, ..., N$, and $t = 1, ..., T$, where N and T denote the cross-section and time dimensions of the panel, respectively; y_{it} is the dependent variable; μ_i represents the country fixed effect; x_{it} is a k-dimensional vector of exogenous variables; and υ_{it} are the errors. The transition function ($g(q_{it}; \gamma, c)$) is a continuous function of the transition variable (q_{it}) and is normalized to be bounded between 0 and 1. We follow González et al. (2005) and others by using the following logistic specification for the transition function:

$$g(q_{it}; \gamma, c) = (1 + exp(-\gamma(q_{it} - c)))^{-1}, \tag{2}$$

where γ is a slope parameter that determines the smoothness of the transition from one regime to another. The parameter $c = (c_1, ..., c_m)'$ is an m-dimensional vector of location parameters for the transition variable that determine the model's regimes. The restrictions $\gamma > 0$ and $0 < c_1 < ... < c_m$ are imposed as identifying restrictions.[1] The equation is estimated using non-linear least squares. The non-linear least squares algorithm is initialized using starting values for γ and c obtained from a grid search.[2]

It is also worth noting that the PSTR model can be seen as a more general specification of the quadratic polynomial model (which has also been used to study non-linearities in the EKC framework). As such, Aslanidis & Xepapadeas (2008) state that using regime-switching models like the PSTR is justified in part by the fact that these polynomial models are simply particular cases of the PSTR.

3.2 Estimation and testing

Prior to estimating the PSTR, we first need to determine whether there is in fact a non-linearity in the relationship between the variables. To do this, we test whether the regime-switching effect is statistically significant. This can be done for equation 1 by testing either $H_0 : \gamma = 0$ or $H_0 : \beta_1 = 0$. However, either of these tests will be non-standard for equation 1, since under either null hypothesis the PSTR model contains unidentified nuisance parameters (see González et al. (2005) for more details). The problem of hypothesis testing in the presence of such parameters is sometimes known as the Davies problem (Davies (1977, 1987)), and we follow the approach of González et al. (2005) and Luukkonen et al. (1988) by testing for linearity using the null hypothesis $H_0 : \gamma = 0$. To address the identification problem,

[1] In practice it is usually sufficient to consider $m = 1$ or $m = 2$, as these values allow for commonly encountered types of variation in the parameters. For $m = 1$ the model implies that the two extreme regimes are associated with low and high values of q_{it}, while for $m = 2$ the transition function attains the value 1 both at low and high values of q_{it} (see González et al. (2005) for more details).

[2] We thank Christophe Hurlin for sharing the code used to estimate the PSTR model.

we replace $g(q_{it};\gamma,c)$ in equation 1 with its first-order Taylor expansion around $\gamma = 0$. This leads to the auxiliary regression:

$$y_{it} = \mu_i + \beta_0'^{*} x_{it} + \beta_1'^{*} x_{it} q_{it} + ... + \beta_m'^{*} x_{it} q_{it}^m + \upsilon_{it}^*, \tag{3}$$

where the parameter vectors $\beta_1^*, ..., \beta_m^*$ are multiples of γ and $\upsilon_{it}^* = \upsilon_{it} + R_m \beta_1 x_{it}$ (and R_m is the remainder of the Taylor expansion). Therefore, testing $H_0 : \gamma = 0$ in equation 1 is equivalent to testing the null hypothesis $H_0^* : \beta_1^* = ... = \beta_m^* = 0$ in equation 3. Denoting SSR_0 as the panel sum of squared residuals under H_0 (linear panel model with individual effects) and SSR_1 the panel sum of squared residuals under H_1 (PSTR model with multiple regimes), we can compute the following two test statistics (Colletaz & Hurlin (2006)):

$$LM = TN \frac{SSR_0 - SSR_1}{SSR_0} \tag{4}$$

$$LM_F = \frac{\frac{SSR_0 - SSR_1}{K_m}}{\frac{SSR_0}{TN - N - mK}} \tag{5}$$

This test can also be used for determining the appropriate order m of the logistic function in equation 2. Granger & Terasvirta (1993) and Terasvirta (1994) propose a sequence of tests for choosing between $m = 1$ and $m = 2$ based on the auxiliary regression. Starting with $m = 3$, test the null hypothesis $H_0^* : \beta_3^* = \beta_2^* = \beta_1^* = 0$. If it is rejected, test $H_{03}^* : \beta_3^* = 0$, $H_{02}^* : \beta_2^* = 0 | \beta_3^* = 0$, and $H_{01}^* : \beta_1^* = 0 | \beta_3^* = \beta_2^* = 0$. Select $m = 2$ if the rejection of H_{02}^* is the strongest one; otherwise select $m = 1$. For more details, see Terasvirta (1994) and González et al. (2005).

Once we have tested for non-linearity, we can then estimate the PSTR in equation 1 by applying the fixed-effects estimator and non-linear least squares (NLS). We first eliminate the individual effects, μ_i, by removing individual-specific means and then apply NLS to the transformed data.

While removing individual means is standard in linear models, it requires more care in the context of a PSTR model. Equation 1 can be rewritten as follows:

$$y_{it} = \mu_i + \beta' x_{it}(\gamma, c) + \upsilon_{it}, \tag{6}$$

where $x_{it}(\gamma, c) = (x_{it}, x_{it} g(q_{it}; \gamma, c))$ and $\beta = (\beta_0, \beta_1)$. Subtracting individual means from equation 6 yields the following:

$$\tilde{y_{it}} = \beta' \tilde{x_{it}}(\gamma, c) + \tilde{\upsilon_{it}}, \tag{7}$$

where $\tilde{y_{it}} = y_{it} - \bar{y_i}$, $\tilde{x_{it}}(\gamma, c) = (x_{it} - \bar{x}_i, x_{it} g(q_{it}; \gamma, c) - \bar{w}_i(\gamma, c))$, $\tilde{u_{it}} = u_{it} - \bar{u}_i$, and $\bar{y}_i, \bar{x}_i, \bar{w}_i, \bar{u}_i$ are individual means, with $\bar{w}_i(\gamma, c) \equiv T^{-1} \sum_{t=1} T x_{it} g(q_{it}; \gamma, c)$. Consequently, the transformed vector $\tilde{x_{it}}(\gamma, c)$ in equation 7 depends on γ and c through both the levels and individual means. As such, $\tilde{x_{it}}(\gamma, c)$ needs to be recomputed at each iteration in the NLS optimization.

To estimate the parameters after subtracting individual means, we apply an NLS framework. First, we use a grid search to obtain an initial (γ,c). Given these values, the βs are estimated by ordinary least squares conditioned on the values of γ and c. In the second step, the parameters γ and c are re-estimated using NLS (see González et al. (2005) and Colletaz & Hurlin (2006) for more details).

After estimating the model, we also need to test the number of transition functions that should be included in the specification. Following González et al. (2005), we do this in a similar manner to the test of linearity but consider an alternative model with an additional transition function. In its simplest form, we assess:

$$y_{it} = \mu_i + \beta_0' x_{it} + \beta_1' x_{it} g_1(q_{it}^{(1)}; \gamma_1, c_1) + \beta_2' x_{it} g_2(q_{it}^{(2)}; \gamma_2, c_2) + v_{it}, \tag{8}$$

where the transition variables $q_{it}^{(1)}$ and $q_{it}^{(2)}$ can but need not be the same. This allows us to test for additional non-linearity not only of our first transition variable but also for a different transition variable. The null hypothesis of no remaining heterogeneity in an estimated two-regime PSTR model is formulated as $H_0 : \gamma_2 = 0$ in equation 8. This test is once again complicated by unidentified nuisance parameters under the null hypothesis. As such, we again replace $g_2(q_{it}^{(2)}; \gamma_2, c_2)$ by a Taylor expansion around $\gamma_2 = 0$. This leads to the auxiliary regression:

$$y_{it} = \mu_i + \beta_0^{*'} x_{it} + \beta_1' x_{it} g_1(q_{it}^{(1)}; \hat{\gamma_1}, \hat{c_1}) + \beta_{21}^{*'} x_{it} q_{it}^{(2)} + ... + \beta_{2m}^{*'} x_{it} q_{it}^{(2)m} + v_{it}^*, \tag{9}$$

where $\hat{\gamma_1}$ and $\hat{c_1}$ are estimates under the null hypothesis. The hypothesis of no remaining heterogeneity can then be restated as $H_0^* : \beta_{21} = ... = \beta_{2m}^* = 0$. We can then compute the test statistics defined in equations 4 and 5 (see González et al. (2005) and Colletaz & Hurlin (2006) for more details). This process continues until we reject the addition of a further regime.

4 Estimation results

4.1 Pre-estimation tests and final specification

The dependent variable of interest in the estimation is copper intensity, defined as kilograms of copper consumed per million units of GDP in constant 2011 US dollars. The relationship we seek to model is between copper intensity and the share of gross capital formation in GDP. To help estimate this relationship, we also employ two transition variables: the urbanization rate and the level of real GDP per capita. There are no statistical tests for choosing the optimal transition variables; our choices are guided by the assumption that these two variables are key measures of economic development. Our priors are that the elasticity of copper intensity to the investment share should increase as a country becomes more urbanized but should decrease as the economy becomes richer. The real price of copper and a time trend are included as control variables. Appendix B provides a list of variables and corresponding data sources. Summary statistics are shown in Table 1.

Both the dependent variable and the explanatory variables need to be stationary to avoid spurious regression results. Given the likely cross-sectional dependence of our data, we proceed with second-generation panel unit root tests, specifically using the modified cross-sectionally augmented Dickey-Fuller (CADF) and the cross-sectionally augmented Im, Pesaran and Shin (CIPS) panel test statistics as proposed by Pesaran (2007). These test results reject the presence of a unit root at the 5 per cent confidence level for all relevant variables. In contrast, the transition variables do not need to be stationary in the PSTR specification.[3]

We also need to determine whether there is in fact non-linearity in the relationship between the variables. As discussed in Section 3, we do this by comparing the sum of squared residuals (SSR) from a linear specification against those from a specification where the regressor interacts with the transition variables via a first-order Taylor expansion around $\gamma = 0$. We compute the non-linearity test for specifications involving just a single transition variable (either the urbanization rate or real GDP per capita) and a specification with both transition variables entering simultaneously. Table 2 shows that the linearity hypothesis is rejected for the specification with the urbanization rate as a transition but not the one with the real GDP per capita. However, we also calculate a second set of test statistics to determine whether non-linearities in the data are better captured by including both transition variables simultaneously. Indeed, Table 2 shows that specifications with both transition variables generate significantly lower SSR compared with specifications with the single transition variable. Similarly, we test for the number of thresholds for each transition variable and find a single threshold to be optimal for both transition variables in our estimation. As a robustness check, we also explore model specifications where the urbanization rate and real GDP per capita are used as simple regressors rather than as transition variables. We find that while they remain statistically significant, the model fit is worse compared with specifications using them as transition variables. Taken together, these results confirm that a regime-switching approach is appropriate to model the data and that the threshold effects can be determined by a combination of the urbanization rate and real GDP per capita.

Given the pre-estimation tests, for this paper we use a PSTR specification with the investment share of GDP as an explanatory variable and the urbanization rate and real GDP per capita as transition variables. We also use the price of copper and a time trend as control variables. The time trend is included to capture the effects of technological progress, which would be expected to increase efficiency and weigh on copper demand over time.

4.2 Parameter estimates

Figure 3 displays the estimated transition function, with each dot in the graph representing an observation in our dataset. For the model with the investment share of GDP, we observe a gradual transition across the urbanization rate. In contrast, the transition function based on real GDP per capita is less smooth, with the transition taking place at a per capita income level of between roughly US$10,000 to US$25,000.

Table 3 shows the estimated coefficients in the PSTR model. The estimated coefficients β_1, β_2 and β_3 cannot be simply summed and interpreted as the generalized or average

[3] This is a relevant issue for our estimation since one of the transition variables—the level of real GDP per capita—is determined to be non-stationary. However, this variable is transformed via the transition function into a value bound between 0 and 1, which is then interacted with the dependent variable. The resulting series that enters the final OLS regression is determined to be stationary, and hence the non-stationarity of the original real GDP per capita series does not pose issues for the estimation.

elasticity of copper intensity to the investment share of GDP, since the transition function evolves over time and between countries. Instead, the signs of these estimated coefficients are of interest. The estimated coefficient β_2 represents the regressor interacting with the transition function based on the urbanization rate. The positive sign on this coefficient implies that as a country becomes relatively more urbanized, the elasticity of copper intensity increases. The estimated coefficient β_3 is negative, implying that as a country becomes richer, the elasticity of copper intensity to the investment share of GDP decreases. This result confirms our prior that developed countries are on average more likely to engage in higher value-added and therefore less copper-intensive activities, such as research and development.

A point estimate for the estimated elasticity of copper intensity to the investment share of GDP can be computed for each period. Figure 4 shows the estimated elasticity between 1970 and 2015 for the 25 countries in our sample based on their actual urbanization rate and real GDP per capita. There is a wide variation between countries at each point in time. Moreover, we also observe that elasticities can change significantly over time for emerging-market economies, since these economies are undergoing fundamental economic transformations. In contrast, elasticities for developed countries are more stable over time.

Coefficients for the control variables are in line with priors. We find a negative relationship between copper intensity and the real price for copper. This negative relationship captures potential substitution between copper and other metals where possible (i.e., aluminum and stainless steel, depending on the application) in response to relative price movements. The estimated coefficient for the time trend is also negative, which can be interpreted as technological advances that increase the efficiency of copper use over time.

The estimated fixed effects for China are about three times larger than for most other countries, suggesting other China-specific factors also boosted its copper intensity over the estimation period. We offer two plausible explanations for these unusually high fixed effects. First, infrastructure investment in China has been unusually high even after controlling for its urbanization process and rapid economic growth. The focus on building up the infrastructure stock by all levels of Chinese government over the past few decades is well documented (Sahoo et al. (2010)). As examples, China now has the longest high-speed rail network in the world, and it has also built up an electricity grid that provides access to almost all of its population. These efforts are very copper intensive in nature. Second, processing trade in China has been prevalent, including the assembly of computers, electronic devices and telecommunication equipment, which require copper parts. It is therefore possible that a portion of the copper consumption in China reflects use in goods that are eventually exported and not consumed in China.

5 Different rebalancing outcomes in China and their likely impact on the global copper market

Following a simple Cobb-Douglas production framework, China's previous economic growth model has been characterized by a heavy reliance on inputs. It benefited from an abundance of relatively cheap labour during the earlier stages of its rapid economic development. The buildup of the capital stock from a low starting level also contributed strongly to growth, especially over the past decade. However, a shrinking labour supply and rapid capital accumulation have raised financial stability and efficiency concerns. It now appears necessary for China to transition away from its traditional growth model. A sustainable growth model going forward involves a greater contribution from total factor productivity growth. Chinese authorities recognize this and have been pursuing policies that encourage such shifts to drive future growth.

An important outcome of this transition in the Chinese economy would be a rebalancing away from investment and toward consumption and services. The shift away from investment-led growth has important implications for Chinese copper intensity. In this section, we outline three possible rebalancing scenarios for the Chinese economy over the next 15 years and explore the path for Chinese copper intensity associated with each scenario.

5.1 Baseline scenario

We draw on Bailliu et al. (2017) and construct a bottom-up outlook for China's potential GDP growth through 2030 (Figure 5). This exercise helps to derive two key assumptions required by our PSTR model to compute Chinese copper intensity going forward.

- *GDP per capita*
  In our baseline scenario we assume that China's GDP growth is roughly in line with the potential growth presented in Figure 5, slowing from 6.9 per cent in 2017 to 4.7 per cent by 2030. The growth profile embedded in our baseline assumption is comparable to current medium- and long-term growth forecasts from other forecasters (e.g., International Monetary Fund, World Bank and Organisation for Economic Co-operation and Development). With still robust growth through 2030, China's per capita GDP doubles from around US$15,000 to US$30,000 (constant 2011 US dollars) over this period.

- *Investment/GDP ratio (pace of rebalancing)*
  The profile in Figure 5 embeds the view of substantial changes in the drivers of economic growth in China, which helps inform our assumption on the pace of rebalancing. Importantly, the contribution to growth from capital accumulation nearly shrinks in half from 4.9 percentage points in the current decade to just 2.7 percentage points over the next decade. The slowdown in capital accumulation implies that the investment share of GDP in China will fall from around 45 per cent currently to around 30 per cent by the end of our scenario horizon (Figure 6). This baseline assumption for China's experience as it rebalances its economy is comparable to that of Japan as its economy transitioned away from a very high investment rate in the 1970s (Figure 7).

- *Other assumptions*
  To complete our scenario, we need to make assumptions regarding the paths of population growth and the urbanization rate in China through 2030. These were obtained from the United Nations Department of Economic and Social Affairs (2017) and United

Nations Department of Economic and Social Affairs (2018), respectively. Table 4 summarizes all key assumptions.[4]

Based on these assumptions, our baseline scenario suggests that Chinese copper intensity has already peaked and is expected to decline steadily through the next decade (Figure 8). Within the framework of our model, the elasticity of copper intensity to the investment share remains fairly constant at around 0.5 over the scenario period, since the urbanization-driven pickup is offset by downward pressure stemming from the higher level of per capita income. In this context, the 15-percentage-point decline in the investment share results in a substantial drop in China's copper intensity. Intuitively, our model helps to identify the major forces that are expected to help shape China's copper intensity through the coming decades (Table 5):

- *Rebalancing*
  We find that China's transition away from investment should reduce copper intensity by about 100 kilograms per US$1 million of GDP through 2030. This reduction reflects the fact that other segments of economic activity such as consumption and services are less copper intensive than investment in fixed assets.

- *Economic development*
  China's still robust GDP growth will likely reduce its copper intensity over the next decade. China is expected to move through the steepest part of the transition function for per capita GDP, and by 2030 the model would sort China into the relatively high-income regime. More precisely, per capita income in China in 2030 would be around 85 per cent of the current per capita income in Japan. As China becomes richer, its elasticity of copper intensity to the investment share of GDP is expected to fall.

- *Urbanization*
  Based on official statistics, China's current urbanization rate is only around 58 per cent. This level is below most other countries even at comparable stages of economic development, with urbanization confined mainly to coastal areas. Going forward, the national urbanization rate should continue to rise as poorer inland areas become more urbanized, pushing up the copper intensity for a given investment share of GDP. In line with United Nations Department of Economic and Social Affairs (2018), the Chinese urbanization rate is assumed to rise to 70 per cent by 2030. Rising urbanization will provide an important offset to other forces pushing down copper intensity over the next decade.

- *Efficiency gains*
  We assume that the time trend proceeds linearly over the scenario horizon, implying constant gains in efficiency that reduce the amount of copper needed for a given level of output.

[4] For our control variables, we assume that our time trend proceeds linearly over the scenario horizon and that prices remain flat.

Overall, our baseline scenario suggests that China's copper intensity has already peaked. We find that copper intensity is predicted to decline from 638 to 393 kilograms per US$1 million of GDP in constant US dollars. However, since GDP growth remains robust, this path of copper intensity still implies a sizable increase in China's expected copper consumption volume through 2030. Specifically, our baseline scenario suggests copper demand will increase by about 4 million tonnes during the next decade (Figure 9). While the increase in our baseline scenario is lower compared with observed increases in the past, it nonetheless is large relative to current global consumption (roughly 20 million tonnes).

It is important to note that our baseline outlook is subject to some key caveats. In particular, our model does not capture the potential endogenous price response to changes in China's copper intensity—our baseline scenario includes the assumption that prices remain flat.[5] Presumably, the still large increases in Chinese consumption would put upward pressure on prices, which could in turn dampen the increase in copper intensity by spurring substitution to other metals (where possible) and by encouraging greater efficiency in copper use. In addition, as discussed in Section 4, the larger fixed effects for China could be interpreted as idiosyncratic features in China that result in the higher copper intensity. While these factors may be relevant over history, they may be less important in the future, which would weigh on China's copper intensity relative to our scenario results.

5.2 Alternative rebalancing scenarios

As discussed in the previous section, China's rebalancing path will have a significant impact on its copper intensity. As such, it is worth exploring the potential effects of faster or slower rebalancing scenarios on China's copper intensity, holding all else equal as in the baseline scenario.

We consider first the case of a more rapid rebalancing scenario. In particular, we assume that China's investment share of GDP falls to 20 per cent by 2030, as compared with the 30 per cent share assumed in the baseline scenario. This investment share is broadly in line with the current structure of the US economy and implies a much more rapid transition than is generally expected for China. Under this scenario, China's copper intensity would be around 18 per cent lower by 2030 relative to our baseline scenario, with demand increasing by just 1.5 million tonnes over the next decade.

We then turn to a slower rebalancing scenario, assuming that China's investment share of GDP remains roughly unchanged around 40 per cent by 2030, remaining significantly higher than that in most other economies. In this case, copper intensity would be around 13 per cent higher by 2030 relative to our baseline scenario. Copper consumption in China would be expected to increase by 5.8 million tonnes over the next decade, which would represent the largest volume increase on record.

Given China's substantial share of current copper consumption, a 15 per cent shock to its copper intensity would have a significant impact on the global copper market (Figure 5). Our alternative scenarios thus suggest that China's rebalancing path will play an important role in shaping the evolution of Chinese—and hence world—copper consumption.

[5] As discussed by Stuermer & Schwerhoff (2015), over the long-term copper (and other commodity) prices are trendless (flat) due to increasing extraction brought about by new extraction technologies.

6 Potential structural changes in the world copper market

While our modelling efforts have significantly improved upon existing approaches to capturing regime-switching behaviour, some important aspects may be missing, particularly given potential structural changes in the global copper market. First and foremost, the historical consumption patterns we have estimated across countries inherently reflect the state of technology over history. Going forward, there are reasons to believe copper consumption may shift with changing trends in industrial practices and consumer behaviour. For instance, with the expected rise in the use of electric vehicles (EVs), copper consumption may be revised upward relative to the past for a given state of development. A study commissioned by the International Copper Association (ICA) in 2017 suggests that hybrid EVs require almost twice the amount of copper as traditional vehicles based on internal combustion engines, whereas battery EVs will require almost four times as much. Initiatives to enhance the efficiency of industry and lower its environmental footprint are expected to provide additional boosts to copper consumption. For China, its "Made in China 2025" initiative is expected to lead to additional copper consumption. The plan encourages technologies that require greater copper intensity. These initiatives include the adoption of Industry 4.0, advanced railway networks, higher efficiency motors and renewable energy technologies (e.g., wind power).[6] Thus, while growth in Chinese copper consumption may be expected to slow due to the rebalancing efforts mentioned earlier, policies and technological trends could provide an offsetting influence.

Resource availability was also not explicitly captured in our model and could affect copper use going forward. Global reserves of copper are not scarce, with global copper resources estimated by the United States Geological Survey (2014) to exceed 5,000 million tonnes (which equals 127 years of production at current rates). Chile and Peru are the most significant producers and exporters of mined copper, with Chile alone accounting for 27 per cent of world mined production (Figure 10). Holding technology constant, extracting these resources could prove to be increasingly costly as less accessible, lower-grade resources need to be developed. However, if history is any guide, technology improves, which lowers extraction costs and keeps prices stable (Stuermer & Schwerhoff (2015)).

Increasing consumption combined with dwindling supply is bound to put upward pressure on copper prices, which in turn might lead to new technologies that serve to substitute away from greater copper use over the long-term. Endogenous technological innovation is not captured in our models, and understanding this process could be the subject of future research not just for copper but for other resources as well. Such processes are highly complex and likely difficult to model, since the specific factors that lead to greater technological innovation are not constant through time and are likely to be non-linear.

Regional spillover effects are also not well understood and could be captured better in future models. In addition to spillovers that occur simply due to the transmission of general consumption shocks across regions, technological spillovers may be important. For instance,

[6] Industry 4.0 is the name given to the current trend of automation and data exchange in manufacturing technologies.

innovations that lead to greater copper intensity in one country may be transmitted to other countries through trade or other linkages (e.g., geographic proximity, cultural connections, business networks). With greater global market integration, the speed of transmission is likely to have increased over time, although verifying this phenomenon through empirical study would be a worthy endeavour.

Finally, patterns of development are subject to change, which may influence copper consumption. Countries such as India have developed with a greater emphasis on the service sector vis-à-vis manufacturing relative to other countries in the past. While our panel setting allows for country-specific effects, relationships are being formed based on overall experiences from a broad pool of countries, and country-specific effects are soaking up a broad range of residual effects that are poorly understood.

7 Concluding remarks

This paper examines how the rebalancing of China's economy will likely affect its copper consumption over the next decade, focusing on the relationship between the copper intensity of GDP and the share of investment in GDP. To address potential non-linearities between copper intensity and the investment share of GDP, we apply a PSTR model on a panel of countries using data over the period 1950–2015, with the urbanization rate and level of real GDP per capita as transition variables to determine regimes in the PSTR. Estimation results suggest that there is indeed a significant relationship between a country's copper intensity of GDP and its investment share. Moreover, they confirm that a regime-switching approach is appropriate in this context and that the threshold effects can be determined by a combination of the urbanization rate and real GDP per capita.

Our baseline rebalancing scenario for China implies that copper intensity in China has already peaked and is expected to decline steadily through the next decade. This anticipated reduction in Chinese copper intensity results from the dampening impact of rebalancing and higher per capita income on copper intensity, which more than offsets the upward pressure stemming from the ongoing process of urbanization. An exploration of alternative rebalancing scenarios confirms that China's rebalancing path will play an important role in shaping the evolution of Chinese—and hence world—copper consumption.

References

Aslanidis, N., & Xepapadeas, A. (2008). Regime switching and the shape of the emission-income relationship. *Economic Modelling*, *25(4)*, 731–739.

Bailliu, J., Kruger, M., Toktamyssov, A., & Welbourn, W. (2017). How fast can China grow? The Middle Kingdom's prospects to 2030. *Pacific Economic Review*, (pp. 1–27). https://doi.org/10.1111/1468-0106.12240.

Canas, A., Ferrao, P., & Conceicao, P. (2003). A new environmental Kuznets curve? Relationship between direct material input and income per capita: evidence from industrialised countries. *Ecological Economics*, *46(2)*, 217–229.

Cole, J., & McCoskey, S. (2013). Does global meat consumption follow an environmental Kuznets curve? *Sustainability: Science, Practice and Policy*, *9(2)*, 26–36.

Colletaz, G., & Hurlin, C. (2006). Threshold effects of the public capital productivity: An international panel smooth transition approach. Document de recherche du LEO / LEO Working Paper No. 2006-01.

Davies, R. (1977). Hypothesis testing when a nuisance parameter is present only in the alternative. *Biometrika*, *64(2)*, 247–254.

Davies, R. (1987). Hypothesis testing when a nuisance parameter is present only in the alternative. *Biometrika*, *74(1)*, 33–43.

Destais, G., Fouquau, J., & Hurlin, C. (2007). Economic development and energy intensity: A panel data analysis. In J. Keppler, R. Bourbonnais, & J. Girod (Eds.), *The Econometrics of Energy Systems* (pp. 98–120). London: Palgrave Macmillan.

Döhrn, R., & Krätschell, K. (2013). Long term trends in steel consumption. RUHR Eonomic Papers No. 415.

Duarte, R., Pinilla, V., & Serrano, A. (2013). Is there an environmental Kuznets curve for water use? A panel smooth transition regression approach. *Economic Modelling*, *31(March)*, 518–527.

Ghoshray, A., & Pundit, M. (2016). The impact of a People's Republic of China's slowdown on commodity prices and detecting the asymmetric responses of economic activity in Asian countries to commodity price shocks. ADB Working Paper No. 493.

González, A., Terasvirta, T., & van Dijk, D. (2005). Panel smooth transition regression models. Quantitative Finance Research Centre Paper No. 165, University of Technology Sydney.

Granger, C., & Terasvirta, T. (1993). *Modelling Nonlinear Economic Relationships*. Oxford: Oxford University Press.

Hansen, B. (1999). Threshold effects in non-dynamic panels: Estimation, testing and inference. *Journal of Econometrics*, *93(2)*, 345–368.

Jones, D. (1991). How urbanization affects energy use in developing countries. *Energy Policy*, *19(2)*, 621–630.

Kruger, M., Mo, K., & Sawatzky, B. (2016). The evolution of the Chinese housing market and its impact on base metal prices. Bank of Canada Staff Discussion Paper No. 2016-7.

Luukkonen, R., Saikkonen, P., & Terasvirta, T. (1988). Testing linearity against smooth transition autoregressive models. *Biometrika*, *75(3)*, 491–499.

McKay, S., Sheng, Y., & Song, L. (2010). China's metal intensity in comparative perspective. In R. Garnaut, J. Golley, & L. Song (Eds.), *China: The Next Twenty Years of Reform and Development* (pp. 73–98). Canberra, Australia: Australian National University Press.

Pesaran, M. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, *22(2)*, 265–312.

Roache, S. (2012). China's impact on world commodity markets. IMF Working Paper No. 12/115.

Roberts, I., & Rush, A. (2010). Sources of Chinese demand for resources commodities. Reserve Bank of Australia Research Discussion Paper No. 2010-08.

Roberts, I., Saunders, T., Spence, G., & Cassidy, N. (2016). China's evolving demand for commodities. URL: `http://www.rba.gov.au/publications/confs/2016/pdf/rba-conference-2016-roberts-saunders-spence-cassidy.pdf`. Paper presented at RBA conference on *Structural Change in China: Implications for Australia and the World*, March 2016.

Sahoo, P., Ranjan, D., & Geethanjali, N. (2010). Infrastructure development and economic growth in china. Institute of Developing Economies Discussion Papers No. 261.

Stuermer, M., & Schwerhoff, G. (2015). Non-renewable resources, extraction technology, and endogenous growth. Mimeo.

Terasvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, *89*, 208–218.

United Nations Department of Economic and Social Affairs (2017). *World Population Prospects: The 2017 Revision*. United Nations.

United Nations Department of Economic and Social Affairs (2018). *2018 Revision of World Urbanization Prospects*. United Nations.

United States Geological Survey (2014). Estimate of undiscovered copper resources of the world, 2013. Fact Sheet 2014-3004, January 2014.

Warell, L., & Olsson, A. (2009). Trends and developments in the intensity of steel use: an econometric analysis. Mimeo.

Zhao, P., & Zhang, M. (2018). The impact of urbanisation on energy consumption. *Urban Climate*, *24(June)*, 940–953.

Appendix A: figures and tables

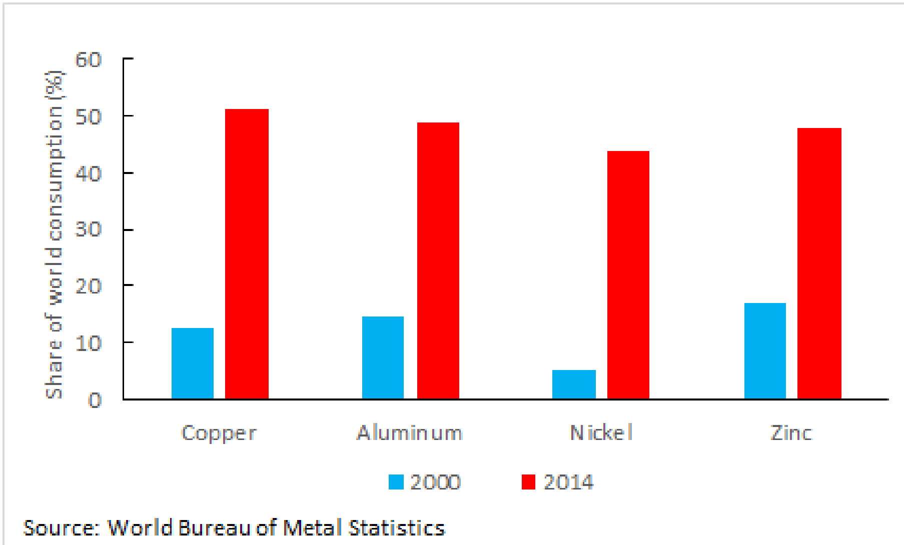

Figure 1: Growth of China's share of the global base metal markets (2000–2014)

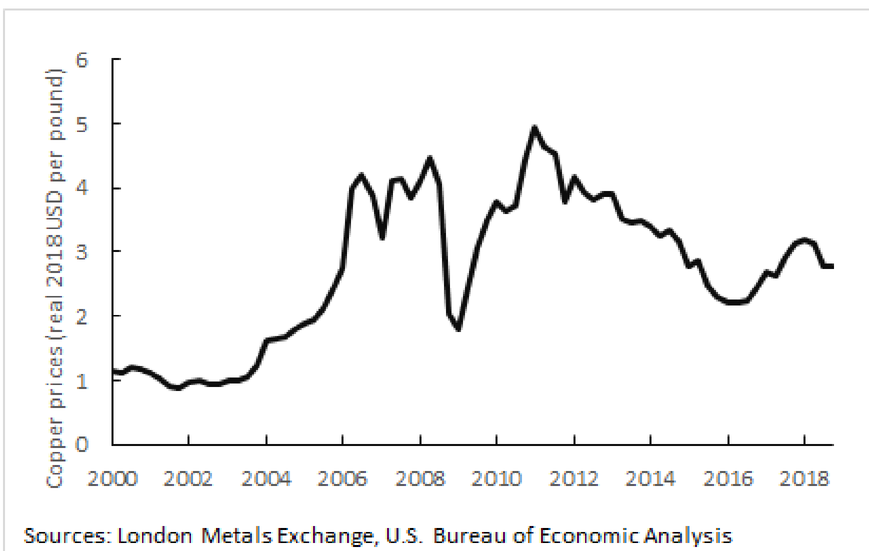

Figure 2: World copper price (2000–2018)

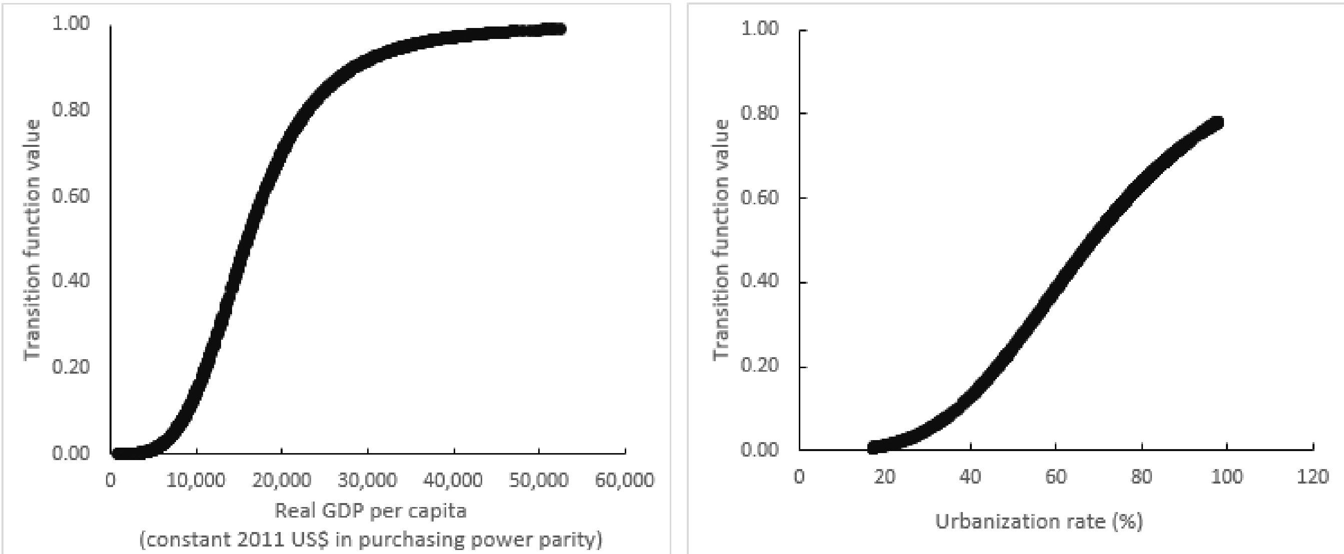

Figure 3: Estimated transition functions

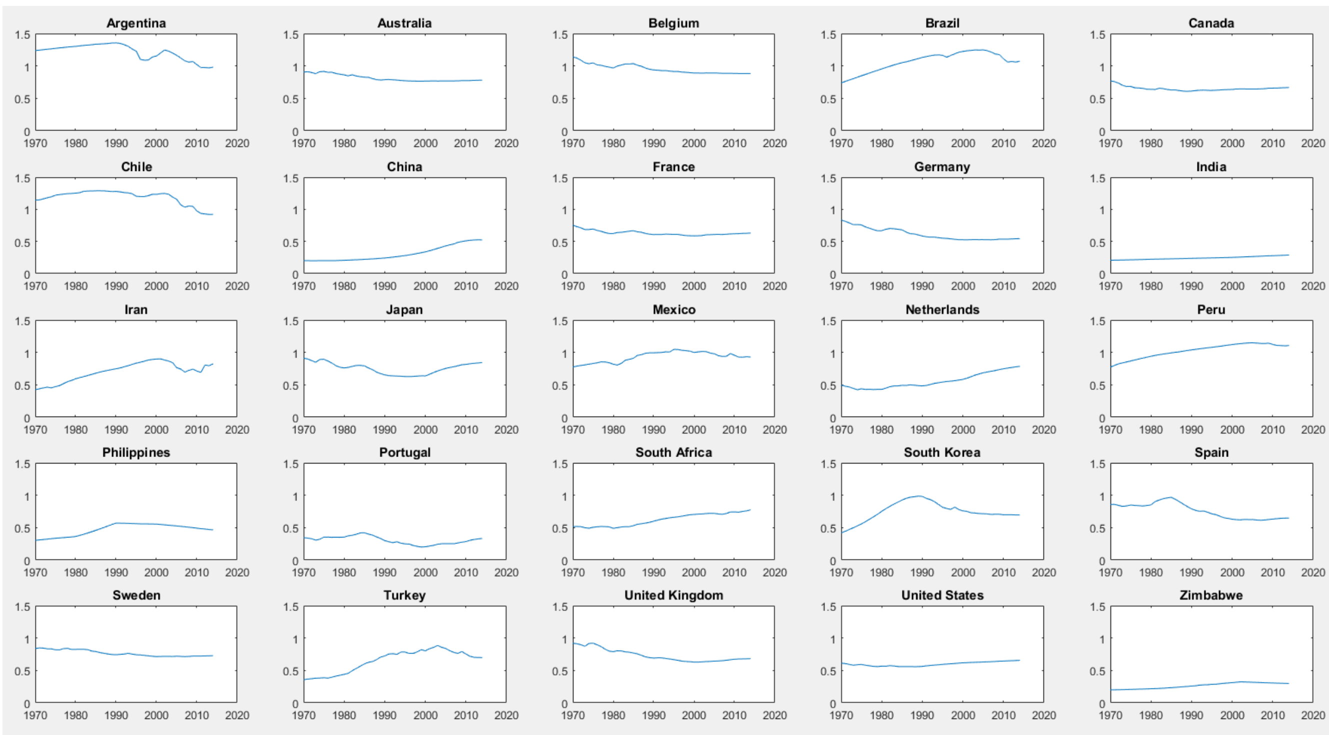

Figure 4: Estimated elasticity of copper intensity to investment share of GDP by country (1970–2016)

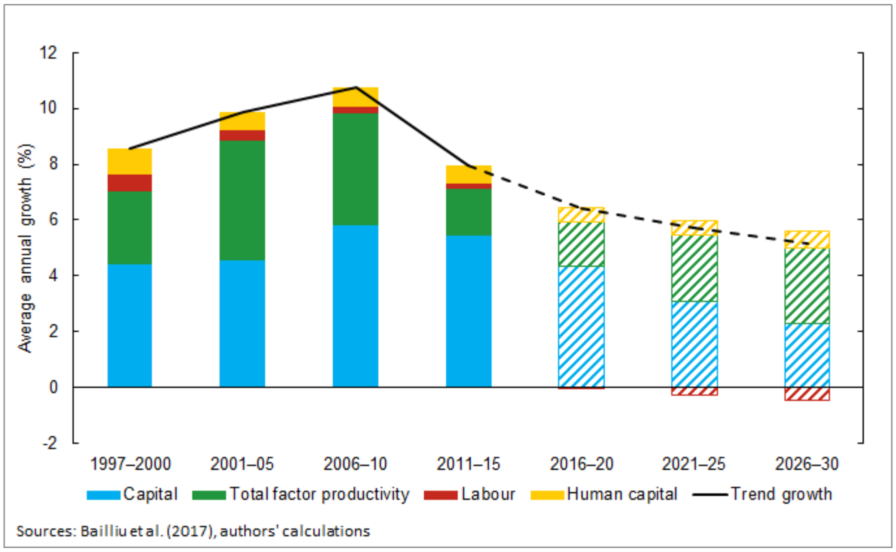

Figure 5: Outlook for China's potential GDP growth

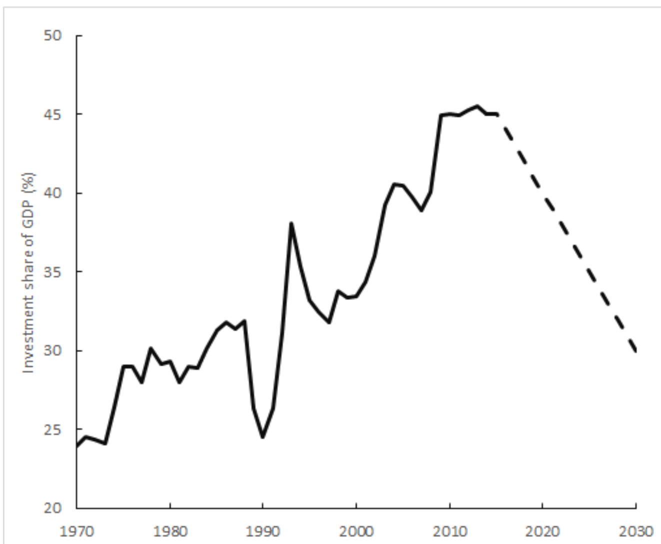

Figure 6: Baseline scenario for Chinese investment share to GDP out to 2030

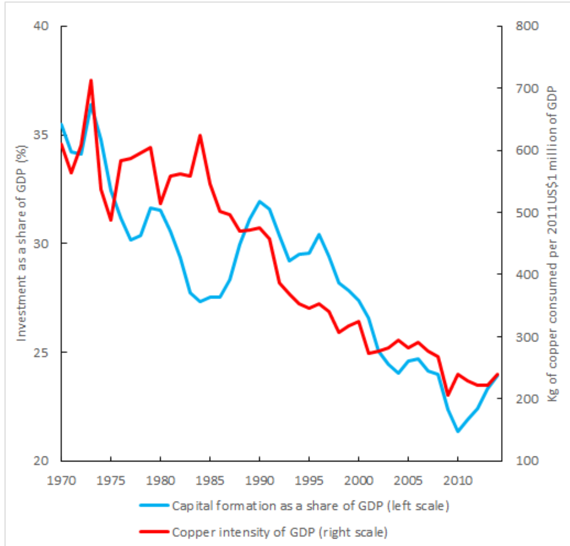

Figure 7: Japanese experience with rebalancing and copper intensity

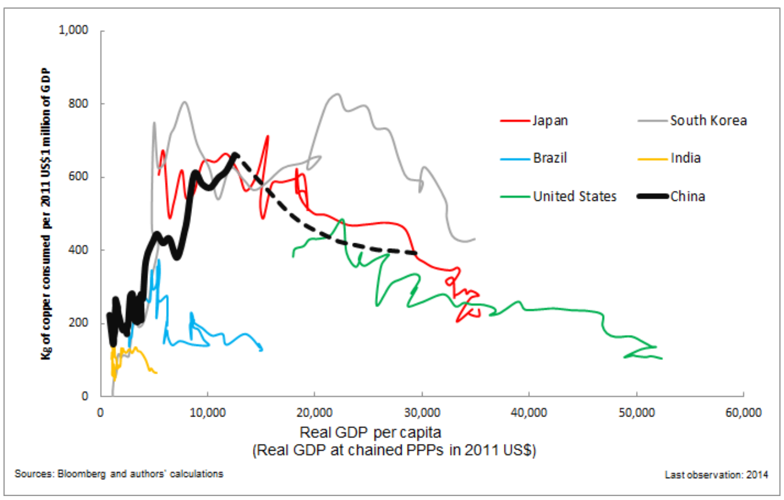

Figure 8: Copper intensity and real GDP per capita by country

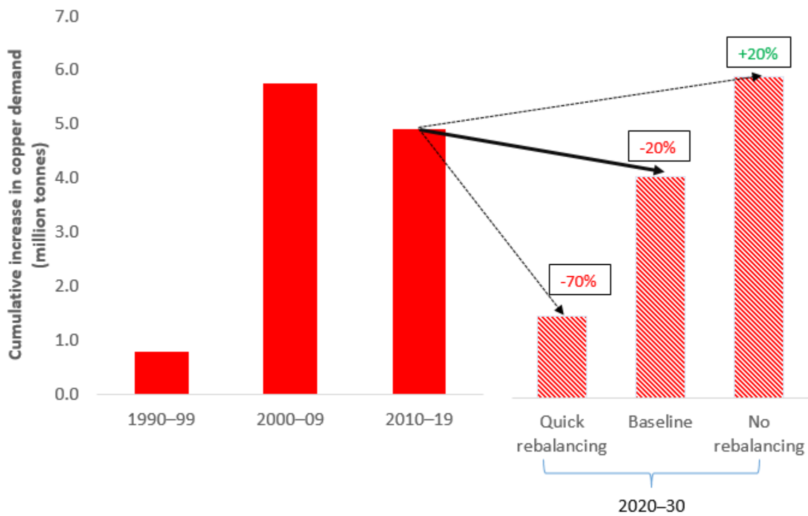

Figure 9: Cumulative increases in China's copper consumption (millions of tonnes)

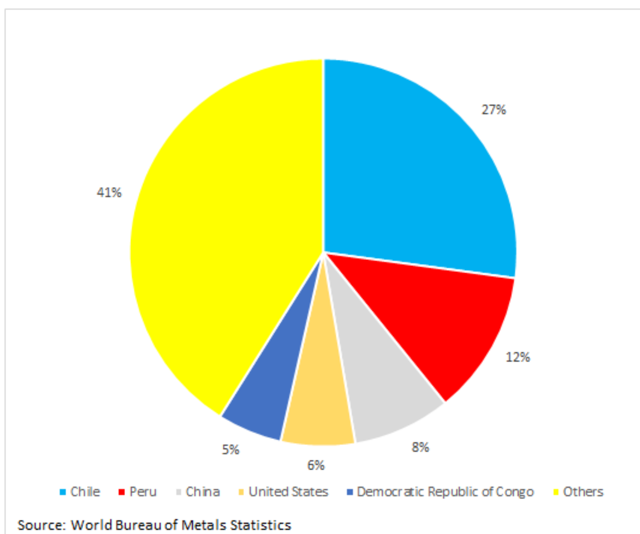

Figure 10: Global mined copper production, share of total

Table 1: Summary statistics

	Mean	Median	Standard deviation
Copper intensity (kg / US$1 million)	297	245	239
Investment share of GDP (%)	23	23	6
Urbanization rate (%)	67	74	20
GDP per capita ($)	17524	14875	12856

Table 2: Homogeneity tests

	Wald test		Fisher test		Likelihood-ratio test	
Transition variable	Statistic	p-value	Statistic	p-value	Statistic	p-value
Urbanization rate (1)	111.0	0.00	40.1	0.00	116.9	0.00
Real GDP per capita (2)	0.9	0.83	0.3	0.83	0.9	0.83
Both (1) and (2)	137.2	0.00	50.8	0.00	146.3	0.00

Table 3: Estimation results

	Coefficient	Standard error
β_1	0.188	0.114
β_2 (marginal effect of urbanization regime)	1.685	0.106
β_3 (marginal effect of GDP regime)	-0.635	0.070
β_4 (price of copper)	-0.014	0.002
β_5 (time trend)	-0.210	0.047

Table 4: Baseline scenario assumptions (per cent)

	2015	2029
Investment share of GDP	45	30
Urbanization rate	56	71
Real GDP growth	7.3	4.7
Population growth	0.5	0

Table 5: Baseline scenario results

	Copper intensity (kg / US$1 million)
2015	638
Rebalancing	-100
GDP per capita	-318
Urbanization	283
Efficiency gains	-99
2030	393

Appendix B: list of countries

Argentina
Australia
Belgium
Brazil
Canada
Chile
China
France
Germany
India
Iran
Japan
Mexico
Netherlands
Peru
Philippines
Portugal
South Africa
South Korea
Spain
Sweden
Turkey
United Kingdom
United States
Zimbabwe

Appendix C: data description

Variable	Source	Description
Copper intensity (kg consumed / US$1 million real GDP)	*Copper consumption*: World Bureau of Metals Statistics	*Copper consumption*: Refined copper consumption with missing values interpolated using a linear/cubic spline
	Real GDP: Penn World Tables version 9.0	*Real GDP*: Expenditure-side real GDP at chained PPPs (in millions of constant 2011 US$)
Real GDP per capita	*Real GDP*: Penn World Tables version 9.0	*Real GDP*: Expenditure-side real GDP at chained PPPs (in millions of constant 2011 US$)
	Population: World Bank's World Development Indicators	*Population*: Total population
Investment share of GDP	World Bank's World Development Indicators	Gross capital formation as a share of GDP
Urbanization rate	World Bank's World Development Indicators	Share of total population living in urban areas
Real copper price	*Copper price*: London Metals Exchange	*Copper price*: Measured in US dollars per pound
	US GDP deflator: U.S. Bureau of Economic Analysis	*US GDP deflator*: GDP implicit price deflator