Migration in the time of crisis: evidence on its effectiveness from Indonesia

M. Sandi[1] and L.A. Winters[2]

Preliminary – Please do not cite without the authors' permission

December 2013

ABSTRACT

This paper explores the role of geographic mobility as a major element of economic adjustment after the occurrence of a spatially-heterogeneous economic shock. At this aim, the high degree of mobility in Indonesia during the years of the East-Asian financial crisis is exploited; estimates are presented on the effectiveness of migration of adults in Indonesia (1998-2000) on the evolution of living standards, as measured by consumption per capita, by 2000 and by 2007/8. The analysis extends the existing research on risk-coping strategies by: (i) investigating the effects of migration and changes in living arrangements as immediate responses to an economic shock. We assess the relevance of mobility not only for migrants themselves or their origin households, but also for migrant-recipient households; (ii) investigating the effectiveness of migration both in the short run and in the long run; (iii) using semi-parametric techniques to estimate the average treatment effect on the treated (ATET) for all households affected by migration. Evidence from Indonesia suggests that migration in times of crisis is an effective strategy on average. While in the short run migrant-receiving households seem to experience lower average consumption growth than non-migrant affected households, in the long run migrant, origin and recipient households all fare at least as well as non-migrant households.

[1] Department of Economics, University of Sussex – E-mail: m.sandi@sussex.ac.uk

[2] Department of Economics, University of Sussex – E-mail: l.a.winters@sussex.ac.uk

Very special thanks to Iftikhar Hussain, Peter Dolton, Barry Reilly and Jeff Smith for beneficial discussions and helpful comments.

1. Introduction

Rural and urban households in developing economies face significant idiosyncratic and common risk, which often translates in high income variability and an inability to smooth consumption. In developing (and risky) environments, formal credit and insurance markets appear to contribute little to reducing income risk and its consequences, especially for the poor, who are often excluded from these markets. The failure to cope with income risk may lead to relatively high levels of transient poverty; this, in turn, may affect nutrition, health and education and lead to inefficient and unequal intra-household allocations. In the long term, this may result in low human capital accumulation and, ultimately, in reduced chances to escape poverty. For these reasons, poor households have developed sophisticated strategies to deal with risk. This paper investigates the consequences of one important mechanism of economic adjustment in response to an economic downturn, namely the geographic mobility of individuals and related changes in households' living arrangements.

We present a case study from the East-Asian crisis in the late 1990s - Indonesia, arguably the hardest hit country in the region. Due to its depth and distribution, the crisis induced massive net migration towards rural areas in Indonesia (Fallon and Lucas, 2002). Following a period of substantial trade liberalization, industrialization and economic growth, the crisis hit Indonesia particularly hard: between 1997 and 1998 the rupiah collapsed, GDP contracted by an estimated 12-15% and a political crisis followed. Only after Soeharto's resignation, in May 1998, the new government announced support for a set of "safety net" (JPS, Jaring Pengaman Sosial) programs: in the July 1998 budget, programs included targeted sales of subsidized rice, work creation programs, scholarships to students and block grants to schools, targeted health care subsidies and community block grants. Nonetheless, a good deal of evidence suggests that relatively little income smoothing was achieved, especially by poorer households, and many households appeared to be vulnerable to income risk during the crisis (see, e.g., Thomas et al., 2004; Strauss et al., 2004; and Thomas and Frankenberg, 2007).

The economic shock was sudden, substantial and, most interestingly for us, geographically heterogeneous, a feature which we presume induced the enhanced levels of internal mobility. We explore the consequences of the migration of adults between districts in Indonesia over the crisis period (1998-2000) for the evolution of living standards. The latter are measured by the evolution of consumption per capita from 1998 to 2000 (which we refer to as the short run) and to 2007/8 (long run). The analysis extends the existing research on risk-coping strategies in the following ways. First, the effects of migration and changes in living arrangements as immediate responses to an economic shock are investigated. The analysis is conducted at the household level, and households are classified according to whether or not (and the way in which) they modified their living arrangements in the years of the crisis. Hence, households affected differently by migratory flows are treated separately, and the relevance of mobility is assessed for migrants themselves, as well as for their origin households and for migrant-recipient households. It is important to stress that we do not treat migration as a once-and-for-all decision because we do not consider whether people migrated in the post-crisis period; rather, we seek to analyse the specific effect of migration that took place in the years of the crisis.

Second, the effectiveness of migration is tested both in the short run (up to 2000) and, more unusually, in the long run (up to 2007/8), ten years after the peak of the crisis. Third, the remarkable heterogeneity in the impact of the crisis across regions in Indonesia and the consequent price dispersion are exploited to assess the effectiveness of mobility in the presence of idiosyncratic economic shocks. The relevance of mobility across districts is studied, and the entire sample is further subdivided in a number of ways in order to better isolate the crisis, both at the social and at the geographic level. Fourth, a number of semi-parametric techniques are used to estimate the average treatment effect on the treated (ATET) for all households affected by migration. In fact, the econometric analysis mostly relies on matching techniques, but we also consider semi-parametric local average treatment effects (LATE) estimators with covariates suggested by Frölich (2007) using five sets of instrumental variables. We calculate with precision the proportion of compliers in our sample, but we cannot retrieve precise LATE estimates due to the small number of compliers.

The evidence suggests that significant implications of migration were already visible by 2000 for most of the migration categories of households: net outflows of migration are associated with statistically significant net gains for both origin households and newly-formed split-off households, while migrant-receiving households see a short-run decrease in living standards (all relative to non-migrant households). By 2007/8, net-exporters of migrants and newly formed split-off households are still found to be positively affected by crisis-period migration and net-importers are found to be no longer negatively affected, suggesting that the declines caused by the increase in household size due to migration did not produce any lasting negative consequences. In addition households that experienced balanced inflows and outflows of migrants showed significant longer-run gains in per capita consumption. These last results suggest that, on average, mobility in times of crisis is beneficial and that the benefits are likely to persist over time: even net inflows of migration into a household do not seem to generate long-lasting negative effects, presumably because further adjustment takes place among these households in the following years (including, possibly, further migration).

In order to check the robustness of the results and to try to isolate some of the ways in which migration had its effects we also conduct the analysis on several different subsamples. With few exceptions, the conclusions from 2000 appear robust to a number of different specifications and they are observed for almost all subsamples analysed in this study. Somewhat less consistency is found for the long run; more precisely, the main results are replicated when the analysis is restricted to the subsample of households that did not move in the pre-crisis period, to the subsample where no new births occurred in the years that followed the crisis, to the subsample of non-rice producers and to the subsample that resided in Java in 1997, arguably the hardest hit region in the country. In each case, these subsamples of households were also more likely than their counterparts to experience migration in response to the crisis. The complementary subsamples to these – namely those households that had already proved mobile in the years prior to the crisis, where new births occurred post-crisis and those that resided outside Java in 1997 – showed less well-defined benefits from their crisis-period migration.

Overall, these results suggest that migration is a positive strategy in times of crisis. It is unlikely to benefit everyone relative to non-migration, but in the short run and on average it generates benefits for the migrants, for the origin households, and for newly created households, whereas existing households that receive migrants apparently suffer relative declines in consumption. In the long run, positive gains persist for several classes of household while no class suffers significant harm on average.

Our paper is organized as follows: in the next section, we review the migration literature. Section 3 presents the data used in the analysis and some descriptive statistics from our data. Section 4 describes our econometric strategy, and Section 5 presents the results. Section 6 carries out some subsample analyses that aim to isolate the crisis in a number of ways. Section 7 concludes the paper.

2. Existing Literature and Purpose of the Study

A large economic literature has developed over the recent years on the importance of geographic mobility in the process of economic development: most notably, in 1997, Lucas defined the issues related to the efficiency of the use of labour and consequences of migration for overall poverty as "of paramount importance" (Lucas 1997, p. 727). In the Asian region, Lucas also showed that there was a steady rise in the proportion of population residing in urban areas in the period between 1950 and 2000. Collier and Dercon (2009) recognized the key role of geographic mobility in the historical experience of most rich economies and the recent experience of fast growing developing Asian economies. They identify five essential characteristics that were key to the economic success of the Asian region: first, a large reduction in the number of people engaged in agriculture; second, a large increase in the population located in urban areas and coastal areas; third, a large reduction of the size of the population living in rural areas, away from the coast or away from urban centres; fourth, a substantial increase in labour productivity in agriculture; and fifth, a substantial increase in overall agricultural production. A clear nexus exists between the first three elements and geographic mobility in the process of economic transformation; the fourth is not necessarily attributable to migration but nonetheless is typically linked to it, since, as discussed by Collier and Dercon, sustained increases in labour productivity have always been strongly associated with the release of labour from the land. Collier and Dercon (2009) conclude that the five characteristics of success – and the migration they imply – will be key for economic development to succeed in the African continent in the next 50 years.

From an historical perspective, therefore, what happened in Indonesia in 1997-8 can be considered as an anomaly. In 1998, the economic scenario in the East-Asian region changed dramatically; the financial (and then economic) shock that hit Indonesia was sudden, spatially heterogeneous and resulted, among other things, in an immediate increase in the migration rates in the country. Fallon and Lucas (2002), in their review of the impact of the financial (and economic) crises of the 1990s on labour markets, household incomes and poverty, note little systematic evidence of internal mobility in response to financial crises. However, they identify Indonesia during the East-Asian crisis as an exception with "massive" net migration to the rural areas. Hence, not only did the economic crisis lead to an increase in the rates of internal migration, but it also resulted in the reversal of the rural-to-urban flows of migration historically linked to economic growth. The combination of these elements and the high geographic mobility in those years makes the East-Asian crisis in Indonesia an ideal case to study the role of geographic mobility as a major strategy for economic adjustment after an economic shock.

Theoreticians have long recognised geographic mobility as a major equalizer and means of economic adjustment, but the existing empirical literature remains patchy. Foster and Rosenzweig (2002) observed that household structure is usually treated as an exogenous factor in much of the empirical and theoretical economic development literature; for this reason, they formulated and tested a structural model of household division, that yields implications for how household size and intra-household inequality interact with exogenous income growth. In 1997, Lucas argued that little empirical evidence was available on the link between structural adjustment and mobility of labour, despite its potential relevance. In 2004, the same author further argued how slight was the evidence on the impact and indirect effects of migration upon inequality. Some recent economic literature has also considered the link between global imbalances and migration by looking at the price of labour: Mortensen (2005) asserts that the theory on earnings and their determination is definitely incomplete, and draws attention to the fundamental empirical fact that there remains substantial heterogeneity across individuals however many controls are included in the wage equation. By doing so, he highlights the need to consider other factors in the analysis of income determination, like sector, size and location of the enterprises where employment takes place. Rosenzweig (2010) likewise found high wage differentials in his recent discussion of the pricing of skills across countries; his results show that variation in skill prices dominates the cross-country variation in schooling levels or rates of return to schooling in accounting for global inequality in the earnings of workers. Hence, Rosenzweig points to the potential of international migration to reduce the global inequality in the earnings of workers worldwide and to increase global efficiency. This conclusion accords with the results from recent simulations that indicate that even relatively small modifications in international migration may have far greater implications for global production than would the complete removal of trade restrictions (Walmsley and Winters, 2005). In his recent review of the literature on labour markets in developing countries, Teal (2011) also points to the crucial role of the price of labour in understanding why people in some countries are so poor relative to others. Further, observing wage differences both across countries and within countries, he argues that the central fact about the price of labour is that it is much more closely related to the place where people live than with what they know as measured by their education. Development, he concludes, has been inextricably linked to movements of labour.

In the case of Indonesia, a number of studies (i.e., Thomas, Beegle and Frankenberg, 2000 and Smith et al., 2002) provide careful analyses of the effects of the East Asian Crisis on the labour market. These studies do not explicitly model the role of migration as an (ex-post) coping strategy (instrument) for Indonesian households during the crisis, but they document the large social and geographic mobility in the country in response to the crisis. Safir and Beegle (2009) attempt to identify the categories of individuals that were most hurt by the crisis and to establish a causal link between the East-Asian economic crisis in the late 1990s and the migration decision in those years; while a significant increase in migration is found over the period, these authors do not find relatively higher migration among the groups of workers that were expected to be mostly hurt by the crisis.

What none of these studies assesses, however, is whether or not (and, possibly, the extent to which) those individuals and families who do move during a crisis are also more able to recover after the crisis and achieve better long-run trajectories. In other words, none of the existing studies assesses the effectiveness of migration as a response mechanism to an economic crisis. This is a relevant question because it mirrors the decision faced by a representative household in a period of economic downturn: whether to produce/receive migration flows while crisis recovery is still ongoing. Moreover, no study to our knowledge has explicitly focussed just on migration as a crisis-response rather than as a once-and-for-all decision; that is, previous studies have not recognised that people can move more than once, or can delay migration decisions until after the crisis-period. In the case of Indonesia, households that experienced migration in the years of the crisis may have experienced it again from 2000-08, while households that did not experience migration during the crisis may have done so from 2000-08.

The Indonesian context looks particularly appealing for rectifying these omissions. We have four observations over sixteen years on most households; the lack of significant correlation between the regional distribution of the crisis and the pre-crisis regional distribution of poverty (Sumarto, Wetterberg, and Pritchett, 1998) makes it less likely that our results are confounded by income-level effects, and the positive relationship between the degree of economic integration of different regions and the impact of the crisis (Ravallion and Lokshin, 2007) means that traditional migration corridors were unlikely to be effective. In such circumstances leaving an origin household during the crisis was riskier than usual and more likely to fail to raise living standards, so that we are less likely than studies of more stable times to identify benefits. In addition, we are also able to analyse the consequences of migration for the destination households of migrants; Thomas and Frankenberg (2007) found that a good deal of the migration that took place during the crisis was return migration which led the poorest from urban areas to move back to rural areas and join family members, to try to exploit the less adverse conditions in rural areas and economies of scale of consumption. It is far from clear that mobility will be to the advantage of these recipient households. Finally, insofar as the immediate effects of mobility in response to the crisis on migrant households' welfare are not obvious, its long-run consequences are even less clear a-priori. These are the issues that we seek to address in this paper.

3. Data and Descriptive Statistics

Our main sources of data are the IFLS (Indonesia Family Life Survey) panel dataset and the National Socio-Economic Survey (SUSENAS) of Indonesia. The IFLS dataset was collected in four waves, namely in 1993, 1997, 2000 and 2007/8. A large amount of information at the individual, household and community level was obtained on a sample representative of around 83% of Indonesian population; the original sample contained over 30,000 individuals from 13 of the 27 provinces in the country. The timing of the survey and the effective tracking of internal migrants (within national boundaries) made this dataset extremely suitable for the present analysis. Following migrants to their destination places and interviewing them there also contributed to the low attrition rate obtained in this dataset: among any part of the original IFLS1 households, 87.6% were interviewed in all waves and 90.3% were either interviewed or died. The SUSENAS, instead, is a national household survey implemented by the government's Central Bureau of Statistics (Biru Pusat Statistics, BPS) and surveys households over the entire national territory. Since 1993, a core questionnaire has been administered yearly, obtaining information on a variety of demographic characteristics of all household members, including their labour market activities and their consumption behaviour. Each survey has a sample size of about 200,000 households and, since 1993, SUSENAS is representative at the level of the kabupaten/kota[3] (i.e., the district). Provinces in Indonesia are further sub-divided into districts; the geographical composition of the districts in Indonesia did not vary significantly in the period from 1993 to 1998. Nonetheless, new decentralisation laws were passed in 1999 and this resulted in splits and changes in the districts during the later years. This, in turn, implies that a substantial increase was observed in the number of districts, that increased from 270 in 1998 to over 350 in 2002 (Ravallion and Lokshin, 2007). Ravallion and Lokshin (2007), that extracted this dataset from the SUSENAS, used a mapping of

[3] Drawing on the SUSENAS, a 10-year panel of district-level data for 1993–2002 was constructed by Ravallion and Lokshin (2007). The complete data set of poverty lines and district-level estimates was made publicly available.

the new kabupaten (i.e., in 2002) into the old aggregates (i.e., in earlier years) that was provided by BPS. This was done in order to compare the same geographical units over time, and it resulted in the availability of information for 297 districts from 1997. This is the information that we extracted from this dataset for our analysis.

In this study, the definition of migration crucially relies on this geographical threshold. An imported migrant is defined as a new household member (not in the household in 1997) who crossed the border of the kabupaten of residence during the 1998-2000 period in order to join the household for six months or more. An exported migrant is defined as an individual who was found in the household in 1997 but no longer resides into the household, who migrated after 1997 for six months or more, and who, by 2000, resides in a different kabupaten. Regrettably, since information on the migration history of individuals was only collected for adults 15 years or older, these definitions (as well as the statistics presented below) ignore any movement of children; for this reason, the actual number of movers, as reflected by these figures, must be underestimated. The reason for imposing a minimum geographical threshold for mobility to be explicitly assessed here is to explore the effectiveness of mobility across regions, which is especially interesting due to the spatial heterogeneity of the economic shock in 1998.

Due to the practice of tracking migrants in the IFLS dataset, the number of observation sampled in the IFLS increased over time; in order to understand the evolution of the IFLS sample from 1993 up to 2000, we first analyse the formation of split-off households without applying any geographical threshold. In the second part of this section, statistics are also presented when the kabupaten threshold is applied to the analysis of migration. In the IFLS dataset, the number of households that were interviewed in 1993 is 7,224. In 1997, 7,620 households were interviewed or re-interviewed; this figure is inclusive of both 1993 origin households as well as new households that split-off from 1993-97. To be precise, in 1997 6,742 origin households (93.3 percent of 1993 households) were re-interviewed; out of them, 5,950 households did not produce any migrants and they did not belong to any extended family with multiple households, whereas 792 households produced migrants and were indeed part of extended families with multiple households. In 1997 an effort was also made to track migrants in their new split-off households; 878 new split-off households were interviewed in 1997, spawning from 792 origin households. Further details of extended families' formation from 1993-97 are provided in table 1.

Table 1 - Extended Families' Formation from 1993-97

Extended Families' Formation before the Crisis		
Origin Families in 1993	No. of Households per family in 1997	Extended Family with Multiple Households by 1997
5,950	1	No
711	2	Yes
76	3	Yes
5	4	Yes
Grand Total		
6,742 (No. of origin households re-interviewed in 1997)	7,620 (No. of origin & split-off households interviewed in 1997)	792 extended families with multiple households

In 2000, 10,435 households were sampled: this number is inclusive of both 1993 origin households as well as new households that split-off from 1993-2000. In 2000, 6,774 origin households (93.8 percent of 1993 households) were re-interviewed; out of them, 4,164 households did not produce any migrants and they did not belong to any extended family with multiple households, whereas 2,610 households produced migrants and were indeed part of extended families with multiple households. These statistics suggest that, in 1997, there were 792 extended families with multiple households, whereas, by 2000, the number of extended families with multiple households increased to 2,610. In addition to re-interviewing origin households, in 2000 an effort was made again to track migrants in their new split-off households; 3,661 split-off households were interviewed in 2000, spawning from 2,610 origin households. This figure is taken from IFLS3 and includes households that split-off both before and during the crisis. Further details on this are provided in table 2.

Table 2 - Extended Families' Formation by 2000

Extended Families' Formation before & during the Crisis		
Origin Families in 1993	No. of Households per family in 2000	Extended Family with Multiple Households by 2000
4,164	1	No
1,837	2	Yes
557	3	Yes
174	4	Yes
34	5	Yes
5	6	Yes
1	7	Yes
1	8	Yes
1	15	Yes
Grand Total		
6,774 (No. of origin households re-interviewed in 2000)	10,435 (No. of origin & split-off households interviewed in 2000)	2610 extended families with multiple households

Our primary interest in this paper is migration during the crisis (and, arguably, in response to it). Out of the 10,435 households interviewed in 2000, 878 were new households that split-off before the crisis (i.e., from 1993-97); these households were considered new split-off households in 1997 (IFLS2) and target households in 2000 (IFLS3). We know from IFLS2+ that 280 new households split-off in the first year of the crisis (i.e., from 1997-1998)[4]; this information is reported here because these households were considered new split-off households in 1998 and target households in 2000 (IFLS3). In 2000, 7,790 target households were interviewed and 2,645 new 2000 split-off households were identified (i.e., households that split-off in the years of the crisis, from 1998-2000); the total number of individuals that were recorded to have moved out of the origin household during the crisis is 9,822. These statistics imply that a total of 2,924 new split-off households were formed over the crisis (and

[4] The low number of new split-off households reported from 1997-98 is also explained by the fact that IFLS2+ only sampled a 25% sub-sample of IFLS households.

interviewed in 2000)[5]. Therefore, from 1993-97, 878 new split-off households formed, an average of approximately 220 per year. From 1997-2000, 2,924 new split-off households formed, an average of roughly 975 per year[6]; these statistics suggest that a considerable increase in households' formation rates took place in Indonesia in the years of the crisis.

Finally, 2,228 extended families with multiple households were created during the crisis, accounting for 5,474 single households (i.e., 5,474 households belonged to these newly formed extended families)[7]; the number of extended families formed in the years prior to the crisis (and found in 2000) is only 683, accounting for 1,839 single households[8]. Even if causation is still difficult to establish, these figures suggest that extended families' formation intensified in Indonesia after the crisis struck. Indeed, some overlapping applies to these groups of households: out of the 2,228 extended families with multiple households formed during the crisis, 288 extended families were already composed by multiple households by 1997. They are still included among "crisis extended-families" since they further enlarged in the years of the crisis, forming new households. The remaining 1,940 extended families were not multiple households' families in 1997 and they enlarged during the crisis. The majority of the split-off households created both in the pre-crisis years and during the crisis was formed by children leaving their parents' household: in 1997, 82 percent of all extended families (653 extended families) were parent-child extended families. In 2000, 2,176 extended families (around 83 percent of the total, 2,610) were in fact parent-child extended families (Witoelar 2005).

Having investigated the dynamics of the IFLS dataset, we next turn our attention to the mobility of adults between kabupaten lasting for at least six months; measuring the effectiveness of this mobility in the years of the crisis is, in fact, the primary goal of our paper.

Mobility of individuals during the crisis did not simply result in the formation of new households; rather, it may have occurred from one IFLS household to another, with individuals joining family members into another target household, without creating any new household; also, IFLS households were often found to receive migrants from other non-IFLS households, thus increasing their size with new individuals that had never been sampled before. Some households did not experience any mobility in the years of the crisis (i.e., nobody arrived and nobody left), but many others engaged in the complex dynamics of mobility in those years, sending and receiving people. Therefore, we gathered up all the information on mobility at the family and individual level presented so far, and we grouped households into six broad categories according to their migration status during the crisis[9]. The first category

[5] One household was labeled as a split-off household both in 1998 (IFLS2+) and 2000 (IFLS3); however, since newly formed households in 1998 (i.e., in IFLS2+) were no longer treated as newly formed households in 2000 (i.e., in IFLS3), this appears as an inconsistency in the data. Hence, in order to avoid double-counting, we only counted this household once in the computation of the overall number of split-off households formed in the years of the crisis.

[6] This average is likely to slightly underestimate the actual degree of mobility observed in those years, since this figure also includes the result from IFLS2+ (i.e., 280 new split-off households in the first year of the crisis), that only covered a fraction of the IFLS respondents. Because IFLS2+ only surveyed part of the IFLS sample, we do not use this IFLS wave directly in the analysis. Also, the fact that split-off households found in the IFLS2+ are treated as target respondents in 2000 explains why target respondents in 2000 are more than the total number of respondents in 1997.

[7] For 5 of these extended families only the split-off household was sampled in 2000; for this reason, the number of extended families with multiple households in the IFLS3 dataset is actually 2,223 (not 2,228).

[8] For 8 of these extended families only the split-off household was sampled in 2000; for this reason, the number of pre-crisis extended families with multiple households in the IFLS3 dataset is actually 675 (not 683).

[9] In our dataset, modifications of the composition of the household due to births and deaths could be distinguished from those resulting from migration. For this reason, the migration categories were constructed purely on the basis of

comprises non-migrant households: their composition was not modified by migration, since no individual was either received from outside the kabupaten or sent out of the household (and outside the kabupaten). Additionally, these households did not move outside their origin kabupaten. The second category includes balanced households: the aggregate size of these households at the end of the period under consideration was the same as it was initially, but they differ from the group of non-migrant households because they received and sent out a number different from zero (and equivalent) of in-migrants and out-migrants[10]. The third category captures net importer households; whether or not these households sent away any members, they were net recipients of migrants, since they received more people from outside the kabupaten than they sent out. The fourth category contains net exporter households; whether or not these households received any members from outside the kabupaten, they were net senders of migrants, since they sent away more people than they received. The fifth category groups new split-off households, which were formed by out-migrants from IFLS target households and settled down in a new kabupaten. Finally, the sixth category includes households that did not send out or receive any individual but, by the end of the period observed, had moved entirely from their kabupaten in 1997.

Table 3 reports information on all households interviewed in 1997 (IFLS2) and 2000 (IFLS3), grouped according to their migration status. An immediate comparison can be made between the statistics relative to the years that preceded the crisis, and those related to the crisis period: the lower proportion of non-migrant households in the 1998-2000 period suggests that mobility across kabupaten intensified in the years of the crisis. Looking at yearly rates of migrant households' formation and assuming that, in both the 1993-97 and the 1998-2000 periods, migrant households' formation is uniformly distributed across years within each period, these figures suggest even more strongly that mobility between kabupaten intensified in the years of the crisis: for example, formation of new split-off households seems to have grown significantly from 1998-2000 in comparison with previous trends. Even for net migrant exporters, that constituted 20.01% of households between 1993-97, and only 16.20% of families between 1998-2000, the yearly rate of formation from 1998-2000 is actually 3.1 percentage points higher than the rate from 1993-97.

Clearly, one may argue that it is difficult to disentangle the time trend from the trend due to the ageing of individuals in the sample or from the effect of the crisis. However, given its magnitude, we interpret the increase in mobility rates after 1997 as high and unusual; at least in part, we interpret this increase in the rate of mobility as attributable to the crisis. This conclusion is supported by a substantial empirical literature from Indonesia, that concludes unanimously that mobility intensified in the country in the years of the crisis. Among others, Gilligan et al. (2000) use two waves of the National Farmers Household Panel Survey (PATANAS) from 1994/95 and 1998/99 to compare monthly migration rates in the two periods from six provinces in Indonesia. They conclude that the monthly migration rate of individuals coming from Jakarta or the capital city of the province or district in which a household is located doubled in the period between August 1997 to April 1999, in comparison with the pre-crisis period (October 1995 to July 1997); also, they find that rural households that benefitted

migration flows of adults, and are not affected by births and deaths. Also, since the migration histories were not available for individuals 15 years old or younger, child emigrants or immigrants are not counted, and the migratory flows of children are excluded from the analysis. This does not threaten our classification of households according to the migration decision of their adult members, because adult emigrants and immigrants could be clearly identified from the data.

[10] Hence their composition remained unchanged only at the aggregate, but not at the individual level.

Table 3 – Migration Status of Target and New Split-Off Households from 1993-1997 and 1998-2000

Import-Export of Migrants and New Households' Formation						
Migration Status	Before the Crisis: 1993-97			During the Crisis: 1998-2000		
	No. of Hhlds	Fraction of Hhlds	Yearly Rate of Migrant Hhlds Formation	No. of Hhlds	Fraction of Hhlds	Yearly Rate of Migrant Hhlds Formation
Non Migrants' Household	5,333	69.99%	-	6,899	66.11%	-
Balanced Migrants' Household	85	1.11%	0.28%	151	1.45%	0.73%
Net Migrants' Importer Household	264	3.46%	0.86%	499	4.78%	2.39%
Net Migrants' Exporter Household	1,525	20.01%	5%	1,690	16.20%	8.10%
New Split-Off Household	291	3.82%	0.95%	1,061	10.17%	5.09%
Moving Household	122	1.60%	0.4%	135	1.29%	0.65%
Grand Total	7,620	100%	-	10,435	100%	-

from the crisis, tree-crop growers, were more likely to host new household members after the crisis. Evidence of urban to rural migration was also found by Breman (2000), using data collected between November 1997 and February 1998 in West Java.

As noted above, Fallon and Lucas (2002) recognised Indonesia as generating "massive" migration during the crisis; Frankenberg et al. (2003) also describe migration of individual household members and changes in household size and composition at the onset of the crisis as key factors in the response of Indonesian households to the crisis. They use data from the IFLS2 and IFLS2+ subsample (sampling approximately a quarter of the full IFLS sample) and they present evidence from the first year of the crisis. Their analysis reveals that significant changes in household size and composition occurred between 1997 and 1998; the evidence points towards the fact that changes in household composition were associated with smoothing during the crisis. Urban households at the bottom of the pre-crisis PCE distribution tended to lose household members, whereas household size increased over the entire distribution of pre-crisis PCE in rural areas and a positive correlation was observed between the pre-crisis PCE and the rate of growth of household size in that year. The authors conclude that, in response to the crisis, Indonesian households and families modified their living arrangements, as dependents generally moved to lower-cost locations and working age family members joined households that were able to absorb additional workers. However, a general problem with this analysis is that the authors do not distinguish changes in household size due to geographic mobility from changes in household size due to births and deaths. Thomas and Frankenberg (2007) further argue that one reaction to the crisis was a modification of living arrangements, as family members moved in together to exploit economies of scale of consumption. The increase in household size was greater among households in rural areas, mirroring both the effect of households joining together within the rural sector and the migration of individuals from urban areas to join households in rural areas. Also, they argue that migration took place from the poorest urban households to join households in rural areas.

Finally, evidence of the intensification of migration flows in Indonesia is also provided by Safir and Beegle (2009); these authors look at the proportion of 1993 IFLS interviewees who migrated in each year between 1994 and 2000, among individuals residing in 1993 in the urban areas of the provinces covered in IFLS2+. They find evidence of a major increase in migration rates after 1997; in the attempt to reduce the effect of migration linked to aging, they also restrict the observation to sub-samples of individuals that were 25-50 years old, and individuals that were 30-50 years, in 1993. They conclude that migration increased substantially for both these sub-groups between 1996 and 1997. Hence, to conclude, although the causal impact of the economic shock on geographic mobility is not explicitly modelled in this study, our analysis builds on the existing empirical literature that suggests consistently that rates of mobility were unusually high in Indonesia in the years of the economic crisis.

Before moving to the econometric strategy of this paper, a few more descriptive statistics are provided. Out of the 10,435 households sampled in IFLS3 (2000), only 10,180 actually responded to the questionnaire. Of these, only 9,799 were located in 1997 in kabupaten sampled by the SUSENAS; this is relevant because the SUSENAS dataset is used to extract information at the kabupaten level used in the econometric analysis. Of the 10,180 households sampled by IFLS in 2000, 9,911 reported complete information on the p/capita monthly HH expenditure from 1997 and 2000, and 8,405 reported this information from 1997 and 2007. This resulted in having all necessary information from 1997, 2000 and 2007 for 8,024 households. Grouping these households by migration category, Table 4 shows the average values of the p/capita consumption indicators used to assess their well-being. Interestingly, from Table 4 it seems like net-exporter households were broadly poorer than other migrant categories in 1997, especially as compared to net importers and moving households. This is consistent with Thomas and Frankenberg (2007), as more people seem to leave from poorer households and join better-off households. Also split-off households seem to come from relatively wealthy families (in fact,

Table 4 – Key Welfare Indicators, by Migration Status of households included in the econometric analysis

Averages of key welfare indicators - by Migration Status	N (total=8,024)	Ln P/capita monthly HH expenditure 1997	Ln P/capita monthly HH expenditure 2000	Ln P/capita monthly HH expenditure 2007	97-00 Growth Ln P/capita monthly HH expenditure	97-07 Growth Ln P/capita monthly HH expenditure
No Migration	5,561	11.532	12.247	13.382	0.716	1.850
Balanced	124	11.608	12.320	13.508	0.713	1.900
Net Importer	412	11.786	12.258	13.522	0.472	1.736
Net Exporter	1,339	11.599	12.394	13.520	0.795	1.921
Split-Off	512	11.833	12.712	13.674	0.879	1.841
Moving	76	11.915	12.516	13.618	0.601	1.702

for these households the data relative to 1997 applies to the origin household of the newly created split-off). By 2000, the level of p/capita consumption of net exporters is greater than that of net importers, arguably as a consequence of the modification in living arrangements; however, by 2007, the levels of p/capita consumption seem remarkably similar between these two categories.

4. Econometric Procedure

In the econometric analysis, we use the same consumption-based measures of living standards, and we try to assess the extent to which these observed differences among migration categories can be interpreted as causal. In other words, we evaluate the impact of migration on the evolution over time of this measure of well-being. Hence, our goal was to estimate the actual (net) effect of migration in Indonesia on consumption growth, respectively during the crisis and ten years later, for those households where migration occurred in the years of the crisis. In order to do so, difference-in-difference matching estimators were used (Heckman et al., 1997, 1998; Abadie, 2005); by conditioning on a rich set of variables, the econometric analysis aims to estimate the ATET (Average Treatment Effect on the Treated) of migration in the years of the crisis on, respectively, the short-run and long-run change in consumption of migrants' households. The decision to model the crisis-period migration choice, as opposed to migration over the entire period under observation, in turn implies that, in our control group, households are included that may have experienced migration from 2000-07/8, before being re-interviewed in the most recent wave of the IFLS panel dataset. By doing so, we avoid treating migration as a once-and-for-all decision; the choice of this comparison group better mirrors the choice actually faced by Indonesian households during the crisis; also, this strategy allows to postulate not on the effect of migration per se, but rather on the effect of migration as an immediate response to an economic crisis. At the end of this section we also discuss why we prefer matching methods to the more usual instrumental variable regression approach to problems like ours.

Given that a number of $(M + 1)$ mutually exclusive states can be observed for the observations in our analysis (i.e., a household can fall into six different migration categories), our econometric strategy relies on the contribution from Imbens (1999) and Lechner (2001), that demonstrate that properties similar to the propensity score property valid in the binary treatment framework hold also in a multiple-treatment framework. In particular, for the multiple-treatment model, Lechner (2001) suggests a matching estimator that is as analogous as possible to the algorithms used in the literature on binary treatment evaluation. Following the notation in Lechner (2002), in a world with $(M + 1)$ mutually exclusive states, the potential outcomes are denoted by $\{Y^0, Y^1, \dots, Y^M\}$, and for every observation only one of these outcomes is observed. Participation in a particular treatment is indicated by the variable $S \in \{0, 1, \dots, M\}$. To take into account all possible treatments, the definitions of average treatment effects developed for binary treatments need to be modified[11]. Focusing on a pairwise comparison of the effects of treatments m and l for the participants in treatment m, Lechner (2002) shows the multiple-treatment version of the ATET, which is the parameter usually estimated in evaluation studies (and that is estimated in the present study):

[11] In line with Lechner (2002), it is assumed for the rest of the discussion that the typical assumptions of the Roy (1951)-Rubin (1974) model are fulfilled (e.g., see Holland, 1986 or Rubin, 1974); hence, dependence or interference between individuals is ruled out.

$$\theta_0^{m,l} = \mathrm{E}\big(Y^m - Y^l | S = m\big) = \mathrm{E}(Y^m | S = m) - \mathrm{E}\big(Y^l | S = m\big),$$

where $\theta_0^{m,l}$ indicates the expected effect for an observation randomly drawn from the sample of participants in treatment m (Lechner 2002, proposition (1)).

The Roy (1951)-Rubin (1974) model states that identification is obtained under the Conditional Independence Assumption (CIA), an untestable assumption that the participation in the treatment and the treatment outcome is independent conditional on a set of observable attributes. Imbens (1999) and Lechner (2001) consider identification under the multiple-treatment version of the CIA; this version of the CIA requires all potential treatment outcomes to be independent of the assignment mechanism for any given value of a vector of characteristics, X, in an attribute space, χ. They formalize the CIA in the multiple treatment framework as follows:

$$Y^0, Y^1, \dots, Y^M \perp\!\!\!\perp S \,|X = x, \forall x \in \chi \,,$$

where $\perp\!\!\!\perp$ indicates independence (Lechner 2002, proposition (2)). Similarly to the binary treatment framework, the common support condition also needs to be satisfied in this case too; this requires that a probability of each treatment to occur be found, for all values of $x \in \chi$. Rubin (1977) and Rosenbaum and Rubin (1983) also demonstrate for the binary treatment framework that, in fact, it is sufficient to condition on the propensity score, i.e., the probability of participation conditional on the observed attributes of the observations. Imbens (1999) and Lechner (2001) show that similar properties apply to the multiple-treatment framework as well. In particular, for the ATET, Lechner (2002, proposition 3) shows the following:

$$\theta_0^{m,l} = \mathrm{E}(Y^m \mid S = m) + E_{p^{l|ml}(X)}\left[\mathrm{E}\big(Y^l \mid P^{l|ml}(X), S = l\big) \,|S = m\right];$$

$$P^{l|ml}(x) := P^{l|ml}(S = l | S = l \text{ } or \text{ } S = m, X = x),$$

where $P^{l|ml}$ are the estimated conditional probabilities on the subsample of participants in m and l; matching on the propensity score $P^{l|ml}$ gives a consistent estimator of the counterfactual mean $\mathrm{E}\big(Y^l | S = m\big)$ in order to implement causal analysis. This proposition suggests that standard nonparametric methods, generally used in the binary treatment framework that condition on an estimated propensity score, can be applied here as well. Further, this also implies that, in order to identify $\theta_0^{m,l}$, only information from the subsample of observations in m and l is sufficient. Based on this conclusion, in our main specification, we treat the multiple-treatment framework similarly to the usual binary framework, and, for each bilateral comparison m and l between non-migrant households and households in any of the categories of households affected by migration, we estimate the binary conditional probabilities $P^{l|ml}$ and use such estimated probabilities to match our observations and implement causal analysis. To be specific, we use observations in the non-migrant category to create a counterfactual for each of the remaining migration categories. An alternative option, also suggested by Lechner (2001), is to specify and estimate a multinomial choice model (e.g., using a flexible multinomial probit), then compute $P^{l|ml}$ from that model and, finally, match observations using these estimated probabilities. In fact, we did this too to check the robustness of our results; in line with the evidence in Lechner (2002), the results from the two alternative strategies are substantially consistent[12].

Hence, following the notation in the evaluation literature, for each bilateral comparison, we defined a binary assignment indicator, D, indicating whether any mobility took place among members of the household; let $D_i = 1$ if mobility occurred and $D_i = 0$ otherwise. The outcome variable was defined as the logarithm of the final level of per capita consumption in the household (i.e., in 2000 and in 2007/8, respectively) minus the logarithm of the starting level of per capita consumption in the household (i.e., in 1997). For each household i, we defined two potential outcomes by treatment status,

$$Y_i^0 = \log(C_{i0t}) - \log(C_{i0t'}) \text{ for } D_i = 0$$

and

$$Y_i^1 = \log(C_{i1t}) - \log(C_{i0t'}) \text{ for } D_i = 1,$$

where t and t' represent, respectively, the years 2000 and 2008 (i.e., short run and long run; in both cases, post-treatment) and 1997 (pre-treatment). After dropping the i subscript (for ease of exposition), our purpose was to investigate the migration effect on the treated (migrant households):

$$ATET = E(Y^1 - Y^0 \mid D = 1) = E(Y^1 \mid D = 1) - E(Y^0 \mid D = 1)$$

Consumption growth after the occurrence of migration was observed for migrant households, hence $E(Y^1 \mid D = 1)$ could be estimated. What was not observed, on the contrary, is the growth of consumption of migrant households had migration not occurred, i.e. the counterfactual consumption growth, $E(Y^0 \mid D = 1)$. For this reason, we used Propensity Score Matching (PSM) to estimate such counterfactual. For such procedure to be valid, as aforementioned, certain assumptions need to hold. The fundamental assumption the PSM relies on is known as the Ignorable Treatment Assignment Assumption (Rosenbaum and Rubin, 1983) or the Conditional Independence Assumption (CIA, Lechner, 2000). Since, in the present study, we are estimating the migration effect on the treated (ATET for migrant households) we do not need CIA to hold for Y^0, but only for Y^1. This is because, by calculating the ATET of migration for migrant households, we need Y^0 of matched non-migrant households to constitute a valid counterfactual for Y^0 of matched migrant households (i.e., what migrant households would have experienced in the absence of migration). Instead, we do not need Y^1 of matched migrant households to constitute a valid counterfactual for Y^1 of matched non-migrant households (i.e., what non-migrant households would have experienced in the presence of migration). In this case, the difference-in-difference (DiD) matching estimator (Heckman et al., 1997, 1998; Abadie, 2005) requires

$$E(\log(C_{0t}) - \log(C_{0t'}) \mid X, D = 1) = E(\log(C_{0t}) - \log(C_{0t'}) \mid X, D = 0),$$

where X is an appropriate set of observable variables unaffected by the treatment. This assumption is stated in terms of the before-after consumption evolution instead of levels; this means that, conditional on X, the migrant households' potential consumption growth Y^0 had they not experienced migration would be the same as non-migrant households' potential consumption growth. We decided to use a DiD

[12] Given their similarity, we do not report the results from the multinomial choice specification here for reasons of space.

matching estimator (instead of a cross-section (CS) matching estimator, where the level of consumption in 2000 would be defined as the outcome variable) primarily because the DiD matching estimator only requires the CIA assumption to hold once unobserved time invariant (separable) components that affect both consumption and the migration choice have been differenced out. The CS matching estimator, on the other hand, requires the CIA to hold without removing such time invariant (separable) components. Indeed, it seems to us this is a relevant issue in the present study: if unobservable characteristics influenced in a random way the decision to produce a migrant, the PSM would allow contrasting consumption levels of similar households, the only difference, in principle, being the production of a migrant in response to the crisis. However, since PSM only matches migrant households and non-migrant households based on observable characteristics, and self-selection coming from unobservable characteristics is indeed likely to apply to the migration decision, this approach would not be able to address the bias coming from time invariant (separable) components. The DiD matching estimator is preferable, insofar as it nets out the effects of any additive factors (whether observable or unobservable) that have fixed (time-invariant) impacts on consumption (such as ability) or that reflect common trends affecting migrant and non-migrant households in the same way (such as changes in prices or weather; Ravallion, 2005). Moreover, Smith and Todd (2005) also find the DiD matching estimator to perform better than the CS matching estimator when treated and nontreated were drawn from different regional labour markets; building on the evidence on the spatial heterogeneity of the pre-crisis economic conditions and, more importantly, on the spatial diversity in the impact of the crisis (e.g., Ravallion and Lokshin, 2007), we thought it preferable to use DiD matching estimator.

Unlike IV estimators, matching estimators do not require inclusion in the participation equation of variables that affect the migration decision but not consumption growth in the absence of migration; it is in fact important not to include variables that might be considered powerful instruments for migration in matching exercises, because instruments can reduce the quality of the match (Bhattacharya and Vogt 2007). Different moving costs or, more simply, different Y^1 (i.e., different payoffs from migration) are both likely to satisfy these criteria, and for this reason they were not included in the participation equation. It is important to stress again that, since we are estimating the ATET for migrant households, we need Y^0 to be independent of D after conditioning on X, but we do not need Y^1 to be independent of D after conditioning on X; in other words, we do not need Y^1, the payoff from migration, to be the same for matched migrant as for matched non-migrant households. In our analysis, we argue that, conditional on X and after removing the time-invariant components, migrant households would have observed the same consumption growth as non-migrant households had no migration taken place, but we do not need to argue that non-migrant households would have observed the same consumption growth as migrant households had migration taken place for them.

Finally, to identify the ATET, the following common support condition must be satisfied too:

$$\Pr(D = 1 \,|X) < 1$$

This is a common support condition, that requires a positive probability of observing nonparticipants at each level of X. This assumption would be violated if, at $X = X_0$, only movers and no non-migrant households were observed, i.e., if $\Pr(D = 1 \,|X = X_0) = 1$. However, as we show in the next section, the common support constraint does not constitute a major issue in our data.

With these considerations in mind, the analysis was carried out at the household level and, in the first stage, the likelihood to fall into one of the migration categories presented in table 3 in the years of the crisis[13] was estimated. In order to do so, our dependent variable captures dynamics of migration flows from 1998-2000 (i.e., including the years of the crisis, 1998-2000); in its functional form, the model was estimated with a series of probit estimators, where the dependent variable was a dummy variable taking up value 1 if the household belonged to one of the 'migrant households' categories and 0 if the household belonged to the 'non-migrant households category' (i.e., five different probit models were estimated: in the first model, the dependent variable took up value 1 if the household was grouped among 'Balanced Migrants' households, and 0 if it was grouped among 'Non Migrants' households; in the second model, the dependent variable took up value 1 if the household was grouped among 'Net Migrants' Importer' households, and 0 if it was grouped among 'Non Migrants' households; the same was done for the remaining categories of households). In the estimation of the participation equation, we exploited the panel dimension of our dataset and we used information measured in the pre-treatment period, i.e., from 1993 and 1997. The migration status of household h at time $t+1$ (i.e., in 2000) was modelled on a set of relevant information calculated at time t (i.e., in the pre-crisis period), likely to influence both the migration decision and the growth of consumption. Hence, the first stage equation can be written as follows:

$$M_{ht+1} = g_0 + g_1 X_{ht} + g_2 V_{kt} + \varepsilon_{ht},$$

where M_{ht+1} is a discrete variable that captures the migration status of household h at time $t+1$ (for $t = 1997$, M_{ht+1} captures migration in the years of the crisis); in each binary comparison between non-migrant households and the remaining migration categories, this variable took up value 0 for Non Migrants' households and value 1 for all other categories of households (i.e., namely, Balanced Migrants' households, Net Migrants' Importer households, Net Migrants' Exporter households, New Split-Off households and, finally, Moving Households). X_{ht} is a vector of individual and household characteristics at time t that may affect Y^0 as well as M_{ht+1}, while V_{kt} is a vector of information at the district (kabupaten) level at time t that may affect Y^0 as well as M_{ht+1}, such as infrastructures' availability. In principle, having controlled for unobservable fixed effects at the household level, this should represent an appropriate set of observable variables, unaffected by the treatment, likely to affect both the migration decision and consumption growth in the absence of the treatment. For households that are similar in all individual and local characteristics described, the decision to produce/receive migrants, or to move entirely, in those years, rather than staying and reacting in a different way needs to be random, in the sense that it depends on factors unrelated to future potential outcomes in the absence of the treatment (i.e., Y^0). For us to retrieve a causal parameter, this is in fact required by the CIA; the plausibility of the CIA in our study and the choice of the conditioning variables are discussed in detail in Appendix III.

In the second stage of the analysis, once the predicted probabilities for the migration decision were estimated, households were matched across migration categories. Different matching methods were considered to construct the counterfactual $E(Y^0 \mid D = 1, p(X) = \hat{p}(x))$. Since the conditioning and outcome variables, $\{x_i, Y_i^1\}_{i=1}^{N_1}$ and $\{x_j, Y_j^0\}_{j=1}^{N_0}$, are observed for the two groups respectively and the

[13] Categories presented in Table 3 are the following: Non Migrants' Household, Balanced Migrants' Household, Net Migrants' Importer Household, Net Migrants' Exporter Household, New Split-Off Household and Moving Household.

propensity scores $\{\hat{p}(x_i)\}_{i=1}^{N_1}$ are estimated for the group of migrant households and $\{\hat{p}(x_j)\}_{j=1}^{N_0}$ are estimated for the remaining households, matching could be performed. A number of matching techniques have been suggested recently. Traditional matching estimators include the nearest neighbour estimator, that uses only the closest observation (in terms of the propensity score $\hat{p}(x)$) in the non-treated group to construct the counterfactual consumption growth for a migrant household; however, using only one non-migrant household is likely to be inefficient. Heckman et al. (1997, 1998) suggest local polynomial matching estimators, comprising kernel matching and local linear matching. Hirano et al. (2003) suggest a weighted estimator, where the inverse of a nonparametric estimate of the propensity score act as weights; such methodology has been recently applied by Chen, Mu and Ravallion (2009) and Yamauchi and Liu (2012). Frölich (2004) finds that local linear regression matching performs well when the control-treated ratio is relatively high; he also proposes matching based on a local linear ridge regression, which is found to work well in his Monte Carlo experiments, is robust to simulation designs and, often, is found to outperform other estimators, including nearest neighbour, kernel and local linear matching. Finally, Ham et al. (2011) apply the local linear ridge regression technique suggested by Frölich (2004); for the first time in the matching literature, they also use local cubic matching, which offers the possibility of decreasing the bias of the estimates – although it increases their variance.

After performing matching in the second stage of the analysis, difference-in differences estimates were calculated in the consumption levels before and after the crisis comparing migrant and non-migrant households. All expenditure aggregates used in this analysis were per capita (hence not directly affected by changes in household size) and they were spatially deflated in all waves of the panel dataset; indeed, spatial deflation of the expenditure aggregates was necessary in order to guarantee the comparability of expenditure aggregates across regions[14]. In order to spatially deflate the expenditure aggregates, poverty lines were found in all periods under observation (namely, 1997, 2000 and 2007) at the provincial level, distinctively for urban and rural areas, and they were all indexed and made comparable (all poverty lines were expressed as a fraction of the Jakarta poverty line). In order to finally obtain the deflated (real) value of consumption in each region, consumption aggregates in nominal values were then divided by the relative indexed value. On the other hand, temporal deflation was not necessary given the aims of this analysis and it was not implemented; this is because the analysis is a Diff-in-Diffs analysis, where the change in consumption over time for some individuals is compared to that of some other individuals; once price dispersion across regions is accounted for, whatever the inflationary trends were over time (1997-2007), they are not relevant in the determination of the result of interest. Also, standard bootstrap techniques were implemented in order to calculate standard errors for the matching estimators, in line with a large applied economics literature[15].

---

[14] Indeed, in the absence of any spatial deflator and in the presence of only nominal values of expenditure, any potential correlation between the movements of migrants and the inflationary trends at the regional level over time may have questioned the validity of the analysis. On the contrary, by spatially deflating the aggregate nominal values of consumption, meaningful comparisons could be carried out on the levels of expenditure of households located in different regions.

[15] Abadie and Imbens (2008) showed that, in general, the use of the bootstrap is not valid for nearest neighbour matching, even when the estimator is root-N consistent and asymptotically normally distributed, with zero asymptotic bias. In our analysis this was not a problem, since a range of estimators were used, the performances (in terms of mean-bias and median-bias minimization) of each of them were compared before using the bootstrap

Finally, we have preferred to use matching methods to tackle our problem rather than the more common panel econometric regression model. This is because the latter requires a number of unattractive assumptions and features. First, a linear regression analysis regressing consumption growth on the migration decision and other control variables measured in 1997 would suffer from potential biases when the implied functional form assumptions are not satisfied[16]. This is especially relevant because these assumptions in turn imply that the effects have to be homogeneous in the population or specific subpopulation (see for example Heckman, Smith, and LaLonde, 1999). In a parametric regression, the estimand (i.e., the treatment dummy) is the ATE; in the case of homogeneous effects, the ATE and the ATET coincide. However, in the case of heterogeneous treatment effects (i.e., ATE≠ATET), a regression that uses both the treated and untreated units but does not interact the treatment dummy with the covariates does not estimate the mean impact of treatment on the treated, but rather a different weighted average of the treatment effects (see Angrist and Pischke (2008) for more details). Such assumptions are clearly undesirable in this context. In the assessment of the effect of migration, it is not clear why one would want to investigate the ATE; to the extent that people decide to move on a voluntary basis and they do so in a rational way, then really the ATET constitutes the more interesting question (i.e., exploring the effect of migration on the trajectories of movers, rather than the effect of migration on the trajectory of an individual taken randomly from population). Matching estimators circumvent these problems, since they are semiparametric and allow for arbitrary individual effect heterogeneity.

The second reason that motivated our choice of matching techniques is the key assumptions that the fixed effect, i.e. the part of the error that is allowed to be correlated with the regressors and captures potentially unobservable confounders, has a constant effect on the outcomes over roughly 15 years; this assumption would be very difficult to justify in this context (especially because of the occurrence of the crisis itself).

Finally, one further alternative to identify the effects of migration in response to the crisis would be to use a parametric instrumental variable approach (e.g. Imbens and Angrist, 1994). However, such an approach requires the set of unattractive assumptions discussed above. Additionally, it also requires an exogenous variable that influences consumption growth over time only by influencing the migration decision, whereas any direct effects are ruled out. In the migration literature such a variable seems very hard to find and, in actual fact, a convincing instrument does not appear to be available in the present context. Bearing these considerations in mind, we did not implement a parametric instrumental variable analysis; nonetheless, we investigated the presence of potential instruments for migration in our data, and we estimated the semi-parametric local average treatment effects (LATE) estimators with covariates suggested by Frölich (2007) using five sets of instrumental variables. The remainder of the paper presents the results of our analysis.

and, finally, the bootstrap was only applied to the best-performing estimator. The nearest neighbor estimator was never found to perform best in this study, hence the bootstrap was not applied to it.

[16] We have, in fact, estimated this model but do not report the results because they are similar to our preferred approach.

5. Empirical Results

a. Propensity score specification

Building on the theoretical literature on the determinants of migration, in the first stage of the analysis we estimated the migration decision in 1998-2000 as a function of a set of variables that are summarized in table 5; variables were broadly categorized as variables at the microeconomic level and variables at the macroeconomic level; the former were taken from the IFLS dataset and provide information at the household level, whereas the latter were extracted from the SUSENAS and provide information at the kabupaten level (the geographical unit of observation for the migration analysis).

Table 5 – Determinants of Migration in Indonesia in response to the East-Asian crisis – Description of Covariates

Description of Covariates included in the estimation of the Propensity Score of Migration	
Microeconomic factors (household level)	Macroeconomic factors (kabupaten level)
• Household characteristics (e.g., household size, occurrence of births/deaths in the pre-crisis period, rural/urban household, own production of rice in 1997, presence of young members, i.e. < 25 years old, presence of recent graduates) • Household head's characteristics (e.g. age, gender, marital status, religious faith and education level) • Migration history of the household between 1993-97 (e.g., whether the household was affected by migration between 1993-97) • Farm ownership, i.e., whether the household owns entirely a farm • Access to credit market in the last year, i.e., if any household member borrowed money in the year prior to the survey • Working status of the adult members of the household (i.e., presence of self-employed individuals, government workers, private workers and family workers) • Average health status (e.g., the average health status as reported by adult household members in 1997) • Presence of relatives away (e.g., presence of parents away from the household for any member, or presence of siblings away from the kabupaten)	• Presence of health facilities • Proportion of urban households • Unemployment rate • Proportion of self-employed individuals • Proportion of individuals employed in the agricultural sector • Proportion of households living under the local poverty line • Gini coefficient

The inclusion of this relatively large number of covariates in the specification aims to satisfy the CIA: if, after controlling for this set of covariates, there are no other factors that influence the choice of the different values of M_{ht+1} as well as the outcomes that would be realized in the absence of migration, this comparison produces a causal effect, namely the average treatment effect on the treated (ATET). The plausibility of the Conditional Independence Assumption is required for this, and its plausibility relies on the richness of the information available; for this reason, in applications exploiting the CIA, X is generally of high dimension, as many covariates are often needed to make the CIA assumption plausible (Huber, Lechner and Wunsch, 2013).

Table 6 presents the evidence on the determinants of the migration decision between 1998-2000; results are presented for all migration categories, and in each case the dependent variable was a binary variable taking up value 1 if the household belonged to the migration category of interest, and 0 if it belonged to the non-migrants household category. In its functional form, the equation was estimated as a probit model in each case, and estimates were calculated only on the sample of 8,024 households that reported information in all periods of the analysis (i.e., households for which control variables could be observed in 1997, and for which the level of per capita consumption could be observed in 1997, in 2000 and also in 2007/8); the same is valid for all the econometric analysis, where all results were calculated on the sample that reported at all times under consideration. In fact, this resulted in the loss of some observations, but it was rendered necessary by the fact that, when comparing the short run and the long run effectiveness of a treatment, the direct comparability of these results is implicit in the analysis; if different households were used to assess the short run and the long run effectiveness of migration in response to the crisis, then these estimates may not be directly comparable.

The evidence suggests that larger households were generally less likely to move as a whole, but they were more likely to observe changes in living arrangements in the years of the crisis; on the other hand, the opposite is valid for households where births took place in the years prior to the crisis. Male-headed households proved less likely to send out migrants, or receive them; arguably, this is best explained by the fact that the migrant in the household is more likely to be the male-head himself, rather than his wife. Age of the head is an inverse U-shaped function of the likelihood to produce/receive migrants, as well as to form new split-off households; instead, education of the head is positively related to the receipt of migrants and to the formation of new split-off households. The presence of young members in the household in 1997 generally made migration during the crisis more likely to occur; migration in the years preceding the crisis were also found to be strong (and positive) predictors of the migration decision between 1998-2000.

The presence of workers from different types of employment (i.e., presence of self-employed workers, or of government workers) did not seem to affect positively the likelihood to experience changes in living arrangements; this was unexpected because different types of Indonesian workers were hit differently by the crisis[17]. More healthy households were generally more likely to experience modifications in living arrangements; similarly, presence of siblings away from the kabupaten for some household members increased the chances to experience migration in those years. Access to the formal credit market did not seem to play any role in the determination of the migration decisions. Finally, from the information collected at the kabupaten (macroeconomic) level, it is mostly interesting to

[17] Smith et al. (2002) show that real hourly earnings of the self-employed in urban areas and of female self-employed in rural areas dropped by at least as much as real market sector wages, whereas real hourly earnings of self-employed males in rural areas remained fundamentally stable.

notice that split-off households formed during the crisis tended to come from wealthier kabupaten, whereas net-exporters and moving households were more likely to be found in more unequal regions (as indicated by the Gini coefficient at the kabupaten level). Linking this result with the evidence that initially better-off and more unequal kabupaten were more vulnerable to the crisis (Ravallion and Lokshin, 2007), this may reflect the fact that the economic crisis further encouraged formation of new split-off households in those years and departure of household members (or of entire households).

Table 6 – Evidence on the Determinants of Migration in Indonesia in response to the East-Asian crisis (1998-2000)

Probit Estimates Migration Decision	Balanced Households	Net-Importer Households	Net-Exporter Households	Split-Off Households	Moving Households
Household Size	0.032	-0.007	0.047	0.062	-0.253
	(1.92)*	(0.58)	(5.49)***	(5.91)***	(5.36)***
Births 1993/97	-0.218	-0.189	-0.157	-0.131	0.398
	(2.34)**	(3.07)***	(4.22)***	(2.58)***	(3.85)***
Deaths 1993/97	-0.127	0.048	-0.027	-0.195	0.097
	(1.02)	(0.62)	(0.49)	(2.33)**	(0.45)
Rice Self-Product	0.005	-0.051	-0.068	-0.010	-0.090
	(0.04)	(0.74)	(1.44)	(0.15)	(0.57)
Male head	-0.157	-0.270	-0.201	0.084	-0.013
	(0.91)	(2.39)**	(2.41)**	(0.68)	(0.05)
Age of Head	0.004	0.005	0.006	0.067	-0.008
	(0.81)	(1.96)**	(3.29)***	(4.65)***	(0.29)
Age of Head (sq)	-0.000	-0.000	-0.000	-0.001	-0.000
	(0.27)	(1.59)	(3.02)***	(3.85)***	(0.01)
Married Head	-0.157	0.156	0.105	-0.105	0.199
	(0.88)	(1.33)	(1.22)	(0.85)	(0.78)
Muslim Head	0.027	-0.020	-0.049	0.186	0.122
	(0.19)	(0.21)	(0.75)	(1.88)*	(0.55)
Head Educ – Elem	0.175	0.050	0.080	0.150	0.089
	(1.29)	(0.65)	(1.46)	(1.83)*	(0.42)
Head Educ – High	0.282	0.052	0.141	0.295	0.102
	(1.78)*	(0.56)	(2.13)**	(3.12)***	(0.45)
Head Educ – Univ	0.594	0.346	0.266	0.489	0.129
	(2.65)***	(2.37)**	(2.45)**	(3.46)***	(0.42)
Young Members	0.392	-0.123	0.351	0.194	0.326
	(2.28)**	(1.49)	(4.84)***	(1.96)**	(1.67)*
Migrant 93/97	0.764	0.659	0.551	0.633	0.501
	(8.74)***	(11.72)***	(13.84)***	(11.79)***	(4.07)***
Health Facil Kabup	-0.450	-0.141	-0.043	0.411	-0.350
	(1.54)	(0.64)	(0.24)	(1.15)	(1.02)
Recent Graduates	0.144	0.087	0.072	0.110	-0.014
	(1.56)	(1.38)	(1.71)*	(1.90)*	(0.09)
Self-Employed	-0.118	-0.014	-0.055	-0.042	-0.299
	(1.17)	(0.23)	(1.24)	(0.68)	(2.18)**
Public Worker	-0.361	-0.128	-0.131	-0.173	-0.262
	(2.35)**	(1.39)	(2.05)**	(2.07)**	(1.37)
Private Worker	-0.152	-0.108	-0.073	-0.106	-0.170
	(1.55)	(1.75)*	(1.69)*	(1.77)*	(1.27)
Family Worker	-0.099	-0.130	0.013	-0.046	-0.683
	(0.82)	(1.64)	(0.26)	(0.61)	(2.11)**
Healthy hhld	0.340	0.016	0.224	0.189	-0.203
	(2.03)**	(0.20)	(3.56)***	(2.11)**	(1.40)
Farm Owner	0.156	-0.036	-0.033	-0.154	-0.009
	(1.40)	(0.52)	(0.71)	(2.22)**	(0.05)
Borrowed Money	-0.043	0.038	0.047	0.023	-0.064
	(0.50)	(0.69)	(1.25)	(0.45)	(0.58)

Parents Away	-0.065	-0.121	-0.041	0.022	0.140
	(0.63)	(1.84)*	(0.86)	(0.33)	(0.75)
Siblings Away	0.102	0.183	0.239	0.113	0.386
	(1.02)	(2.96)***	(5.63)***	(1.85)*	(2.48)**
Rural hhld	-0.160	-0.068	0.035	-0.113	-0.153
	(1.38)	(0.95)	(0.71)	(1.64)	(1.02)
Kabup Pr Urban	-0.009	0.077	-0.189	0.060	0.805
	(0.03)	(0.36)	(1.27)	(0.29)	(1.72)*
Kabup Pr Unempl	2.416	2.009	0.360	-1.156	-1.268
	(0.72)	(0.95)	(0.23)	(0.56)	(0.28)
Kabup Pr SelfEmpl	-1.935	0.046	-1.189	-1.287	-0.944
	(1.74)*	(0.07)	(2.53)**	(1.91)*	(0.68)
Kabup Pr Agric	1.431	0.260	0.297	0.693	1.765
	(2.48)**	(0.69)	(1.18)	(1.89)*	(2.10)**
Kabup Pr Poor	-0.420	0.027	0.138	-1.160	-0.694
	(0.66)	(0.07)	(0.51)	(2.68)***	(0.71)
Kabup Gini coeff	0.002	-0.216	-1.319	0.212	2.266
	(0.00)	(0.36)	(3.03)***	(0.39)	(2.16)**
Constant	-2.551	-1.651	-1.439	-4.662	-2.194
	(3.59)***	(3.60)***	(4.33)***	(7.44)***	(2.09)**
Pseudo-R^2	0.1490	0.0851	0.0782	0.1445	0.2406
N	5,685	5,973	6,900	6,073	5,637

* p<0.1; ** p<0.05; *** p<0.01

b. Balancing Tests and Propensity Score Estimates of the ATET of Migration

In line with the general agreement in the economic literature, the previous section showed that departure/receipt of an household member, creation of a new split-off household, as well as move of an entire household are not random events. Based on this, comparing consumption growth of migrant households and non-migrant households is expected to result in a positive consumption growth effect for the migrant households simply because, inter alia, households headed by individuals who attended university and where more students were found in 1997 were also more likely to use migration as a coping mechanism from 1998-2000. Therefore, such crude comparisons would lead to biases for the 'causal effects' of migration that have to be corrected. Indeed, all possible parametric, semi- and nonparametric estimators of causal effects that allow for heterogeneous effects are implicitly or explicitly based on the idea that, in order to find the effects of being in one state and not in another, outcomes from observations from both states with the same distribution of relevant characteristics should be compared (Lechner 2009).

A number of specifications are used in the present study to produce such comparisons, and both traditional matching estimators and weighted smoothed matching estimators are used and compared: namely, the nearest neighbour matching estimator is used, as well as caliper matching, which imposes a value for the maximum distance of controls from the treated observation. By doing so, estimates should remain more stable, since this prevents from matching a large number of neighbours in most of the cases. For it to be binding in each specification, we impose a caliper of 0.2% in all cases. A number of Kernel estimators are also used: Kernel matching is implemented both using an Epanechnikov curve and a Gaussian curve, with or without imposing the common support and with the imposition of different bandwidths (therefore being more or less strict in the weight given to observations further away from the propensity score). These elements are combined in a number of ways. However, in each case, the ocular and sensitivity method is used to choose the bandwidth. Finally, Mahalanobis-metric matching and Local Linear Regression matching are also performed.

Among the presented specifications, a series of diagnostic tests are implemented at the aim of selecting the best-performing estimator in terms of bias minimization and balancing of the covariates between treated and controls. For all specifications, the extent of balancing of the covariates between the two samples (migrant and non-migrant households) before and after matching the samples are compared; for each explanatory variable, t-tests for the equality of means in the two samples are computed, respectively before and after the matching. These t-tests are based on a regression of the variable on a treatment indicator; before matching, an unweighted regression was run on the entire sample, whereas after matching the regression is weighted using the matching weight variable and based on the sample on-support (i.e., the matched sample). Additionally, for each explanatory variable, the standardized percentage bias is also calculated; such measure is defined as the percentage difference of the sample means in the treated and non-treated sub-samples (full or matched, namely before and after matching) as a percentage of the square root of the average of the sample variances in the treated and non-treated groups (see Rosenbaum and Rubin, 1985). Finally, overall measures of imbalance in the covariates are also computed: the Pseudo R^2 from the participation equation is calculated on all the included variables both before and after the matching. The mean and median bias as summary indicators of the distribution of the absolute bias are also measured, both before and after implementing matching.

Kernel estimators are generally found to outperform other estimators[18]; in particular, the matching estimator that is found to perform best in this context (in terms of bias reduction) is the kernel estimator with replacement, when an Epanechnikov curve is used, and the common support and the bandwidth at 0.5% are imposed. Most importantly, when this estimator is used to perform matching, no statistically significant difference in the distribution of the covariates is found between the two samples after the matching: this conclusion applies to all covariates in each binary comparison. In all the t-tests implemented, the null hypothesis of equality of the means between the two samples cannot be rejected and no variable is found to be still significantly unbalanced ($p<=0.10$) after matching. For all bilateral comparisons, a likelihood ratio test of the joint significance of the regressors is also implemented, both before and after matching; in each case, we test the null hypothesis of joint significance of the regressors. The null hypothesis of joint significance of the regressors could never be rejected before matching; on the contrary, the null hypothesis is always strongly rejected after matching. In all bilateral comparisons, the median bias observed after matching is below 3%; although a clear indication for the success of the matching procedure is not generally available, in most empirical studies a standardized bias below these thresholds is seen as satisfactory (Caliendo and Kopeining, 2008). More details on the performance of this matching estimator are presented in Appendix II.

Charts 1, 2, 3, 4 and 5 report the distributions of the common support condition in each bilateral comparison, after matching was performed: the evidence suggests that the large majority of observations are found to be on common support, hence they can be compared in a meaningful way.

[18] Support to the use of kernel estimators is provided by Frölich (2004); in his Monte Carlo analysis, he compares alternative matching estimators on a broad set of data generating processes using a mean squared error criterion. Among other things, this work suggests that, over a range of possible data generating processes, kernel matching performs well. He also concludes that the nearest neighbour matching performs particularly poorly in his study.

Charts 1, 2, 3, 4 and 5 – Presence of Common Support between Migrant and Non-Migrant Households (for each bilateral comparison, in the control group are Non-Migrant households only)

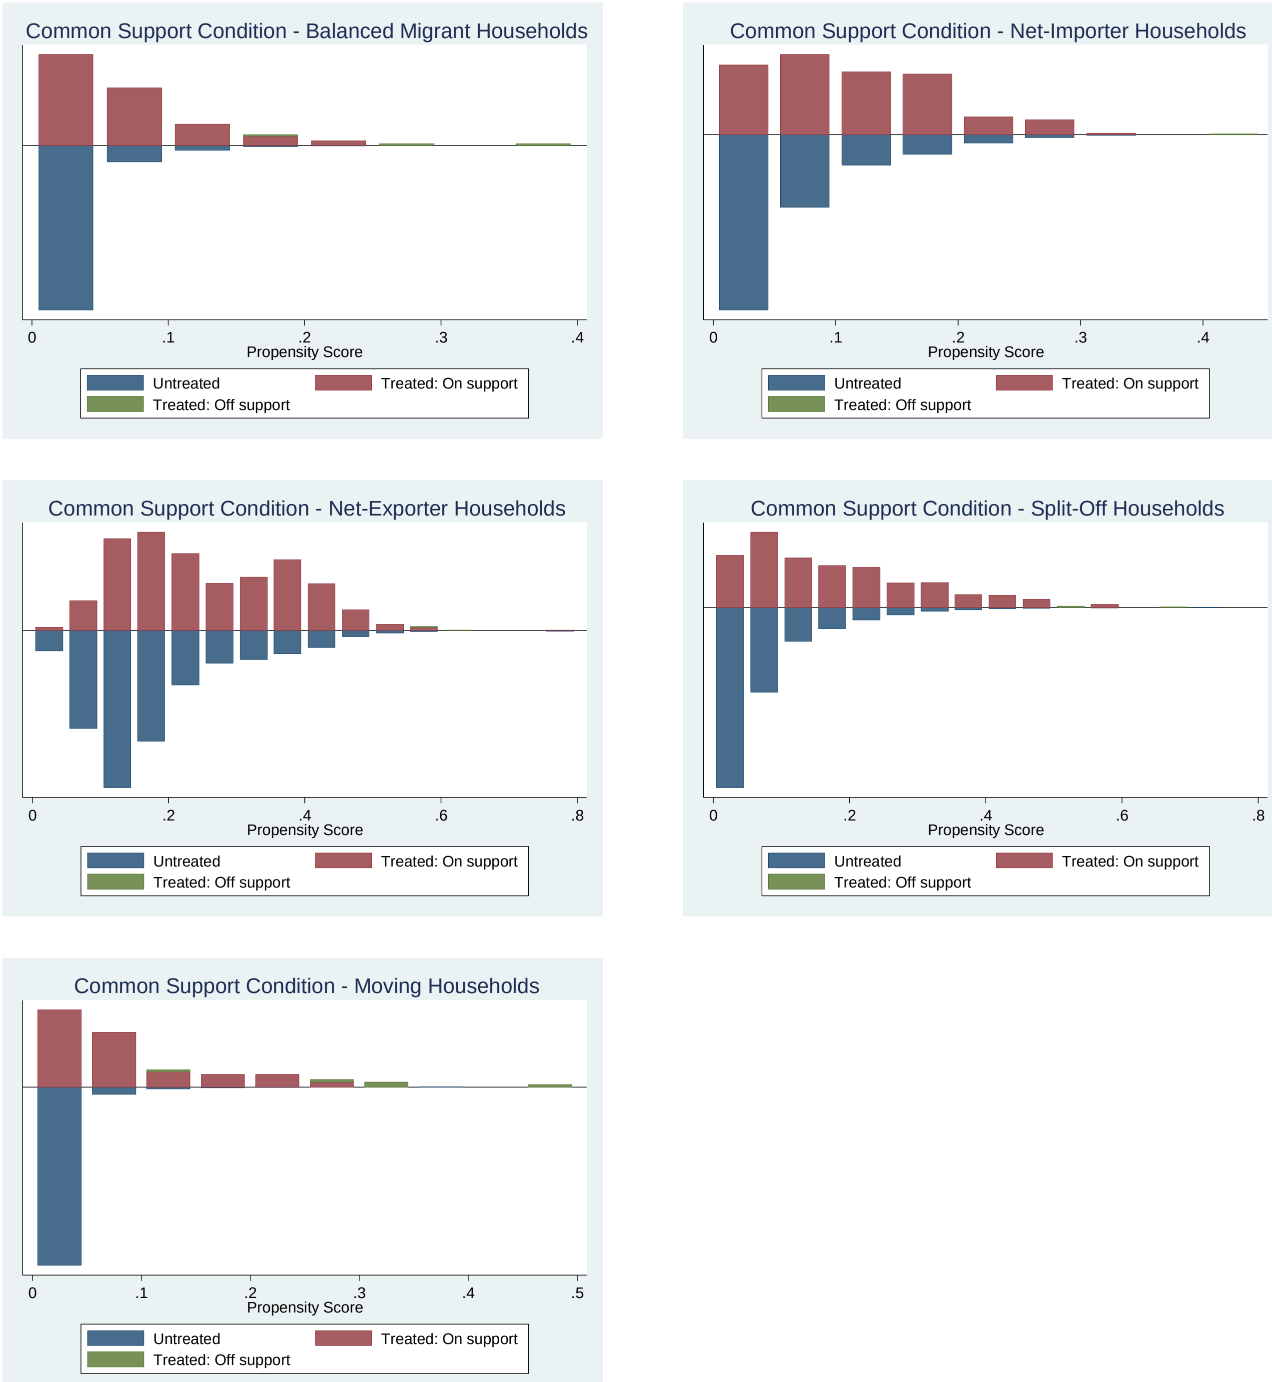

As a consequence of these results, we selected the kernel estimator with replacement, when an Epanechnikov curve is used, and the common support and the bandwidth at 0.5% are imposed as our preferred estimator for all binary comparisons and then used bootstrap techniques with replacement to calculate standard errors; to be precise, 1000 bootstrap replications were carried out in each case. One major reason for applying the bootstrap to the same estimator for all binary comparisons is to make the derived results as comparable as possible across migration categories. Table 9 shows the evidence on the difference-in-differences in consumption levels in the short run, as well as in the long run. No significant difference is observed between balanced migrant households and non-migrant households by 2000; interestingly, a positive and significant difference is observed instead by 2007/8, indicating that gains from the modification of living arrangements during the crisis took some time to materialize for this group of households. For net-importers, once the propensity scores were estimated and matching was performed, significant negative consequences are found by 2000, whereas no statistically significant effect of migration is found for them by 2007/8. For net-exporters of migrants, positive gains from migration appear strong and constant over time; positive and significant gains from migration are also observed for split-off households, although the net gains seem to decrease over time. Finally, no significant gains from migration are found for households that proved mobile during the crisis; however, a word of caution is needed here due to the small number of treated observations.

Table 9 – ATET of Migration

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Short Run (97-00)	0.033	-0.199	0.072	0.202	-0.071
	(0.49)	(5.08)***	(2.97)***	(4.10)***	(0.73)
Long Run (97-07)	0.161	-0.039	0.076	0.092	0.069
	(2.15)**	(0.82)	(2.91)***	(2.05)**	(0.69)
Treated	124	412	1,339	512	76
Untreated	5,561	5,561	5,561	5,561	5,561
N	5,685	5,973	6,900	6,073	5,637

* p<0.1; ** p<0.05; *** p<0.01

In the next section, we explore in a number of ways the heterogeneity in the effects of migration across different households that may lie behind the evidence presented in Table 9; given that migration may occur for a variety of reasons and different outcomes are related to it in different circumstances, we find this heterogeneity an element of major interest. To be precise, given that the focus of the analysis of this paper is on the migration response to the economic crisis that hit Indonesia in the late 1990s and the effect of this response on the evolution of living standards in the future (rather than on migration per se), we present evidence on the implications of migration for a number of different subsamples of population, trying to capture, in a number of ways, the diversity in the effects of migration in response to the East-Asian crisis in Indonesia. First, we consider the semi-parametric IV estimation of local average treatment effect (LATE) suggested by Frölich (2007) to retrieve the effect of migration for those that, at the onset of the crisis, may have found it relatively easier to experience migration. Second, we try to isolate the inherent different propensity to move of different households by using information on their past migration history. Third, we address concerns on the natural aging of the sample and the life-cycle element of migration by using information on the occurrence of births in the years after the crisis. Finally, we concentrate our attention on those that, respectively from a social and a geographic perspective, were hit hardest in Indonesia by the crisis; namely, on households that were not rice-producers at the time of the crisis, and on households that resided in Java in 1997.

6. Further Investigation and Subsamples Analysis

a. Semi-parametric IV estimation of local average treatment effect (LATE)

In the estimation of the impact of migration on the evolution of living standards, an alternative strategy to address the potential endogeneity of migration is Frölich's (2007) semi-parametric instrumental variable estimator to estimate the treatment effects for the compliers[19]. We consider five versions of the LATE estimator using five sets of instrumental variables, namely (i) whether the household head resided at age 12 in a different district from the district of birth, (ii) whether any household member resided at age 12 in a different district from the district of birth, (iii) whether the household experienced any migration in the years prior to the crisis (1993-97), (iv) the presence of siblings away from the district in 1997, and (v) the number of siblings away from the district in 1997[20].

In fact, we acknowledge that a number of limitations are attached to each of these instruments; as discussed in section 4, an entirely persuasive instrument does not seem to be available in our analysis. None of these instruments is randomly assigned; to make them proper instruments, it is necessary to condition on some covariates. If we did not include conditioning variables in our analysis, there could be common factors that influence both the instrument and consumption growth. Frölich (2007) shows that, when a set of assumptions is identified (for this, see Frölich 2007), his conditional LATE estimator is identified. We define a binary instrument $Z = 0$ if, e.g., no migration was experienced in the 1993-97 period and $Z = 1$ otherwise, and we denote compliers as $T = c$. These are households that will experience migration during the crisis if $Z = 1$, i.e., if they experienced migration from 1993-97. For this type of households we estimate the average treatment effect of migration. Following the notation above, we estimate $\gamma^{ATE} = E(Y^1 - Y^0 \mid T = c)$. In addition to this, we also estimate the proportion of compliers in our data, although it is not possible to identify them in our data.

Finally, it is important to stress that the average treatment effect for the compliers $\gamma^{ATE} = E(Y^1 - Y^0 \mid T = c)$, is a very different parameter than the average treatment effect on the treated $ATET = E(Y^1 - Y^0 \mid D = 1)$; the former parameter applies to a group of households that would react to an exogenous shift in Z, but individual compliers cannot be identified; on the contrary, the ATET calculated in the main specification applied to a clearly defined fraction of our sample. Using all proposed instruments, we estimated Frölich's (2007) LATE estimator for each bilateral comparison. For reasons of space and given the similarity of our conclusions in all cases, here we only report (and discuss) results derived from our fourth proposed instrument, i.e, the presence of siblings away from the district in 1997. We discuss these results because the estimated proportion of compliers was highest for this proposed instrument.

As shown in Table 6, presence of siblings in a different district in 1997 is a strong predictor of the migration decision during the crisis; this effect is always positive and, with only one exception, statistically significant. This variable was included in the participation equation in the main specification because presence of siblings away from the district in 1997 may affect consumption growth not only directly through migration, but also by increasing the likelihood to receive

[19] A recent paper that also applies this methodology to assess the effectiveness of mobility is Ham et al. (2011).
[20] Siblings of any household member are considered here.

remittances; therefore it may affect consumption growth in a way that goes beyond migration per se. This is a possibility and, for the present IV analysis, it constitutes a limitation which we acknowledge[21].

Bearing these consideration in mind, we estimated Frölich's (2007) LATE estimator using this instrument for migration for each bilateral comparison. In each bilateral comparison, we calculated a local logit regression where the dependent variable is equal to one if any member of the household reported the presence of siblings away from the district in 1997, and equal to zero otherwise. Therefore, in each bilateral comparison, the two groups here are defined by the instrument, not by the actual migration decision. As noted above, we also calculated with precision the proportion of compliers in our sample: we estimate the proportion of the sample who are compliers to be 0.5% in the comparison between balanced-migrant households and non-migrant households (the base category), 2.1% in the comparison between net-importers and the base category, 5.5% in the comparison between net-exporters and the base category, 1.5% in the comparison between split-offs and the base category, and 0.2% in the comparison between moving households and the base category. Further, Table 10 reports our estimates of the LATE of migration for compliers for all bilateral comparisons. In order to draw inference, 1000 bootstrap replications were implemented for the entire procedure in each case. The very large standard errors reported in Table 10 are best explained by the small fraction of compliers in each case, which does not allow a precise estimate of the LATE in our analysis. In the recent literature, this conclusion was reached in a number of occasions. Notably, Frölich and Lechner (2006) conclude that LATE estimates are highly variable for individual labour markets; also Ham et al., (2011) conclude that the LATE estimator requires a large amount of data and also in their analysis too few compliers are available to estimate the LATE of migration with any precision.

Table 10 – LATE of Migration for Compliers

LATE of Migration for Compliers	Balanced LATE	Net-Importers LATE	Net-Exporters LATE	Split-Offs LATE	Moving LATE
Short Run LATE 1997-2000	7.034	1.573	0.635	2.424	17.978
	(488.10)	(325.87)	(0.52)	(56.70)	(63.51)
Long Run LATE 1997-2008	0.466	0.159	-0.017	0.036	1.150
	(93.01)	(14.11)	(0.72)	(57.67)	(36.65)
N	5,685	5,973	6,900	6,073	5,637

* *p*<0.1; ** *p*<0.05; *** *p*<0.01

[21] However, for each bilateral comparison, we also estimated a reduced-form equation, where the consumption growth (namely, by 2000, and by 2007/8) was estimated as a function of a set of conditioning variables. To be precise, the same variables used for the matching procedure in the main specification were used here as covariates. When this equation was estimated, our proposed instrument did not seem to have any direct effect on consumption growth. The estimated coefficients for our proposed instrument in the consumption growth equations were never found to be statistically significant, with only one exception, namely the bilateral comparison between net-exporters and non-migrant households, where the coefficient was marginally significant – with an estimated t-statistic equal to 1.67.

b. Past Migration History

One potential source of heterogeneity in the assessment of the effects of migration in response to the crisis stems from the inherent different propensity to move of different households: some households may have been naturally more mobile than others in those years, irrespectively of the crisis itself. This is relevant because migration in response to the crisis may have not produced the same benefits among more mobile households as among households less likely to move. One way to capture the different propensity to move in those years is by observing the migration decisions of households in the years prior to the crisis. In fact past migration history was controlled for in the participation equation in the main specification, hence the role of past migration is taken into account in the determination of the results from our main specification. However, by simply controlling for this information, one could not detect the presence of heterogeneous effects of migration in response to the crisis between households that proved mobile in the years prior to the crisis and those that did not. In order to unveil this heterogeneity, the matching estimator was performed separately within subsamples defined such that in each subsample all households have the same level of past migration (in this study, 'past' is defined as from 1993-97). Then, within each subsample, the effects of the crisis-period migration was analysed.

Exploiting the panel dimension of the data, only households that did not experience any migration in the 1993-97 period were now included in the analysis (i.e., households that did not prove mobile in the years preceding the crisis); this means that, out of the total sample of 8,024 households that reported complete information, only 5,549 were kept. Since the same definition of migration was still used and the same categories were defined to group households according to their migration decision, these 5,549 households are households that did not migrate out of the kabupaten in the years from 1993-97, neither did they produce/receive any migrants that crossed the kabupaten border in those years. Within this subsample of households, 1000 bootstrap replications were applied again to the same matching estimator as in the main specification (i.e., a kernel estimator with replacement, using an Epanechnikov curve, imposing the common support and applying a bandwidth at 0.5%); this estimator was used again in order to ensure comparability of the results. The results of this analysis are reported in Table 11, respectively for the effect of migration on the evolution of living standards in the short run and in the long run.

Table 11 – ATET of Migration – No 93/97 migration subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Short Run (97-00)	-0.050	-0.234	0.102	0.245	-0.080
	(0.41)	(4.18)***	(3.44)***	(3.45)***	(0.69)
Long Run (97-07)	-0.066	-0.127	0.089	0.117	0.108
	(0.54)	(2.09)**	(2.81)***	(1.95)*	(0.77)
Treated	40	193	712	217	46
Untreated	4,215	4,341	4,341	4,294	4,294
N	4,255	4,534	5,053	4,511	4,340

* p<0.1; ** p<0.05; *** p<0.01

The evidence presented in Table 11 is generally consistent with the results calculated on the entire sample of households: however, point estimates tend to be larger in all cases, suggesting that for this subsample of households the consequences of mobility in the years of the crisis were more significant. More specifically, substantial consistency was found for both net-exporters and split-off households. On the other hand, some interesting differences are observed in the long run effectiveness of migration for both balanced households and net-importers of migrants. For this subsample of households, balanced migrant households no longer seem to derive net gains from migration by 2007/8; net importers of migrants still observe a decline in their well-being, but, in this case, the decline seems to persist and to be significant also in the long run. Even though the differential in consumption growth between net-importers and non-migrant households seems to decrease in the long run, a significant negative difference is found in both periods. On the other hand, net-exporters of migrants still benefit from the departure of one (or more) member(s) both immediately and in the long run, suggesting that the benefits from migration are stable over time for these families and implying no substantial change in the effect of migration between the short run and the long run. A significant gain is still observed in both periods also for newly formed split-off households, suggesting again long-lasting beneficial effects of the decision to move in the years of the crisis. Finally, no significant impact of mobility was found for the last (small) group of households, moving households, neither by 2000 nor by 2007/8. For the category of balanced migrant households and that of moving households a concern still applies because of the low number of treated observations; the small size of the treated sample also resulted in problems of multicollinearity, and caused some observations to be dropped from the analysis. This explains the lower number of untreated observations reported in Table 11 in the calculation of the ATET of migration for certain categories of households.

A similar analysis was also implemented on the sample of households that, on the contrary, experienced geographic mobility in the period 1993-97, thus proving more mobile regardless of the crisis; within this subsample of households, matching techniques were applied again in order to compare households that migrated again in the years of the crisis (or that produced/received migrants) with comparable households that did not do so. Standard errors were again derived from 1000 bootstrap replications and the results of this analysis are presented here; most of the evidence showed in Table 12 is consistent with the evidence presented above. In particular, the evidence is consistent with Table 9 for balanced households, that exhibit significant gains from changes in living arrangements during the crisis by 2007/8. Net-importers of migrants are still found to be negatively affected by the net receipt of migrants by 2000, whereas no significant effect is observed, on average, by 2007/8. It is interesting to notice that, in comparison with the same coefficient in Table 11, the percentage difference between net-importers' consumption growth and non-migrant households' consumption growth by 2000 appears smaller among these households than among households that were immobile in the pre-crisis period. Arguably, these results suggest that the negative effect of the net receipt of migrants was stronger among immobile households than among mobile households. Consistent results are found for moving households, for which no significant gains from migration emerged, neither in the short run nor in the long run. For balanced households and moving households, the usual concern applies to the results due to the size of the treated sample[22]. Finally, for both net-

[22] The small size of treated sample resulted again in problems of multicollinearity, causing some observations to be dropped from the analysis. This explains the lower number of untreated observations reported in Table 12 in the calculation of the ATET for moving households than in the calculation of the ATET for all other migration categories.

exporters and split-off households, the point estimates of the ATET of migration appear smaller than for the subsample that did not prove mobile from 1993/97; in the short run the ATET of migration does not seem significant for net exporters, whereas by 2007/8 a lack of significance applies to the ATET of migration for split-off households.

Table 12 – ATET of Migration – 93/97 migration subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Short Run (97-00)	0.008	-0.188	0.050	0.150	0.062
	(0.08)	(3.16)***	(1.26)	(2.00)**	(0.24)
Long Run (97-07)	0.216	0.050	0.077	0.097	0.015
	(2.15)**	(0.69)	(1.75)*	(1.38)	(0.07)
Treated	84	219	627	295	30
Untreated	1,220	1,220	1,220	1,220	981
N	1,304	1,439	1,847	1,515	1,011

* p<0.1; ** p<0.05; *** p<0.01

Overall, the results presented in Table 11 and Table 12 provide evidence that migration in response to the crisis produced tangible effects, especially among immobile households, both in the short run and in the long run; among these households, that did not prove mobile until the crisis struck, those who moved, as well as net exporters of migrants, derived significant benefits from the migration decision. Net recipients of migrants also felt the consequences of migratory inflows as they experienced lower p/capita consumption growth both by 2000 and by 2008. On the other hand, the smaller effects of migration that appear from Table 12 for the subsample that proved more mobile in the pre-crisis period may stem from the fact that households that proved to be more mobile may have managed to cope with the crisis whether they moved this time or not; this hypothesis seems also suggested by the (aforementioned) difference between the ATET for net importer households presented in Table 11 and that presented in Table 12[23].

In order to investigate this further, we looked at the evolution in consumption levels separately for households that proved mobile in the years preceding the crisis and those that did not. This information is displayed in Table 13, where the mean growth in (ln) per capita consumption level is presented for each bilateral comparison made in the econometric analysis above and separately for mobile and immobile households in 1993-97. All statistics presented in Table 13 were derived only from households on common support, since this is the same sample that the estimates above apply to. What emerges from the comparison of consumption growth rates of 1993-97 immobile households and those of 1993-97 mobile households is that the latter seemed to fare worse than the counterpart that did not prove mobile in 1993-97. The observed average levels of consumption are often lower

[23] I.e., -23.4 percentage points as opposed to -18.8 percentage points in the short run, and -12.7 percentage points as opposed to +5 percentage points in the long run – although the last coefficient was not significantly different from zero

among mobile households in the 1993-97 period than among others. This suggests that, at the onset of the crisis, households that revealed more mobile in the years prior to the crisis generally found themselves in a worse position compared to the less mobile counterpart in 1993-97.

Table 13 - ΔlnC p/capita by migration category – Analysis disaggregated by pre-crisis (1993/97) migration status

ΔlnC $(p/capita)$	Balanced Migrant Households		Net-Importer Households		Net-Exporter Households		Split-Off Households		Moving Households	
	Untreat	Treat	Untreat	Treat	Untreat	Treat	Untreat	Treat	Untreat	Treat
93/97 Immobile Short Run ΔlnC	0.734	0.765	0.733	0.468	0.733	0.861	0.732	1.01	0.732	0.571
93/97 Mobile Short Run ΔlnC	.655	.688	.658	.475	.655	.72	.655	.783	.628	.647
93/97 Immobile Long Run ΔlnC	1.888	1.845	1.884	1.692	1.884	2.004	1.882	1.998	1.882	1.803
93/97 Mobile Long Run ΔlnC	1.729	1.926	1.729	1.775	1.729	1.828	1.729	1.726	1.689	1.549

However, this is especially true of households that did not move during the crisis; in fact, by looking at those households that experienced migration in the years of the crisis, the differences are considerably smaller. This suggests that migration in the years of the crisis contributed in narrowing the difference between the changes in consumption of pre-crisis mobile households and those of pre-crisis immobile households; in fact in some cases the difference was not simply reduced, but even overturned. For example, among net-importers of migrants, higher levels of consumption were observed among pre-crisis mobile households than among pre-crisis immobile households, both in 2000 and in 2007; given that living standards of net-importer households were generally found to be negatively affected by the net receipt of migrants in the years of the crisis, this result is of particular relevance: it implies that, in this case, more mobile households managed to cope better with the crisis, and, therefore, they experienced lower negative consequences from receiving new members into the household. On the contrary, for households that did not prove mobile in the years prior to the crisis, the negative shock attached to the net receipt of migrants was felt more acutely.

c. Life-Cycle Events and Family Dynamics

Given that the causal impact of the crisis on the migration decision was not modelled explicitly in this study, a concern was raised that the increase in migration rates observed during the crisis may not be attributable to the crisis, but rather stem from the natural aging of our sample. This would amount to saying that, in fact, the geographic mobility observed in 1997-2000 was not driven by the economic shock, but rather by personal reasons unrelated to the crisis, such as creation of new families.

The long time period spanned by the IFLS panel dataset (i.e., 1993-2008) allows exploring, to some extent at least, if this was indeed the case. One way to test this is to restrict the analysis to those households that did not experience the birth of any new children in the years following the crisis; in fact, if the formation of a new household in the 1997-00 period preceded the formation of a new family, and was not driven by the occurrence of the economic shock, then observing new births in the household from 2000-08 would not be unlikely. On the other hand, birth of new children would be less likely to occur subsequent to migration in response to an economic shock. The extent to which family formation and life-cycle events are likely to influence the results above was tested and results are presented in Table 14, Table 15, Table 16 and Table 17.

In fact, occurrence of births in the household may appear endogenous to the migration decision during the crisis, as well as in the determination of the long run consumption growth. The possible endogeneity of occurrence of births is acknowledged; however, we believe that these tests can still be informative on the presence of pure life-cycle elements in the determination of the migration decision and the consumption growth's outcome observed here. In this line of thought, households were divided between those where no new births occurred in the period 2000/07, and those where (any number of) births happened after 2000. In both cases, the sample was also further restricted to the households that did not prove mobile in the years prior to the crisis; the motivation for this is that households that did not prove mobile in the years prior to the crisis and that did not experience any new births after the crisis are arguably the most likely to have migrated in the years of the crisis because of the crisis. Consequently, 'migrant' and 'non-migrant' households were matched again within subsamples and matching estimates of the difference-in-differences in consumption over time were calculated again.

1000 bootstrap replications were implemented again to the kernel estimator chosen in the main specification in order to obtain standard errors; the results are showed in Table 14 and Table 15, respectively for the subsample of households where no births occurred between 2000/07, and the subsample where neither births occurred between 2000/07 nor mobility in the pre-crisis period. The evidence seems to confirm that migration in response to the crisis generated positive gains for those that moved. Table 14 only shows significant and positive gains for split-off households; the ATET of migration associated to net-exporters of migrants is still positive, but marginally not significant.

Table 14 – Long Run ATET of Migration – No 00-07 births subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Long Run (97-07)	0.137	-0.013	0.050	0.113	0.203
	(1.47)	(0.22)	(1.58)	(1.66)*	(1.58)
Treated	87	277	974	247	42
Untreated	3,511	3,511	3,511	3,468	3,511
N	3,598	3,788	4,485	3,715	3,553

* p<0.1; ** p<0.05; *** p<0.01

On the other hand, what emerges from Table 15 is that, in line with the evidence found for the whole sample, significant and positive gains in the evolution of consumption were experienced by both net-exporters of migrants and split-off households. A positive and significant coefficient is also associated to the ATET of moving households in this case; however, the small number of treated observations imposes caution in the interpretation of this result.

Table 15 – Long Run ATET of Migration – No 00-07 births subsample & No 93/97 migration

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Long Run (97-07)	-0.054	-0.075	0.064	0.269	0.373
	(0.25)	(0.92)	(1.67)*	(2.98)***	(2.06)**
Treated	22	131	521	110	25
Untreated	2,727	2,811	2,811	2,775	2,268
N	2,749	2,942	3,332	2,885	2,293

* p<0.1; ** p<0.05; *** p<0.01

Drawing the attention on those households among which births were observed between 2000-08 and, therefore, where migration was more likely to occur regardless of the crisis, a different picture emerges. Table 16 and Table 17 show the results of these analyses, respectively for households where births occurred in the period 2000-07 and for households where births occurred in the period 2000-08 but no migration occurred in 1993-97. In each case, 1000 bootstrap replications were still applied to derive the standard errors. The general evidence from these subsamples of households is that, regardless of pre-crisis mobility, no significant differences are observed between the evolution of consumption of 'migrant' households and the evolution of consumption of 'non-migrant' households.

Table 16 – Long Run ATET of Migration – 00-07 births migration subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Long Run (97-07)	-0.082	-0.077	0.051	0.099	-0.026
	(0.59)	(0.93)	(1.04)	(1.46)	(0.12)
Treated	37	135	365	265	34
Untreated	2,050	2,050	2,050	2,050	1,329
N	2,087	2,185	2,415	2,315	1,363

* p<0.1; ** p<0.05; *** p<0.01

One possible explanation for this is that a higher number of births may be associated with the migration event: in fact, here we are only controlling for the occurrence of births, and not for the number of births, implying that we may be comparing 'migrant' households where several children were born after the crisis with households where only one child was born. If this was the case, then this would explain (at least in part) why the evidence of the positive effect of migration on the evolution of living standards in the long run is not confirmed for this subsample of households. However, this hypothesis was tested and, for each bilateral comparison, a t-test was carried out on the equality of means in the number of births after 2000 between the treated and the untreated on common support. The average number of births among matched non-migrant households was around 1.25 births per household from 2000-07, among matched net-importer households the average was around 1.4, among matched net-exporter households it was approximately 1.23 births and, finally, 1.18 births were reported on average among matched split-off households. The average number of births among matched non-migrant households was compared to the average number of births in each other group, and only for net-importer households the null hypothesis of equality of means could be rejected at 5% level. The group of net-importers of migrants is also the only category for which a significant (and negative) coefficient is associated to the ATET of migration. For the remaining two migration categories, a negative difference (although not significant) was found between their average number of births after 2000 and that among non-migrant households; this is suggestive of a (weakly) negative correlation between migration and births in the subsequent years. However, this evidence rules out the hypothesis that the lack of positive gains from migration for net exporters and split-off households stems from a higher number of births in migrant households. Rather, the lack of significance in these coefficients may simply stem from the fact that, for these households, migration was less likely to be driven by economic considerations and more by personal and family considerations; in a similar scenario, it is not surprising that lower gains are observed, as measured by the evolution of p/capita consumption level in these households.

Table 17 – Long Run ATET of Migration – No 93/97 migration & 00-07 births subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Long Run (97-07)	-0.177	-0.245	0.013	0.144	-0.059
	(0.82)	(2.19)**	(0.18)	(1.53)	(0.15)
Treated	18	62	191	107	21
Untreated	1,164	1,530	1,530	1,519	1,003
N	1,182	1,592	1,721	1,626	1,024

* p<0.1; ** p<0.05; *** p<0.01

d. Migration and the Crisis

In the attempt to further isolate the crisis, the final section builds upon the empirical literature that has recognized a tremendous diversity in the impact of the crisis, both at the social level (for example, Frankenberg, Smith and Thomas, 2003) and at the geographic level (for example, Ravallion and Lokshin, 2007). Among the major channels of transmission of the crisis to the people, one was certainly the price of many commodities. Annual inflation was estimated by the Central Statistical Bureau to be about 80 percent for 1998. The change in the price of rice – the single most important commodity, as measured by the budget share, for most Indonesians (Friedman and Levinsohn, 2002) – was even more dramatic: the price of rice spiralled upwards by almost 200% within few months, the highest increase registered at that time among food commodities. In Central Java, in 1997/98 the price of rice even tripled, and poor households maintained their rice consumption only by cutting back on other foods, resulting in a steep drop in micro-nutrient intake among poor children (FPSA 2003). In fact, as pointed out by McCulloch (2008), one of the reasons behind the equitable economic trend of the country from the early 1980s until the crisis in 1998, was the focus of the Soeharto government on agriculture and the promotion of industrialization by ensuring the stability of food prices, most importantly the price of rice. For almost 25 years, after the food crisis in 1974, rice production was heavily subsidized by the Soeharto's government; such subsidies aimed in particular at developing infrastructures and disseminating new technologies and seeds in the 1970s and 1980s (Timmer 2005, and McCulloch 2008). Moreover, the market for rice in Indonesia was strongly influenced by the role of Bulog, the state logistics agency, that had the monopoly on imports of rice, and that was also in charge of maintaining domestic prices at around the world price. Also as a result of the political incentives to act in this direction, Bulog managed to keep the price of rice in Indonesia in line with the world price up till the late 1990s (Dawe 2008). The crisis led to the end of this state of affairs for a number of reasons; Soeharto resigned in 1998 after 31 years of presidency, no more subsidies for rice production were granted and the price of rice was eventually allowed to float freely and was no longer stabilized at around the world price. Friedman and Levinsohn (2002) and Ravallion and Lokshin (2007) have pointed out that these dynamics in the level of prices likely hurt the poor more than others.

However, rice can be also an important source of income for the poor. Many of the poor in Indonesia may have produced rice while consuming it, and some of them may have been net sellers of rice, thus receiving income from its sale (as well as from other food crops). In this connection, Friedman and Levinsohn (2002) argue that, in the assessment of the impact of the crisis and the rampant inflation that spread in Indonesia in 1998, omitting to account for household self-production of food is likely to change dramatically the results of the analysis, especially for rural households.

With these considerations in mind, rice-producers that were found in 1997 were excluded from the analysis, whatever the amount of their rice production was; arguably, these households may have been more protected from the rampant inflation in the price of food commodities, and their inclusion in the analysis may have influenced the results. In a recent study, Chen (2010) has also used a similar information, i.e., pre-crisis hectares of wetland, to measure the isolation from the crisis of rice producers. Non-rice producers may have responded to the rise in prices in several different ways, but they were certainly not protected from such inflation by their own production of rice. The kernel matching estimator and 1000 bootstrap replications were implemented again in each case, and the evidence on the effectiveness of migration in response to the crisis for the subsample of non-rice producers is presented in Table 18. Both in the short run and in the long run, the results presented on the subsample of non-rice-producers provide support to those found for the entire sample in the main specification. The only significant discrepancy applies to balanced households that, in this case, did not experience significant gains by 2007/8. For all other results, this consistency in the evidence is of particular relevance, since this experiment isolates the crisis better than the main specification[24].

Table 18 – ATET of Migration – No Rice Own Production subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Short Run (97-00)	-0.005	-0.185	0.068	0.182	-0.023
	(0.06)	(4.02)***	(2.44)**	(3.35)***	(0.21)
Long Run (97-07)	0.112	-0.015	0.073	0.087	0.080
	(1.30)	(0.31)	(2.53)**	(1.67)*	(0.72)
Treated	98	334	1,075	422	67
Untreated	4,358	4,358	4,358	4,358	4,358
N	4,456	4,692	5,433	4,780	4,425

* p<0.1; ** p<0.05; *** p<0.01

Finally, at the geographical level, Ravallion and Lokshin (2007) argued that initially poorer areas saw lower proportionate impacts of the crisis, while the well-integrated areas proved to be more vulnerable to the crisis. In this line of thought, an interesting case study is the island of Java, where, for a number of reasons, it is plausible to believe that the crisis was felt more than elsewhere. In fact, Sumarto, Wetterberg, and Pritchett (1998) concluded that the urban areas on the island of Java - among the wealthiest regions in the country until 1997 - were the epicentre of the financial and modern sector crisis, whereas traditionally poorer natural resource exporting areas seemed to derive

[24] Evidence is not presented for rice-producers, since too few were found to permit a meaningful analysis for most of the migration categories.

benefits from the crisis. Jakarta and the island of Java were also the regions where the crisis was felt earliest; considering that the crisis started in the Jakarta foreign exchange financial market, this is in fact not surprising, given that Jakarta is in Indonesia the most integrated city to the global financial system, whereas West Java is the manufacturing centre of the country. Additional support to this is provided by Akita and Alisjahbana (2002): using a Theil index based upon district-level GDP and population data, they present evidence at the national level that regional income inequality increased significantly from 1993 to 1997, mainly due to an increase in within-province inequality, particularly in Riau, Jakarta and West and East Java. In 1997, the within-province component represented about 50% of regional income inequality. In line with the inequality-reducing effect of the crisis found by Ravallion and Lokshin in 2007, in 1998 regional income inequality dropped to its 1993-94 level. However, in contrast to 1993-97, three-quarters of the 1998 decline was due to a change in between-province inequality, with the Java-Bali region playing a prominent role. The authors conclude that the crisis seemed to have hit particularly urban Java and urban Sumatra. Finally, Studdert, Frongillo and Valois (2001) also assert that food security in Java was unarguably compromised by the economic crisis, providing further support to the view that the crisis was strongly felt in this region.

With these considerations in mind, the kernel matching estimator and the bootstrap were applied to the subsample of population that was residing on the island of Java in 1997. 1000 bootstrap replications were implemented again in each case, and the evidence on the effectiveness of migration in response to the crisis is presented in Table 19; the evidence found in Java is consistent with the evidence from the entire sample in the main specification. The only significant discrepancy applies again to balanced households from Java that, similarly to non-rice producers, did not experience significant gains in their consumption growth in the long run.

Table 19 – ATET of Migration – Java subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Short Run (97-00)	0.039	-0.166	0.127	0.205	-0.115
	(0.39)	(3.13)***	(3.67)***	(3.16)***	(0.82)
Long Run (97-07)	0.169	0.013	0.126	0.149	0.129
	(1.59)	(0.22)	(3.59)***	(2.55)**	(0.93)
Treated	73	252	774	345	45
Untreated	3,290	3,290	3,290	3,263	2,876
N	3,363	3,542	4,064	3,608	2,921

* p<0.1; ** p<0.05; *** p<0.01

On the contrary, Table 20 presents the result of this analysis for the rest of population in Indonesia; this evidence generally fails to confirm the results found in Java (as well as in the main specification). The signs of the coefficients remain unchanged and significant effects of migration are still found for both net importers of migrants and split-off households by 2000; in these cases, results are significant at any conventional level and point estimates are even higher than those found in Java. However, the short run ATET of migration for net exporters of migrants is no longer significant, suggesting that the net decrease in household size due to migration did not play a key role for households in areas that were hit less hard by the crisis; in the long run, none of the coefficients are significant. This suggests that migration, especially outflows of migration, in response to the crisis were mostly relevant in Java, where the crisis was felt immediately and most substantially.

Table 20 – Bootstrapped ATET of Migration – Outside Java subsample

ATET	Balanced ATET	Net-Importers ATET	Net-Exporters ATET	Split-Offs ATET	Moving ATET
Short Run (97-00)	-0.048	-0.222	0.022	0.222	0.036
	(0.41)	(3.28)***	(0.55)	(2.58)***	(0.20)
Long Run (97-07)	0.041	-0.062	0.034	0.036	0.070
	(0.33)	(0.73)	(0.77)	(0.45)	(0.37)
Treated	51	160	565	167	31
Untreated	2,271	2,271	2,271	2,271	2,271
N	2,322	2,431	2,836	2,438	2,302

* p<0.1; ** p<0.05; *** p<0.01

7. Conclusion

This paper has investigated the consequences of mobility across districts and changes in living arrangements that took place in Indonesia in the years of the East Asian crisis, arguably the hardest hit country by the economic shock in the late 1990s. Given the high degree of heterogeneity in the magnitude of this shock across regions in Indonesia, the high rates of internal mobility in the country in those years and the richness of data available, the East-Asian crisis in Indonesia constitutes an ideal context to examine the effectiveness of mobility across regions and modifications in living arrangements in response to an economic shock. Using a semi-parametric approach, the effects of migration on the rate of consumption growth are investigated both by 2000, immediately after the end of the crisis, and by 2007/8, ten years after the peak of the crisis. The evidence from 2000 suggests that significant implications of migration are immediately visible in most of the cases: net outflows of migration are often associated with net gains, both for the origin household and for the newly formed household. These conclusions are suggested by the positive and significant coefficients associated to the ATET of migration for both net-exporters of migrants and for newly-formed split-off households. On the contrary, net receipt of migrants results in an immediate decrease in living standards in the household. By 2007/8, net-exporters of migrants and newly formed split-off households are still found to be positively affected by migration; net gains from mobility in the years of the crisis are still visible ten years after the onset of the crisis. On the other hand, net-importers are no longer found to be negatively affected by the net receipt of migrants during the crisis, suggesting that the distress caused by the increase in household size due to migration was only temporary and did not produce any lasting negative consequences.

Since migration, in this study, is not treated as a once-and-for-all decision and people are allowed to move from 2000-2008, the results from 2007/8 are of particular relevance, because they attest to the important role of migration in times of crisis, rather than migration per se. What the evidence suggests is that, on average, mobility in times of crisis is beneficial and benefits are likely to persist over time. Net inflows of migration are unlikely to generate long-lasting negative effects, since, arguably, further adjustment is likely to take place among these households in the following years in response to the initial distress. Finally, in the attempt to isolate in a number of ways the economic crisis, and the migration that occurred in response to it, we performed the semi-parametric IV estimation of local average treatment effect (LATE), and the analysis was also conducted on several different subsamples. Due to the small fraction of compliers in our data, precise estimates of the LATE of migration for compliers could not be obtained. However, the conclusions on the short-run effectiveness of migration appear robust to a number of different specifications and they are observed for almost all subsamples analysed in this study. A weaker degree of consistency is found for the evidence from 2007/8. More precisely, consistency in the results is found especially when the analysis is restricted to the subsample that did not move in the pre-crisis period, to the subsample where no new births occurred in the years that followed the crisis, to the subsample of non-rice producers and to the subsample that resided in Java in 1997, arguably the hardest hit region in the country. In each case, these subsamples of households were more likely to experience migration in response to the crisis than their counterpart. Interestingly, weaker consistency in the effects of migration is found for the counterfactual subsamples, namely the subsample of households that already proved mobile in the years prior to the crisis, the subsample where new births occurred after the years of the crisis and the subsample that resided outside Java in 1997.

8. References

Abadie, A. (2005) "Semi-parametric Difference-in-Differences Estimators," Review of Economic Studies, 72, 1-19.

Abadie, A. and Imbens, G. (2008) "On the Failure of the Bootstrap for Matching Estimators," Econometrica, 76, 1537-1557.

Akita, T., A. Alisjahbana, (2002). "Regional Income Inequality in Indonesia and the Initial Impact of the Economic Crisis." Bulletin of Indonesian Economic Studies, 38(2), pp. 201-222.

Angrist, J., and S. Pischke (2008), Mostly Harmless Econometrics: An Empiricists' Companion, Princeton University Press, Princeton, NJ.

Bhattacharya, J. and Vogt, W.B. (2007). "Do Instrumental Variables Belong in Propensity Scores?", NBER Technical Working Papers 0343, National Bureau of Economic Research, Inc.

Breman, Jan (2000). “The Impact of the Asian Economic Crisis on Work and Welfare in Village”. (mimeo)

Caliendo, M. and Kopeinig, S. (2008), SOME PRACTICAL GUIDANCE FOR THE IMPLEMENTATION OF PROPENSITY SCORE MATCHING. Journal of Economic Surveys, 22: 31–72.

Chen, Daniel L. (2010) “Club Goods and Group Identity: Evidence from Islamic Resurgence during the Indonesian Financial Crisis." Journal of Political Economy 118(2):300-354.

Chen, S., Mu, R. and Ravallion, M. (2009) “Are there lasting impacts of aid to poor areas?” Journal of Public Economics , 93(3-4), pp.512–528

Collier, P., Dercon, S., (2009). “African agriculture in 50 years: Smallholders in a rapidly changing world.” Expert Meeting on How to Feed the World in 2050, Food and Agriculture Organization, 24–26 June 2009, Rome.

Dawe, David (2008) ‘Can Indonesia trust the world rice market?’, Bulletin of Indonesian Economic Studies 44 (1): 115–32, in this issue.

Peter R. Fallon & Robert E. B. Lucas, (2002). "The Impact of Financial Crises on Labor Markets, Household Incomes, and Poverty: A Review of Evidence," World Bank Research Observer, World Bank Group, vol. 17(1), pages 21-45.

Foster, A., & Rosenzweig, M. (2002). “Household division and rural economic growth.” Review of Economic Studies, 69(4), 839–869.

FPSA (Food Policy Support Activity) (2003) A note on rice policy, Food Policy Support Activity, Development Alternatives Inc., Jakarta, mimeo.

Frankenberg, Elizabeth; James P. Smith and Duncan Thomas. (2003). “Economic Shocks, Wealth and Welfare,” J. Human Res. 32:2, pp.280–321.

Friedman, Jed and James Levinsohn. (2002). "The Distributional Impacts of Indonesia's Financial Crisis on Household Welfare: A "Rapid Response Methodology," World Bank Econ. Rev. 16, pp. 397-423.

Frölich, M. (2004), “Finite Sample Properties of Propensity-Score Matching and Weighting Estimators,” Review of Economics and Statistics, 86, 77-90.

Frölich, M. (2007), “Nonparametric IV estimation of local average treatment effects with covariates,” Journal of Econometrics, vol. 139(1), pages 35-75, July.

Frölich, M. and Lechner, M. (2006), “Exploiting Regional Treatment Intensity for the Evaluation of Labour Market Policies,” IZA working paper No. 2144.

Gilligan, Daniel O., Hanan Jacoby and Jaime Quizon (2000). “The Effects of the Indonesian Economic Crisis on Agricultural Households: Evidence from the National Farmers Household Panel Survey (PATANAS)” (Part of a joint study by the World Bank and the”. Report by the Center for Agro-Socioeconomic Research and the World Bank.

Ham, J., Li, X., Reagan, P. (2011) Matching and semi-parametric IV estimation: A distance-based measure of migration, and the wages of young men. Journal of Econometrics 161: 208–227

Heckman, J., Ichimura, H. and Todd, P. (1997), “Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Program,” Review of Economic Studies, 64, 605-654.

Heckman, J., Ichimura, H. and Todd, P. (1998), “Matching as an Econometric Evaluation Estimator,” Review of Economic Studies, 65, 261-294.

Heckman, J,, R, Lalonde, J, Smith, (1999), "The Economics and Econometrics of Active Labor Market Programs," In Handbook of Labor Economics, Vol, 3, ed, O, Ashenfelter, and D, Card, Chapter 31, Amsterdam: Elsevier

Hirano, K., Imbens, G. and Ridder, G. (2003), “Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score,” Econometrica, 71, 1161-1189.

Holland, P.W., “Statistics and Causal Inference,” Journal of the American Statistical Association 81:396 (1986), 945–970, with discussion.

Huber, M., M. Lechner, and C. Wunsch (2013): “The Performance of Estimators Based on the Propensity Score,” Journal of Econometrics.

Imbens, G. W., “The Role of the Propensity Score in Estimating Dose-Response Functions,” NBER technical working paper no. 237 (1999). Also in Biometrika 87 (2000), 706–710.

Imbens, G., and J. Angrist (1994), “Identification and Estimation of Local Average Treatment Effects,”Econometrica, Vol. 61, No. 2, 467-476.

Lechner, M., (2000). “An evaluation of public-sector-sponsored continuous vocational training programs in East Germany”. Journal of Human Resources 35, 347–375.

Lechner, M., (2001). “Identification and Estimation of Causal Effects of Multiple Treatments under the Conditional Independence Assumption” (pp. 43–58), in M. Lechner and F. Pfeiffer (Eds.), Econometric Evaluations of Active Labor Market Policies in Europe (Heidelberg: Physica/Springer, 2001a).

Lechner, M, (2002), “Program Heterogeneity and Propensity Score Matching: An Application to the Evaluation of Active Labor Market Policies,” The Review of Economics and Statistics, 84(2): 205-220, May

Lechner, M., (2009). Long-run labour market and health effects of individual sports activities. The Journal of Health Economics 28, 839–854.

Lucas, Robert E.B. (1997) “Internal Migration in Developing Countries,” Handbook of Population and Family Economics, Chapter 13, page 721-798

Lucas, R.E.B. (2004), ”International migration regimes and economic development”, Report for the Expert Group on Development Issues (EGDI), Swedish Ministry of Foreign Affairs.

McCulloch, Neil (2008): RICE PRICES AND POVERTY IN INDONESIA, Bulletin of Indonesian Economic Studies, 44:1, 45-64

McKenzie, David (2011) “Beyond Baseline and Follow-up: The Case for more T in Experiments”, World Bank Policy Research Working Paper no. 5639.

Mortensen, D.T., 2005. “Wage Dispersion: Why Are Similar Workers Paid Differently?” MIT Press, Cambridge Massachusetts

Ravallion, M. (2005). Evaluating anti-poverty programs. Policy Research Working Paper No. 3625. Washington, DC: World Bank.

Ravallion, M. (2008). Evaluating anti-poverty programs. In Handbook of development economics, Vol. 4, ed. T. Schultz and J. Strauss. Amsterdam: North Holland.

Ravallion, M. and M. Lokshin, (2007). “Lasting Impacts of Indonesia’s Financial Crisis,” Economic Development and Cultural Change, 56(1): 27-56.

Rosenbaum, Paul R. (2002). Observational Studies, 2d ed. New York: Springer.

Rosenbaum, P. R. and Rubin, D. B. (1983), “The Central Role of the Propensity Score in Observational Studies for Casual Effects,” Biometrika, 70, 41-55.

Rosenbaum, P.R. and Rubin, D.B. (1985), "Constructing a Control Group Using Multivariate Matched Sampling Methods that Incorporate the Propensity Score", The American Statistician 39(1), 33-38.

Rosenzweig, M. R. (2010), Global Wage Inequality and the International Flow of Migrants, in R. Kanbur and M. Spence (eds.), Equity and Growth in a Globalizing World , World Bank, Washington DC

Roy, A. D., “Some Thoughts on the Distribution of Earnings,” Oxford Economic Papers 3 (June 1951), 135–146.

Rubin, D. B., “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies,” Journal of Educational Psychology 66 (1974), 688–701.

Rubin, D. B., “Assignment to Treatment Group on the Basis of a Covariate,” Journal of Educational Statistics 2 (1977), 1–26.

Safir, A., & Beegle, K. (2009). Migration After the 1997 Indonesian Crisis: Did Economic Losers Move? (not published version).

Sianesi, B. (2004). “An Evaluation of the Active Labour Market Programmes in Sweden”, The Review of Economics and Statistics, 86(1), 133-155.

Smith, J., Todd, P., (2005). Does matching address LaLonde’s critique of nonexperimental estimators? Journal of Econometrics 125, 355–364.

Smith, J. P.; Thomas, D.; Frankenberg, E.; Beegle, K. and Teruel, G. (2002) “Wages, Employment and Economic Shocks: Evidence from Indonesia,” J. Population Econ. 15, pp. 161–93.

Strauss, J., Beegle, K., Agus Dwiyanto, Yulia Herawati, Daan Pattinasaramy, Elan Satriawan, Bondan Sikoki, Sukamdi, and Firman Witolear (2004) ‘Indonesian living standards before and after the financial crisis’. Singapore: Institute of Southeast Asian Studies.

Studdert LJ, Frongillo EA, Valois P. Household food insecurity was prevalent in Java during Indonesia’s economic crisis. J Nutr (2001);131:2685–91

Sumarto, Sudarno, Anna Wetterberg, and Lant Pritchett (1998), ‘The Social Impact of the Crisis in Indonesia: Results from a Nationwide Kecamatan Survey’, SMERU Report, December, Social Monitoring & Early Response Unit, Jakarta.

Teal, F. (2011). “The Price of Labour and Understanding the Causes of Poverty”. Labour Economics 18 (2011) S7-S15

Thomas, D.; Beegle, K.; Frankenberg, E.; Sikoki, B.; Strauss, J. and Teruel, G. (2004) ‘Education in a Crisis’, Journal of Development Economics 74 (1): 53-85.

Thomas, D. and Frankenberg, E. (2007) ‘Household responses to the financial crisis in Indonesia: longitudinal evidence on poverty, resources and well-being’. in (A.Harrison, ed.), Globalization and Poverty, pp. 517–68, Chicago: University of Chicago Press.

Thomas D., Frankenberg, E. and Beegle, K. (2000). "Labor Market Transitions of Men and Women During an Economic Crisis: Evidence from Indonesia." Following: Thomas, Duncan, E. Frankenberg and K. Beegle. (2003). ” Labor market transitions of men and women during an economic crisis: Evidence from Indonesia”, in R. Anker, B. Garcia and A. Pinelli (eds.) Women in the labour market in changing economies: Demographic issues, Oxford University Press.

Timmer, C. Peter (2005) Operationalizing pro-poor growth: a country case study of Indonesia, Poverty Reduction and Economic Management, World Bank, Washington DC, mimeo.

Walmsley, Terrie L. & Winters, L. Alan, 2005. "Relaxing the Restrictions on the Temporary Movement of Natural Persons: A Simulation Analysis," Journal of Economic Integration, Center for Economic Integration, Sejong University, vol. 20, pages 688-726.

Witoelar, Firman. (2005). “Inter-household Allocations within Extended Family: Evidence from the Indonesia Family Life Survey.” Discussion Paper no. 912, Yale University Economic Growth Center.

Yamauchi, F., Liu, Y. (2012) “Impacts of an Early Stage Education Intervention on Students’ Learning Achievement: Evidence from the Philippines”, Journal of Development Studies

APPENDIX I - Descriptive Statistics of the Sample in 1997 – Units are percent for binary variables. For continuous variables, mean values and standard deviation (in brackets) are listed

Table 21 – Average Value of Variables at the Micro-economic (Household) level in 1997

Description of Variables and Summary Statistics - by migration category	Non-Migrant Households			Balanced-Migrant Households			Net Migrant-Importer Households			Net Migrant-Exporter Households			Split-Off Households			Moving Households		
	N	Mean	Std Dev	N	Mean	Std Dev	N	Mean	Std Dev	N	Mean	Std Dev	N	Mean	Std Dev	N	Mean	Std Dev
Household size	5802	5.49	2.57	130	6.59	2.49	429	5.48	2.49	1400	6.41	2.44	538	7.15	3.07	77	3.75	1.73
Number of births since 1993	5802	0.28	0.53	130	0.19	0.43	429	0.16	0.43	1400	0.26	0.51	538	0.27	0.55	77	0.44	0.70
Number of deaths since 1993	5802	0.11	0.34	130	0.10	0.39	429	0.13	0.35	1400	0.12	0.35	538	0.09	0.31	77	0.05	0.22
Consumed any self-produced rice in 1997	5802	0.22		130	0.21		429	0.18		1400	0.20		538	0.17		77	0.12	
Male head	5802	0.84		130	0.77		429	0.77		1400	0.83		538	0.86		77	0.84	
Age of head	5802	48.50	28.53	130	49.40	13.00	429	52.95	47.46	1400	49.96	28.21	538	51.25	11.59	77	37.87	12.20
Age of head (squared)	5802	3166.22	26118.48	130	2607.95	1292.81	429	5051.24	47973.21	1400	3290.64	26581.19	538	2760.16	1241.92	77	1580.96	1078.11
Married head	5802	0.84		130	0.77		429	0.78		1400	0.84		538	0.85		77	0.84	
Muslim head	5802	0.90		130	0.88		429	0.87		1400	0.89		538	0.92		77	0.94	
Head has elementary education	5802	0.54		130	0.46		429	0.49		1400	0.52		538	0.46		77	0.39	
Head has higher education	5802	0.24		130	0.32		429	0.26		1400	0.27		538	0.34		77	0.45	
Head has college/university degree	5802	0.04		130	0.10		429	0.07		1400	0.06		538	0.09		77	0.09	
Presence of <25 yrs HH members	5802	0.87		130	0.94		429	0.79		1400	0.94		538	0.93		77	0.84	
HH affected by migration between 1993-97	5802	0.22		130	0.68		429	0.52		1400	0.47		538	0.58		77	0.40	
Health facilities in the kabupaten	5802	0.99		130	0.97		429	0.98		1398	0.99		537	0.99		77	0.95	
Presence of recent graduates	5797	0.26		130	0.39		429	0.29		1400	0.34		538	0.42		77	0.17	
Presence of self-employed	5771	0.62		130	0.58		427	0.59		1396	0.62		537	0.59		77	0.32	
Presence of public workers	5771	0.11		130	0.13		427	0.13		1396	0.13		537	0.17		77	0.12	
Presence of private workers	5771	0.48		130	0.47		427	0.43		1396	0.47		537	0.53		77	0.53	
Presence of family workers	5771	0.18		130	0.19		427	0.15		1396	0.21		537	0.16		77	0.01	
Healthy HH (self-reported by HH adults)	5770	0.87		130	0.94		427	0.86		1395	0.92		537	0.92		77	0.82	
Farm business owned (100%)	5801	0.34		130	0.36		429	0.30		1400	0.35		537	0.25		77	0.13	
Any member borrowed money in last year	5771	0.43		130	0.43		427	0.43		1395	0.48		537	0.51		77	0.38	
Any member has parents away from HH	5771	0.75		130	0.73		427	0.66		1395	0.76		537	0.78		77	0.90	
Any member has siblings away from kabupaten	5771	0.66		130	0.77		427	0.75		1395	0.77		537	0.79		77	0.90	
Rural HH	5800	0.58		130	0.47		428	0.50		1400	0.57		538	0.39		77	0.31	
Proportion of urban population	5596	0.41		124	0.44		415	0.45		1345	0.40		513	0.52		76	0.63	
Proportion of unemployed 15-60	5596	0.04		124	0.04		415	0.04		1345	0.04		513	0.04		76	0.05	
Proportion of self-employed	5596	0.27		124	0.26		415	0.26		1345	0.27		513	0.24		76	0.22	
Proportion in agricultural sector	5596	0.26		124	0.26		415	0.24		1345	0.26		513	0.21		76	0.16	
Proportion living under poverty line	5596	0.09		124	0.08		415	0.08		1345	0.09		513	0.07		76	0.06	
Gini coefficient	5596	0.26		124	0.27		415	0.27		1345	0.26		513	0.28		76	0.30	

APPENDIX II – Covariate Balancing and Performance of the Kernel matching estimator

Table 22 – Indicators of Covariate Balancing, Before and After Matching, by Migration Category

Migration Category VS Non-Migrants (1)	No. Treated (2)	No. Nontreated (3)	Treated as % of Nontreated (4)	Probit ps.-R^2 Before (5)	Probit ps.-R^2 After (6)	Pr > χ^2 After (7)	Median Bias Before (8)	Median Bias After (9)	No. Lost to CS After (10)
Balanced Households	124	5,561	2.22	0.149	0.006	1.000	11.3	1.8	3
Net-Importer Households	412	5,561	7.41	0.085	0.001	1.000	8.8	0.8	1
Net-Exporter Households	1,339	5,561	24.08	0.078	0.001	1.000	4.8	0.5	3
Split-Off Households	512	5,561	9.21	0.144	0.002	1.000	16.8	0.9	3
Moving Households	76	5,561	1.37	0.241	0.015	1.000	35.5	2.6	5

Notes:
(1): Migration category matched to the non-migrant category.
(2): Number of treated (that is, households that fell into that migration category from 1998-2000).
(3): Number of potential comparisons (that is, that did not experience any migration event in the years 1998-2000).
(4): Treated as percentage of potential comparisons.
(5): Pseudo-R^2 from probit estimation of the conditional probability to receive the treatment; this gives an indication of how well the covariates X explain the probability to produce/receive a migrant, or to move.
(6), (7), (9), and (10) are post-matching indicators based on the kernel estimator with replacement (Epanechnikov curve and 0.5% bandwidth).
(6): Pseudo-R^2 from a probit of D on X on the matched samples, to be compared with (5).
(7): P-value of the likelihood ratio test after matching. The joint significance of the included covariates is always rejected (Before matching it was never rejected at any significance level, with Pr > χ^2 = 0.0000 always).
(8), (9): Median absolute standardized bias before and after matching, median taken over all the covariates. Following Rosenbaum and Rubin (1985), for a given covariate X, the standardized difference before matching is the difference of the sample means in the full treated and nontreated subsamples as a percentage of the square root of the average of the sample variances in the full treated and nontreated groups. The standardized difference after matching is the difference of the sample means in the matched treated (that is, falling within the common support) and matched nontreated subsamples as a percentage of the square root of the average of the sample variances in the full treated and nontreated groups:

$$B_{before}(\mathrm{X}) \equiv 100 \cdot \frac{\overline{X}_1 - \overline{X}_0}{\sqrt{[V_1(X)+V_0(X)]/2}}, \; B_{after}(\mathrm{X}) \equiv 100 \cdot \frac{\overline{X}_{1M} - \overline{X}_{0M}}{\sqrt{[V_1(X)+V_0(X)]/2}}$$

Note that the standardization allows comparisons between variables X and, for a given X, comparisons before and after matching.
(10): Number of treated individuals falling outside of the common support (based on a bandwidth of 0.5%).

APPENDIX III - Plausibility of the Conditional Independence Assumption (CIA) and the choice of conditioning variables

As stated in Section 4, for the DiD matching estimator to be valid, it requires

$$E(\log(C_{0t}) - \log(C_{0t'}) \mid X, D = 1) = E(\log(C_{0t}) - \log(C_{0t'}) \mid X, D = 0),$$

where X is an appropriate set of observable variables unaffected by the treatment. In this equation, the CIA assumption is stated in terms of the before-after consumption evolution instead of levels; this means that, conditional on X, the migrant households' potential consumption growth Y^0 had they not experienced migration would be the same as non-migrant households' potential consumption growth. This is a fairly strong condition, although it is important to bear in mind that, since the ATET (not the ATE) is being estimated, Y^1, and therefore the migration consumption gain $\Delta = Y^1 - Y^0$, is allowed to depend on D (Ham et al. 2011). In their discussion of the plausibility of the CIA in the context of their study, Ham et al. (2011) consider the case in which two observations have the same potential consumption gain if no migration takes place (Y^0). The reason why one individual moves and the other does not lies in the different Y^1, which is greater than Y^0 for the individual that decides to move and smaller than Y^0 for the individual that decides not to move. However, the authors also point out that, in this scenario, the dependence of Y^1 on D ought to be ruled out if Y^0 depends on Y^1. In our analysis, an example of a possible dependence of Y^0 on Y^1 is the case where household i's members are aware of the level of p/capita consumption after the receipt/departure of a migrant (i.e., Y^1), and are able to use this information to influence the level of p/capita consumption in the absence of migration (i.e., Y^0). We hope and believe that this was generally not the case and, if anything, the influence of Y^1 on Y^0 was small; two considerations support this claim. First, in the context of this study, neither Y^0 nor Y^1 were clearly visible to the members of household i until either of them realized; if it would have been hard to know this information a-priori in normal circumstances, the nature of the crisis experienced by Indonesian households in the years under observation makes the dependence of Y^0 on Y^1 even more implausible. Second, even if Y^0 and Y^1 were known by the members of household i, there is no apparent way in which Y^1 would determine Y^0; for example, in the bilateral comparison between net-exporters and non-migrant households, one could argue that, e.g., in an extended family context, knowledge of the potential gain from sending a migrant away from household i may enable household i to put pressure on the remaining households in the extended family and, thus, lead household i to either send the migrant away or, alternatively, to receive compensation (e.g., remittances) from other households for not sending away the migrant. Although this is possible in principle, it is not clear why, in practice, this should happen, i.e., in a risk- or income-sharing context, it is not clear why one would prefer to compensate for not allowing mobility, rather than just allowing it[25]. If this actually occurred in some cases, we would not expect this to be general practice among Indonesian households in the years

[25] It is important to notice that, in this study, all the discussion focuses on mobility of adults (>15 years old), hence, in principle, not dependents, but rather individuals capable to produce income. This is relevant because it makes even less clear why one would not encourage a potential worker to leave the initial location and move to a more favourable labour market. Instead, (ruling out child labour) such gain from moving would not apply to a child that was to move to a better labour market.

of the crisis. Therefore, even in this case, we would expect the effect of the potential dependence of Y^0 on Y^1 to be small. For these reasons, we believe we can allow the dependence of Y^1 on D in this context.

From a methodological point of view, this is relevant because it implies that, among the conditioning variables, it suffices to control for all factors expected to affect both the migration decision (i.e., D) and the potential consumption growth in the absence of migration (i.e., Y^0), while we do not need to control for factors that affect both Y^1 and D. Looking at, e.g., the bilateral comparison between split-off households and non-migrant households, such factors would include pull factors from the destination kabupaten (Ham et al., 2011); however, since we allow the dependence of Y^1 on D, we do not need to control for such factors.

The CIA requires detailed knowledge of the elements that determine participation, as well as access to proper data in order to capture those factors that are likely to influence both participation and outcome (Sianesi, 2004). Hence, the plausibility of the CIA should be discussed in relation to the richness of the available dataset, as well as the process of selection into migration. In fact, in the present study, the choice of the covariates X can benefit from the abundance of information contained in the IFLS and the SUSENAS datasets; from the former, relevant information was extracted at the household level, and from the latter, information was extracted at the kabupaten level. Detailed knowledge of the elements that determined participation, i.e., the migration decision in Indonesia in those years, mostly derives from three elements: first, the large quantitative evidence in the economics literature on the determinants of migration. Building on such literature, particularly bearing in mind the contribution of the New Economics of Labour Migration (NELM) theory, a set of information at the household level were included in the analysis[26]. Secondly, building on the evidence on the effects of the East Asian Crisis on the labour market and geographic mobility in Indonesia (e.g., Thomas, Beegle and Frankenberg, 2000; Smith et al., 2002; Frankenberg et al., 2003; Thomas and Frankenberg 2007; Safir and Beegle, 2009), a set of information that turned out to be relevant in this specific context was also included in the analysis. Finally, a careful study of the nature of the crisis itself, and the high degree of its spatial heterogeneity (Ravallion and Lokshin, 2007; Sumarto, Wetterberg, and Pritchett, 1998) also drove the choice of the conditioning variables.

A rich set of household's characteristics, household head's characteristics and household members' characteristics were included in the analysis, all of which calculated in the pre-crisis (and therefore, pre-treatment) period. By doing so, we control for a set of demographics, for several dimensions of human capital and for a number of direct indicators of heterogeneity across households that are important determinants of both the migration decision and the socio-economic prospects of the household. Concerning demographics, elements such as age, gender and health conditions are known to be strong determinants of the socio-economic outcome of the household. For this reason, this information was included in the analysis. Concerning human capital information, information on the educational achievement of the household head was included, as well as on the presence of recent graduates in the household. Information on the sector of primary occupation for all adult members of the household was also considered: the sector of occupation was correlated with the extent to which individuals were affected by the crisis (see, e.g., Safir and Beegle, 2009), and therefore this was likely to influence the chances of mobility to take place, as well as the evolution of living standards in the

[26] The NELM theory was first not to consider migration to be an individual decision, but rather a choice made at the household level

household. Information on the ownership of a farm by the household was included in the analysis as a proxy of vulnerability of the household to the crisis; indeed, in the NELM view of migration, minimizing risk is a key driver of the decision to send a migrant away from the household. Vulnerability is also likely to influence, more broadly, the economic decisions of the household, hence, by extension, the evolution of its living standards. Hence, we thought it sensible to include an indicator of the level of vulnerability of the household. A similar rationale applies to the choice to include in the equation information on access to the formal credit market. We also include information intended to capture, ceteris paribus, the propensity of the household to react to the crisis through migration, as well as the evolution of its living standards. In fact, the recent migration history of the household is captured by a dummy variable that distinguishes households that revealed more mobile in the years that preceded the crisis from those that did not do so; information is also included on the occurrence of births and deaths in the pre-crisis period (as presence of a newly born child may make migration less likely to occur), presence of young members (< 25 years old) and, finally, presence of siblings away from the kabupaten.

Furthermore, building on the evidence on the specific dynamics of the crisis in Indonesia, variables such as own production of rice in 1997 and a rich set of macroeconomic variables at the kabupaten level are also included in the analysis. The former is included building on the evidence that rice producers not only have been isolated, but they may have even benefited from the crisis in Indonesia (Chen 2010). The latter are included with relation to the local labour market conditions and their inclusion is motivated by the strong evidence found on the geographic diversity of the crisis, where richer and more well-connected but also more unequal regions of the country were found to be more hardly hit by the crisis (Ravallion and Lokshin, 2007); hence, e.g., the proportion of population living under the local poverty line, the proportion of population employed in the agricultural sector, the proportion of self-employed population, the proportion of urban households and the Gini coefficient were all among the variables included in the equation; all these variables were calculated at the kabupaten level. Finally, within the kabupaten, we also controlled for whether the household was located in urban or rural areas in 1997. Omitting to control for this geographical information, would likely result in ignoring some fundamental determinants of both mobility and evolution of living standards in Indonesia in those years. Therefore, information on the kabupaten of residence in 1997 was attached to each household and it was included in the analysis; instead, since we do not need to control for the factors that affect both Y^1 and D (e.g., pull factors of migration), we attached to each household only information on the kabupaten of origin.

To conclude, in this analysis the CIA is not required, unlike in most of the cases, to hold in terms of a once-and-for-all decision: migrating versus never migrating. Instead, the analysis requires the CIA to hold only in terms of a decision as to move, or to wait (and, possibly, move in later years) in response to the crisis. After eliminating the effect of temporarily invariant components (e.g., ability), for households that are similar in all individual and local characteristics described, the decision to produce/receive migrants, or to move entirely, in those years, rather than staying and reacting in a different way needs to be random, in the sense that it depends on factors unrelated to future potential outcomes in the absence of the treatment (i.e., Y^0). As aforementioned, examples of such factors are Y^1, moving costs, or labour market conditions at destination for split-offs and moving households. The sensitivity of the results presented in Table 9 with respect to deviations from this identifying assumption was checked using Rosenbaum bounds (Rosenbaum 2002); although results are not reported here, the evidence appeared to be robust to the presence of omitted confounding factors.