The rise of benchmark bonds in emerging Asia

By Eli Remolona and James Yetman[1]

Abstract

The most liquid of bonds – often referred to as benchmarks – serve as the focus of price discovery, and in doing so, help to improve the efficiency of financial markets. In this paper, we summarise our ongoing research to see if benchmark bonds exist in emerging market economies and to assess whether there is a connection between their existence and the maturities where authorities have designated official benchmarks. Can we find bonds that act like benchmarks in terms of their liquidity and price movements? We consider four emerging market economies in Asia: Indonesia, Malaysia, the Philippines and Thailand. These are countries where the governments have actively promoted the development of benchmark bonds and have issued those bonds in size to foster their liquidity. To varying degrees, these authorities have also spread their chosen benchmarks across wide maturity ranges in an effort to build benchmark yield curves. To identify the de facto benchmarks, we propose measures of liquidity and price discovery that are appropriate to the sparse data available for these markets. We find that the existence of de jure and de facto benchmarks often do coincide.

JEL classification: G10, G12, G14

Keywords: benchmark bond, price discovery, liquidity, principal component analysis, re-openings, informational public good

[1] Bank for International Settlements, Representative Office for Asia and the Pacific, 78th Floor, Two IFC, 8 Finance Street, Central, Hong Kong SAR; Remolona: emr6@williams.edu; Yetman: james.yetman@bis.org. We thank Anamaria Illes, Jose Vidal Pastor, and Jimmy Shek for excellent research assistance and Mark Knezevic and Nuttathum Chutasripanich for expert advice. We also thank Terence Chong, Frank Packer, seminar participants at the BIS and conference participants at the Bank of Korea-BIS conference on "Asia-Pacific fixed income markets: evolving structure, participation and pricing" for comments. Any remaining errors are our own. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements.

1. Introduction

Benchmark bonds can play an important role in ensuring that fixed income markets function well. For example, price discovery tends to occur in benchmark bonds: their prices react first to new information, and they serve as a means of conveying that information to other bonds (Dunne et al (2002)). Hence, while the discovery of information may be concentrated in a single bond, the information would be quickly conveyed to prices across the market more broadly (Wooldridge (2001)). This process requires that the benchmark be more liquid relative to other bonds.

How exactly a bond becomes a benchmark is less clear: in some economies, including the ones we focus on, authorities designate specific issues as benchmarks and issue them in size, partly in the hope that they will act like benchmarks. If a bond starts to act like a benchmark, network externalities reinforce its role, ensuring that it is likely to remain a benchmark for some time.

Across the fixed income market as a whole, government securities issued in a given currency are most likely to act as benchmarks for private sector securities issued in the same currency. This is partly because governments are relatively creditworthy borrowers, which simplifies the pricing of their securities. The existence of benchmarks, however, improves the functioning of the wider bond market; hence governments face incentives to encourage their development, and support their continued existence where they are already present. The International Monetary Fund and World Bank (2001) have advised governments to consolidate issuance in a few maturities across the maturity spectrum to support the evolution of a benchmark yield curve. And when some advanced economies were facing shrinking government debt markets on account of large budget surpluses, authorities sought to maintain the volume of gross issuance in certain securities to support the benchmarks (McCauley and Remolona (2000)).

The question we are seeking to address is whether benchmark sovereign bonds exist in the economies we study and whether there is any association between their existence and the official designation of benchmark bonds by authorities. While the authorities in Indonesia, Malaysia, the Philippines and Thailand have officially designated benchmark bonds, and have issued them in size to foster their liquidity, it is market participants who are the final arbiters on whether they will actually serve as benchmarks. In the US Treasury market, the "on-the-run" bond of a given maturity becomes the benchmark without the need for any official designation.

In the largest government bond markets, benchmark bonds seem to have naturally emerged as the market has developed, and generally exist at four maturities: two, five, 10 and 30 years (McCauley and Remolona (2000)). In some emerging market economies (EMEs), by contrast, the authorities seem keen to give markets a helping hand, designating specific issues as benchmark bonds and issuing them in size. Table 1 provides a summary of the maturities the authorities have chosen for their designated benchmarks.

To assess the existence of benchmark bonds, we use daily observations on prices, yields and bid-ask spreads for sovereign bonds in the four economies. We are unable to obtain data on trading activity, and the data we do have are missing for a large number of trading days. Given the sparseness of the data, we propose novel measures of liquidity and price discovery. Overall, we find that the de jure benchmarks are often also the de facto benchmarks.

In the following section, we discuss ways to measure liquidity with the sparsely available data and summarise our preliminary results. In Section 3, we propose a way to measure price discovery and summarise preliminary results. In Section 4, we consider the emergence of benchmark bonds in our sample of Asian EMEs by looking at the success of de jure benchmarks over time. Section 5 then concludes.

De jure benchmark maturities for four jurisdictions

Designated benchmark maturities shaded in green

Table 1

	Maturities in years										
	2	3	4	5	7	10	15	20	25	30	50
Indonesia (2010-2018)				■		■	■	■			
Malaysia		■		■	■	■	■	■		■	
Philippines	■	■	■	■	■	■		■	■		
Thailand (FY 2016-2018)				■		■	■	■		■	■

Sources: For Indonesia, Ministry of Finance; for Thailand, Public Debt Management Office; for Malaysia and the Philippines, exchanges of emails with regulators and market participants.

2. Liquidity and benchmarks

2.1 Measuring liquidity

We first seek to identify possible benchmark bonds based on their relative liquidity properties. We take three different measures of liquidity: the average of each of the bid-ask spread and the yield within a month, and the number of days within the month for which price quotes for the bond are available. A benchmark bond would be expected to have a low bid-ask spread, a low yield (since its liquidity would command a premium) and have quotes available relatively frequently. We then combine these measures by summing their rank orders across these three liquidity measures (from most to least liquid). The most liquid bond is the one with the lowest rank-order sum. We then compare the length of time for which the most liquid bond remains so with what we would expect to occur by chance, taking into account the number of bonds in the maturity and if any are equally most liquid. The underlying idea is that once a bond becomes a benchmark bond, and hence more liquid than comparable bonds, the network externalities of a benchmark should preserve its liquidity for multiple months. We carry out this exercise for each country for eight different symmetric non-overlapping maturity buckets, each centred on one of the maturities listed in Table 1.

We assess how persistently a bond is most liquid by calculating the probability P_t of observing the outcome that we would observe under the null hypothesis that the most liquid bond is randomly assigned each period. If there is a benchmark,

we would expect to find a downward-step-function (in logs). Then, when the benchmark switches to a new bond, the probability would jump and the step function would start anew. Graph 1 illustrates this probability under two hypothetical cases: for a bin with four bonds where the benchmark changes every 10th period, and where the most liquid bond is randomly chosen every period (right-hand panel).

Hypothetical probabilities that a bond remains most liquid for multiple periods

Logarithmic scale — Graph 1

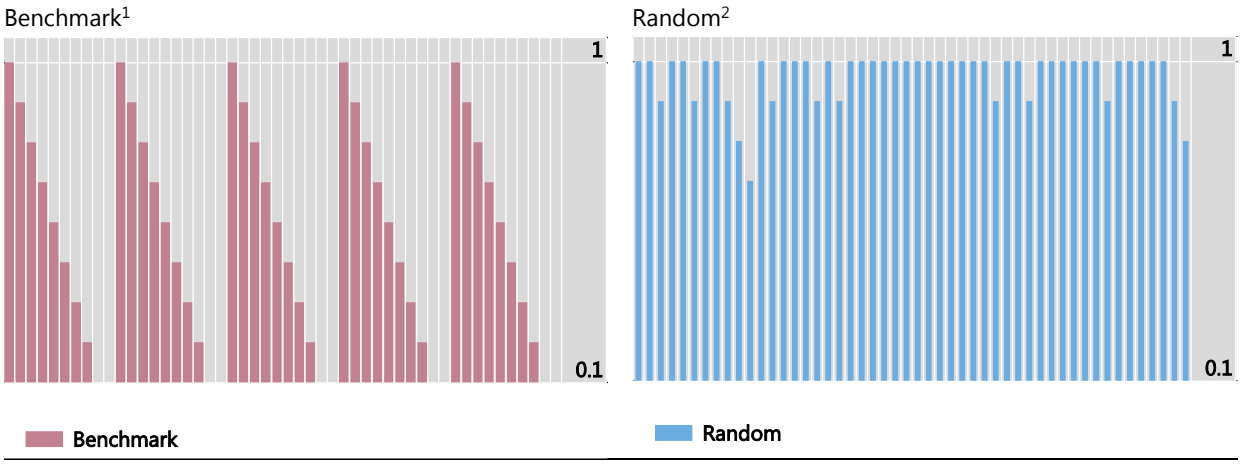

[1] P_t for a bin with four bonds under the assumption that a benchmark bond is chosen every 10th period and remains the most liquid as long as it is the benchmark. [2] Example of P_t for a bin with four bonds under the assumption that the most liquid bond is randomly assigned each period.

Source: Authors' calculations.

2.2 Results on liquidity

When we apply this test across the different countries and maturity bins, we obtain the results presented in Graph 2. For Indonesia (Graph 2a), liquidity persistence consistent with benchmark behaviour is most evident at the 10-year, 15-year and 20-year maturities. These maturities also happen to be de jure benchmark maturities. However, the 5-year maturity is also a de jure benchmark, but liquidity persistence is less evident here. Graph 2b provides the same visual assessment for Malaysia. Here again there is evidence of benchmark behaviour at most maturities, especially the de jure benchmark maturities. For the Philippines, Graph 2c shows liquidity persistence to be most evident at the 20-year maturity, which is a de jure benchmark. There is also some evidence of benchmark behaviour at the 15-year maturity, which is not a de jure benchmark. Finally, for Thailand, Graph 2d shows considerable evidence of benchmark behaviour at most maturities up to 20 years, whether or not these are de jure benchmark maturities. Taken together, there is limited consistency between the benchmark-like behaviour in terms of liquidity persistence and those maturity bins where there is supposed to be a de jure benchmark bond.

Is the most liquid bond persistent in Indonesia?

Graph 2a

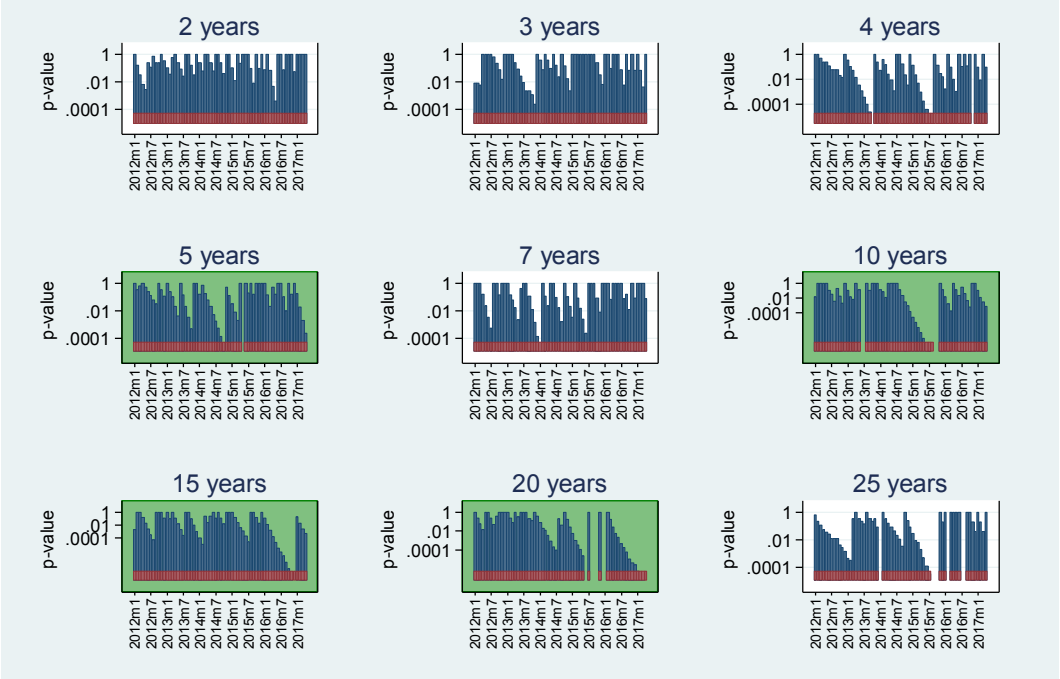

Is the most liquid bond persistent in Malaysia?

Graph 2b

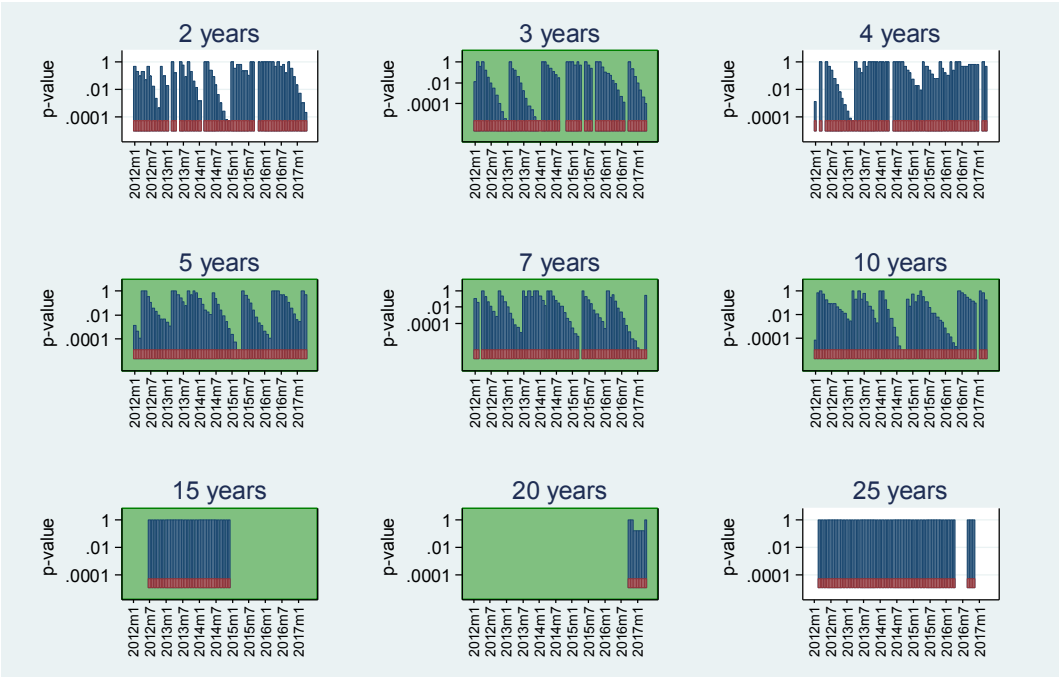

[1] The solid red bar indicates inclusion in the sample (ie at least three bonds in a given maturity bin). The blue bars indicate the probability of the most liquid bond being most liquid for the observed number of recursive periods if this was randomly assigned. Green shading indicates maturities with designated benchmarks. Note the log vertical scale.

Sources: Authors' calculations

Is the most liquid bond persistent in the Philippines?

Graph 2c

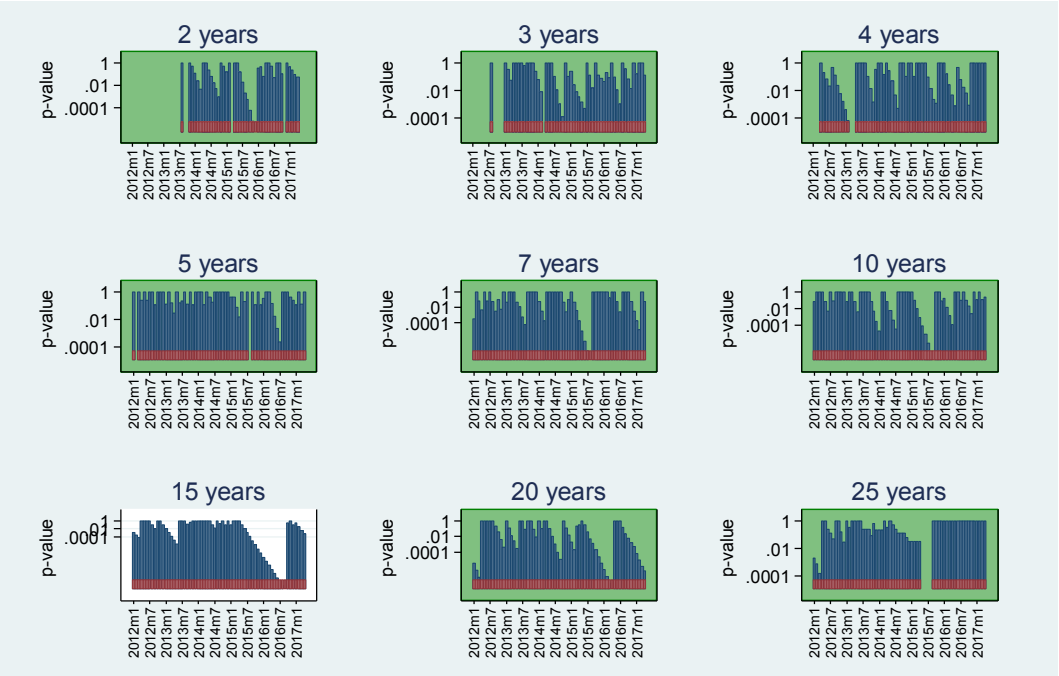

Is the most liquid bond persistent in Thailand?

Graph 2d

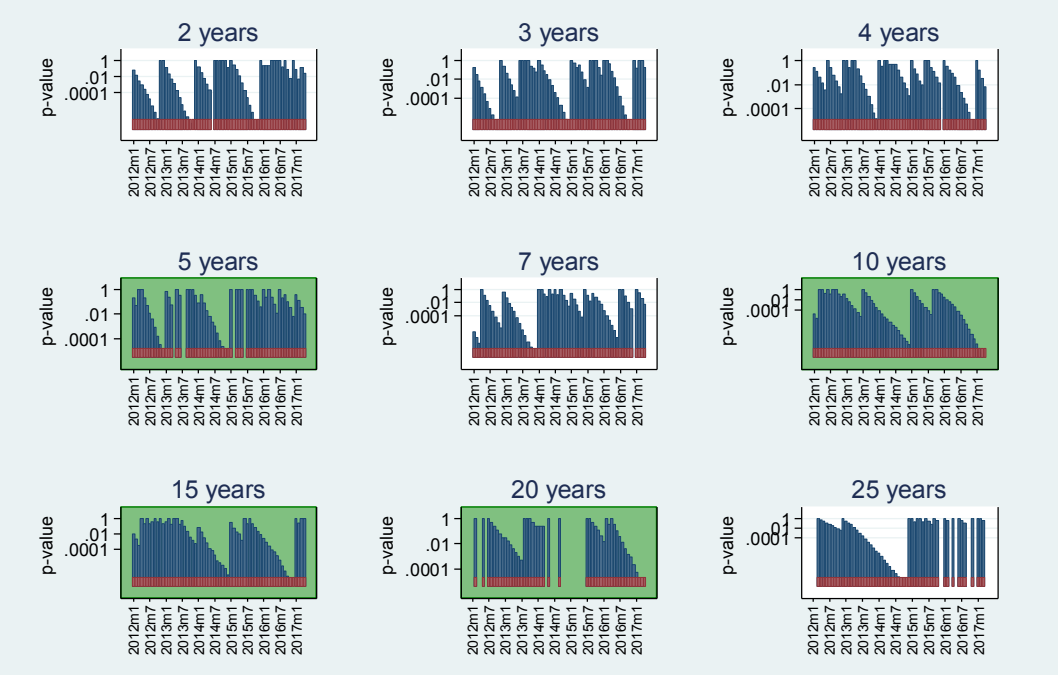

[1] The solid red bar indicates inclusion in the sample (ie at least three bonds in a given maturity bin). The blue bars indicate the probability of the most liquid bond being most liquid for the observed number of recursive periods if this was randomly assigned. Green shading indicates maturities with designated benchmarks. Note the log vertical scale.

Sources: Authors' calculations

3. Price discovery and benchmarks

3.1 Measuring a bond's role in price discovery

If a bond acted as a benchmark in the price discovery process, price movements should be disproportionately affected by systematic factors compared with the price movements of other bonds of similar maturity (Dunne et al (2007)). To asses this, we conduct a principal component analysis (PCA) on the daily change in the prices of bonds of similar remaining maturities and interpret the weight on the first PC as a measure of the importance of the systematic factor.

We face one key challenge in implementing the PCA on our data: our panel is unbalanced, with prices unavailable for some bonds on some days in nearly all months. However, a PCA requires a balanced panel. To remedy this, we interpolate the data panel in such a way that the first PC is not affected by the interpolated factors. We do so with the following steps:

(i) replace missing values with the average change in the yield for the bond on other days in the same month;

(ii) run the PCA based on the now-balanced panel, and extract the fitted values based on the first PC for each of the series;

(iii) replace the original interpolated values with the fitted values and repeat step (ii) until the estimation converges (defined by the fraction of the explained variance increasing by less than 0.00001 from one iteration to the next); and

(iv) compute the load on the first PC as the Pearson correlation coefficient between the first PC and the original data series (with missing values left empty).

Once the procedure has converged at step (iii), interpolated values are approximately equal to their fitted values in the final iteration. Given that the influence of a single data point on the PC is a function of the difference between the data series and its fitted value, this ensures that the missing values have virtually no effect on the first PC; its value at each point in time is instead determined by the non-missing values. Also, when we compute the loading on the first PC at step (iv), this is calculated as a function of only the actual available data, and is not affected by the interpolated data.

For each PCA exercise, we compare the loading on the first PC for the candidate benchmark bond identified on the basis of its liquidity characteristics (above) with other bonds in the same PCA. Note that the PCA is based on different, complementary, information than was used to construct the candidate benchmark bond in the previous section. If the candidate benchmark bond identified in the previous exercise also has the highest loading on the first PC, we interpret this as much stronger evidence that a benchmark exists than if the candidate benchmarks differ across the two measures much of the time.

Loadings on first PC for Indonesia

Graph 3a

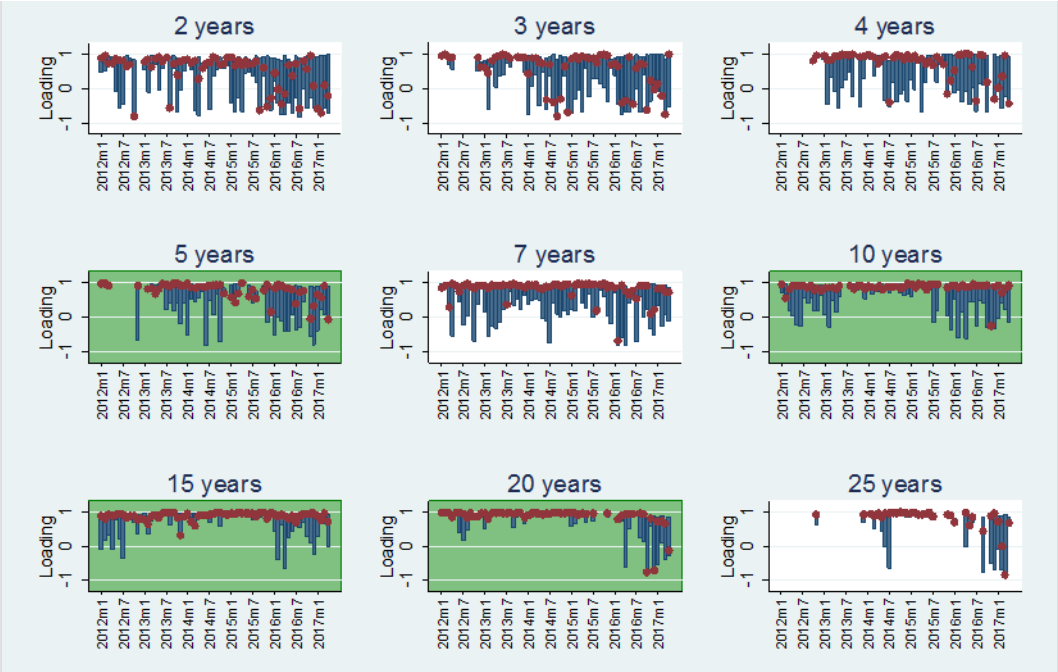

Loadings on first PC for Malaysia

Graph 3b

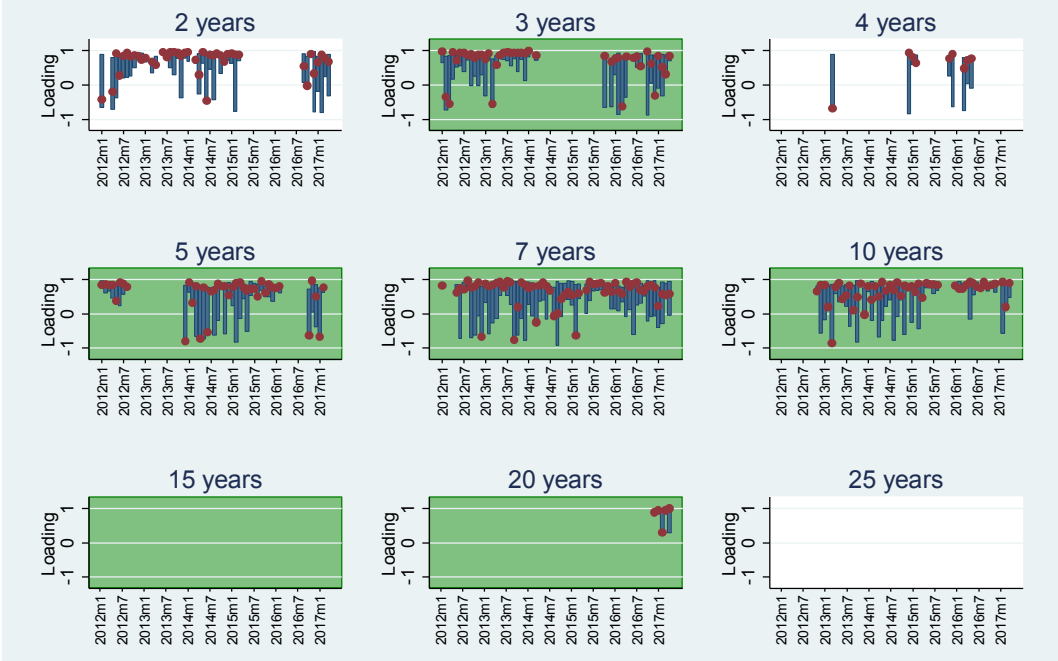

[1] Blue lines indicate the range of loadings on the first principal component. Red dots indicate the loadings of the candidate benchmark bond identified on the basis of liquidity. Green shading indicates maturities with designated benchmarks.

Sources: Authors' calculations

Loadings on first PC for the Philippines

Graph 3c

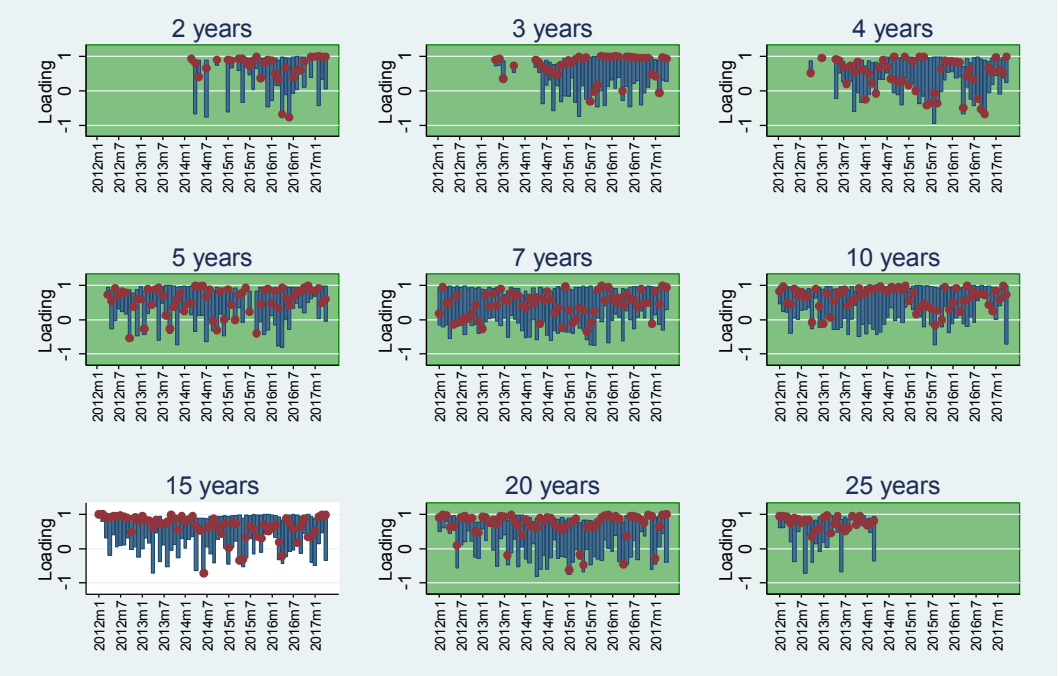

Loadings on first PC for Thailand

Graph 3d

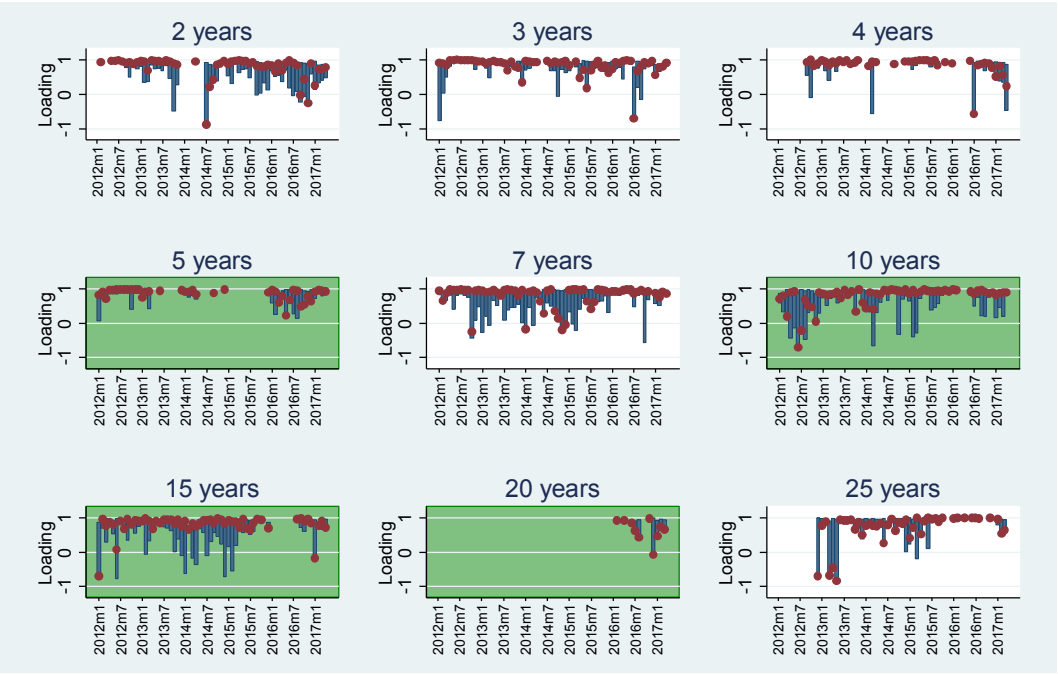

[1] Blue lines indicate the range of loadings on the first principal component. Red dots indicate the loadings of the candidate benchmark bond identified on the basis of liquidity. Green shading indicates maturities with designated benchmarks.

Sources: Authors' calculations

3.2 Results on price discovery

Our results on price discovery are inconclusive, especially when we consider the consistency of these results with those on liquidity. Graph 3 illustrates the results for evidence of price discovery. Blue bars each month display the range of loading on the first PC across all bonds, while red dots indicate the factor loading for the candidate benchmark, identified in terms of its persistent liquidity characteristics. In an ideal world, with a perfectly functioning benchmark, we would expect the red dot to always be at the highest point of the bar, and to have a factor loading of one.

In at least some cases, the candidate benchmark comes close to this ideal. Taking Indonesia, for example, the factor loading for the candidate benchmark has either the highest loading or close to it most months for bonds at the 10- or 15-year maturity. (Both of these are also de jure benchmarks, as indicated by green shading). For Malaysia, the correlation between candidate benchmarks and high loadings is much smaller at all horizons. For the Philippines, the greatest evidence of a link is at the 20-year maturity. Finally, for Thailand, the strongest evidence is at the 10-year maturity.

4. Have benchmarks emerged in emerging Asia?

The primary question of the paper is whether benchmark bonds have emerged in emerging Asia. We provide one piece of tentative evidence that relies on our results on liquidity. We count up the number of months in a given year for which there is a benchmark based on just liquidity characteristics (defined by the probability as displayed in Graph 2 falling below 0.01 for a given benchmark candidate). We also count the number of months for which there is sufficient data to be included in our sample. This latter number reflects the availability of bond information on Bloomberg (our data source) and can be interpreted as a proxy of overall bond market development.

Frequency of benchmark behaviour based on liquidity persistence — Graph 4

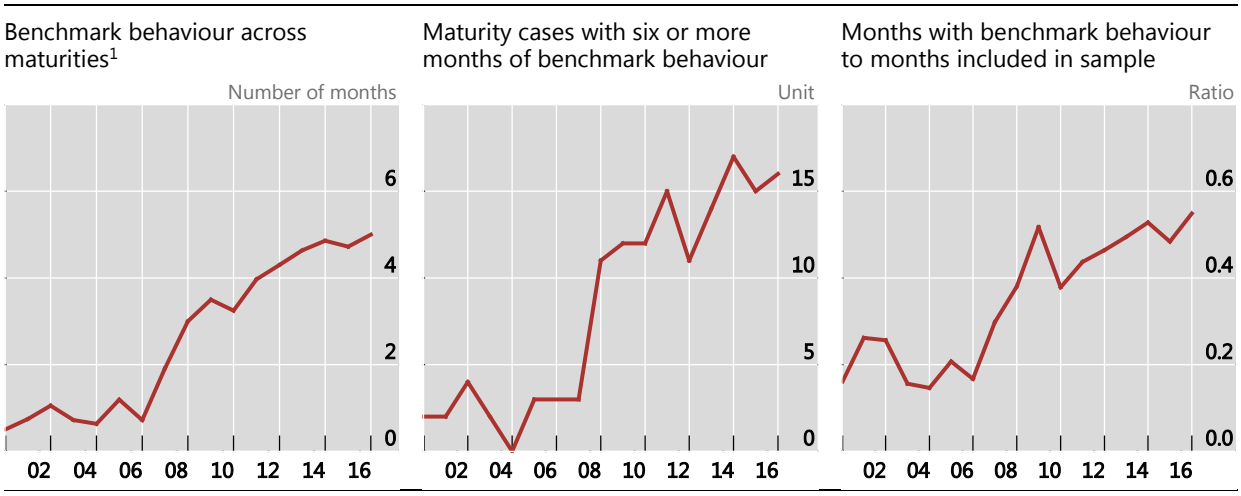

[1] Yearly average.

Source: Authors' calculations.

We summarise these counts, by year, in Graph 4. Across all the maturities, the average number of months with benchmark liquidity behaviour is less than 2.0 for all years before 2008, between 3.0 and 4.0 for the following four years, and more than 4.0 thereafter (left-hand panel). Similarly, the number of maturities with benchmark behaviour for at least six months of the year has increased from the low single digits for 2000–2007 to 11-15 for the next four years and 14-17 thereafter (centre panel). One possible explanation for this behaviour is that these counts are biased by the increased availability of bond market data, as markets have developed, since this is the upper bound on the number of months of benchmark behaviour in our study. However, we can demonstrate that this is not the case. Ignoring cases with zero observations in the sample – since the ratio is then undefined – the ratio of months with benchmark behaviour to total months in the sample has a clear upward trend over the sample as well, increasing from around 0.2 before 2007 (indicating that we find evidence of benchmark behaviour for around 20% of months in the sample) to around 0.55 at the end of the sample (benchmark bond behaviour is present in more than half of all months in the sample; right-hand panel). Put differently, the ratio is positively correlated with the number of months present in the sample (ρ =0.45).

We can also compare the evidence of benchmark behaviour against maturities in which there are de jure benchmarks, as listed in Table 1. We focus on the final three years of the panel and compute the average number of months with benchmark behaviour and the average ratio (where this is defined) for all maturities with de jure benchmarks and all other maturities, for each country. With the exception of the country where the data is least informative (ie the Philippines, for which there is only one non-de jure benchmark maturity available for comparison), we find consistent evidence of more benchmark behaviour at maturities where there are designated de jure benchmarks (as seen in Table 2). The average number of months is higher, as well as the ratio, in all other cases. Again, excluding the Philippines, the smallest differences are for Thailand, although in level terms it has the greatest evidence of benchmark behaviour overall.

Comparing across countries more generally, the de jure benchmark maturities of Indonesia and Thailand have de facto benchmark bonds present for a larger number of months each year than the other countries in our sample (6.08 and 7.00 respectively, compared with less than five for Malaysia and the Philippines). Thai non-benchmark maturity bonds generally display more benchmark-like behaviour by these metrics than the de jure benchmark maturities for the other countries.

Comparison of liquidity between de jure benchmark maturities and others

2014–2016

Table 2

	Average number of months		Average ratio	
	De jure benchmark maturities	Other maturities	De jure benchmark maturities	Other maturities
Indonesia	6.08	3.40	0.62	0.36
Malaysia	4.67	1.67	0.48	0.17
Philippines	4.54	5.67	0.43	0.52
Thailand	7.00	6.13	0.63	0.55

[1] "Average number of months" is the average number of months for which there is evidence of benchmark behaviour, as illustrated in Graph 4, for the 2014–2016 period. "Average ratio" is the average ratio of the number of months for which there is evidence of benchmark behaviour to the number of months included in the sample, for the 2014–2016 period. These are each calculated separately across maturities for which there are de jure benchmark bonds (as indicated in Table 1) and other maturities.

Source: Authors' calculations.

5. Conclusions

So what effect does government policy have in creating benchmark bonds? In some EMEs, governments have chosen bonds that they wish market participants would accept as benchmarks. In this paper, we consider four such economies in Asia. The authorities in Indonesia, Malaysia, the Philippines and Thailand, in particular, have each designated specific government bonds as de jure benchmarks and have issued those bonds in size to foster their liquidity.

Given that the markets we are focusing on are still in the process of development, the available data is sparse. We have proposed statistical tools that can be used to assess benchmark behaviour in spite of this. To assess liquidity, we combined rankings based on relative yields, bid-ask spreads and the number of days for which quotes are available. To assess the importance of a bond in price discovery, we conduct a PCA and use the loadings on the first PC as a measure of that importance. We then cross-checked the resulting evidence of benchmark behaviour with whether there was a de jure benchmark at that maturity.

Based on preliminary findings, there is some evidence of the existence of persistently highly liquid bonds around the maturities for which benchmark bonds are officially designated. Further, the identified highly liquid bonds often have high factor loadings on the first PC of returns on bonds with similar maturities, consistent with them playing an important role in price discovery.

In ongoing work, we will look more closely at the evidence of benchmark behaviour at the level of the individual designated benchmark bond. We are working to identify the specific bonds that have been designated as de jure benchmarks and will investigate whether those bonds act like benchmarks in reality.

References

Dunne, P, M Moore and R Portes (2002): "Defining benchmark status: an application using euro-area bonds", National Bureau of Economic Research, no w9087.

Dunne, P, M Moore and R Portes (2007): "Benchmark Status in Fixed – Income Asset Markets", Journal of Business Finance & Accounting, vol 34, no 9–10), pp 1615–34.

IMF and World Bank (2001): "Developing benchmark issues", Developing Government Bond Markets: A Handbook, Chapter 4, pp 117–50.

McCauley, R and E Remolona (2000): "Size and liquidity of government bond markets", BIS Quarterly Review, pp 52–60, November.

Wooldridge, P (2001): "The emergence of new benchmark yield curves", BIS Quarterly Review, pp 48–57.