The intensity of the Rwandan genocide: Measures from the gacaca records*

Marijke Verpoorten[†]

University of Leuven

September 2011

Abstract

This article illustrates how fine continuous and categorical measures of genocide intensity can be derived from the records of the Rwandan transitional justice system. The data, which include the number of genocide suspects and genocide survivors across 1484 administrative sectors, are highly skewed and contain a non-negligible number of outlying observations. A number of genocide proxies are subjected to Principal Component Analysis (PCA) to obtain a genocide index, and the effect of survival bias on this index is reduced by augmenting the set of genocide proxies subjected to PCA with the distance from an administrative sector to the nearest mass grave. Finally, the administrative sectors are divided into distinct categories of low, moderate and high genocide intensity by means of Local Indicators of Spatial Auto-Correlation (LISA) that allow identifying significant high-high and low-low clusters of genocide intensity.

1 Introduction

The micro-level research on armed conflict has exploded over the past decade. Besides a steady increase in the number of studies, there have been considerable improvements in

*I received helpful comments from Koen Decancq, Giacomo De Luca, Romain Houssa and Pieter Serneels. I am indebted to Bert Ingelaere for bringing to gacaca data to my attention. I owe thanks to the Rwandan National Service of Gacaca Jurisdiction and the Rwandan National Census Service for making available the data used in this study. All errors and opinions expressed remain my own.

[†]Research scholar of the fund for Scientific Research - Flanders, Belgium (FWO), Center for Economic Studies, Centre for Institutions and Economic Performance LICOS, KULeuven. e-mail: marijke.verpoorten@econ.kuleuven.be

methodology. In particular, scholars have increasingly devoted attention to identifying rich micro-level measures of conflict intensity (e.g. Restrepo, Spagat, and Vargas (2006), Raleigh, Linke, Hegre, and Karisen (2010)). This is no coincidence because the identification of micro-level causes and consequences of armed conflict stands or falls with the conflict intensity measure used.

This article aims to promote the use of rich micro-level conflict intensity measures in two ways. First, it provides easy access to and a critical evaluation of the data released by the gacaca courts, i.e. the transitional Rwandan justice system in charge of judging 1994 genocide suspects. Second, it presents fine continuous and categorical measures of various aspects of the genocide, e.g. genocide participation and the genocide's death toll.

The gacaca records include four types of information: (1) the number of accused persons living in the country; (2) the number of genocide survivors living in the country; (3) the number of accused persons who are not living in the country; (4) the number of persons who committed genocide and who passed away. The two latter types are only available at the district level. The first two types of information are available for 1484 sectors, which are, after the cells, the lowest codified administrative unit in Rwanda¹.

The data was released in 2007 in pdf format on the website of gacaca². After converting the data into spreadsheet format, I subject them to a critical examination. In particular, I evaluate their overall reliability through a comparison with data from other sources, including the number of persons imprisoned (Office of the Prosecutor (2002)), an estimate of the number of perpetrators by Straus (2004), and a 2006 census of genocide survivors (Government of Rwanda (2008)). Such a critical examination is required because, as gacaca proceeded, its operation was criticized for lack of objectivity due to (political) manipulation (Ingelaere (2009), Longman (2009), Pitsch (2002), Wolters (2005)).

After this overall data quality check, I transform the data in several ways in order to provide researchers with a menu of different measures that capture (i) genocide participation, (ii) genocide survivorship defined as the survival of close relatives of genocide victims, i.e. widowed, orphaned or disabled genocide survivors, and (iii) the genocide death toll among Tutsi. The genocide participation proxies are useful for researchers studying the determinants of the involvement of the civilian population in the execution of the genocidal campaign; the genocide survivorship proxies may be used to study the legacy of genocide, e.g. the impact

¹At the time of the genocide, Rwanda was divided in 10 prefectures, 145 communes, 1565 sectors and more than 9000 cells.

²<http://www.inkiko-gacaca.gov.rw/>

of genocide on social capital in communities that are home to both surviving victims and perpetrators of the genocide; and the genocide death toll may be used to analyze and explain the intensity of ethnic cleansing.

The availability of fine measures of genocide intensity can mean an important step forward in the research on the causes and consequences of mass killings. Given that the proposed genocide intensity measures are at the sector level, they can be matched with existing data that use the sector as a sample unit, e.g. the 1984-1991 FSRP/DSA³ agricultural household surveys, the 2000/2001 and 2005/2006 Integrated Household Living Conditions Surveys, and the 1991 and the 2002 population censuses. In addition, the sector level genocide measures can contribute to the sampling process of new surveys that seek to stratify their household sample across high and low genocide intensity areas, or they can be related to a number of geographical measures, e.g. community distance to the nearest main road, important urban centers or national borders, or population density.

So far, there have been a number of empirical micro-level studies on both the causes and consequences of the Rwandan genocide, but to the best of my knowledge, only Yanagizawa (2010) has used the gacaca data described in this article, in particular for studying the impact of radio broadcasting of hate messages on the involvement of the civilian population in the killings. The transformation of the data presented here goes at least four steps further, by (1) providing other measures of genocide intensity, besides participation, (2) identifying outliers and anomalous values, (3) correcting for survival bias, and by (4) deriving categorical measures of genocide intensity in a non-arbitrary way.

The next section provides an overview of the gacaca data as well as a first data quality check. Section 3 explains how a set of meaningful measures can be derived from the gacaca records. Section 4 includes a technical discussion on the skewness and outlyingness of the data. Section 5 constructs a genocide index by subjecting a set of genocide proxies to Principal Component Analysis (PCA). Section 6 derives a categorical variable for genocide intensity using Local Indicators of Spatial Association (LISA). Section 7 concludes.

³Food Security Research Project /Division des Statistiques Agricoles

2 The Gacaca Records

2.1 Overview

In 2005, the gacaca courts were stepping in the first phase of their activities, i.e. the phase of collecting information. During weekly sessions with compulsory attendance of all community members, lists were made of victims, suspects and survivors⁴. Part of the results achieved during this phase were made public in the course of 2007. The released sector level data include the number of genocide suspects in a sector, classified in three groups, and the number of genocide survivors, classified in five groups⁵.

- Genocide suspects
 - Category 1: accused of planning, organizing or supervising the genocide, or committing sexual torture
 - Category 2: accused of killings or other serious physical assaults
 - Category 3: accused of looting or other offences against property
- Genocide survivors
 - Widowed
 - Orphaned
 - Disabled
 - Male
 - Female

The exact legal definitions of the suspect categories can be found in the appendix. The first category of alleged genocide perpetrators has to be referred to national criminal courts, while the gacaca courts are charged with judging the two remaining categories. However, if a third category offender and the victim have agreed on an amicable settlement, the offender is no longer prosecuted by the gacaca court. A person cannot be classified in several categories at the same time, therefore if someone stole (Category 3) but also killed (Category 2), he is classified in the higher category (Category 2).

⁴Attendance was initially voluntary, but after problems with low attendance in the pilot phases, the law was revised, making attendance compulsory (Longman (2009)).

⁵The exact legal definitions of the suspect categories can be found in appendix.

There are no legal definitions for genocide survivors. According to a former prosecutor involved in Gacaca, the survivors recorded by the Gacaca information round are Tutsi who were living in the sector at the time of the genocide and survived but can also include Hutu widows (or widowers) who were married to Tutsi⁶. This view corresponds to the perception of Molenaar (2005) in his in-depth study of the gacaca process.

2.2 Reliability

From the gacaca records, it can be calculated that the nationwide total of category 1 and 2 suspects is close to 510,000. Given that on average 20% of suspects are acquitted, this would mean that category 1 and 2 count approximately 400,000 genocide perpetrators. But, adding about 100,000 perpetrators who passed away by 2005, their number increases again to about half a million (Government of Rwanda (2005)). This implies an active participation in the genocide of almost 20% of the adult Hutu population in 1994, or 40% of the adult male Hutu population⁷.

Is this a plausible figure? Compared to the work of Straus (2004), who puts forward an estimate of 175,000 to 210,000 perpetrators, this is at the high end. Straus (2004) underpins his estimate with detailed fieldwork in five administrative communes (out of the 145) and in-depth interviews with prisoners. From his fieldwork and interviews he takes a best estimate of 30-35 perpetrators per administrative cell over the course of the genocide and multiplies this with the number of cells in Rwanda in which genocide took place (5,852). Despite the large effort undertaken in collecting first-hand data, it is difficult to assess the reliability of the estimate put forward by Straus (2004) mainly because of a large number of untestable assumptions underlying the estimate.

The number of genocide suspects emerging from the gacaca is also at the high end compared to the number of detainees and accused persons not detained. In 2000, the government held 109,499 detainees on genocide charges, while the number of accused persons not detained was 49,066 (Office of the Prosecutor (2002)).

According to critics of the gacaca courts, at least three reasons may have caused over-reporting of the accused. First, late Human Rights Watch adviser Alison Des Forges argued that the concession programme, which requires the naming of all those who participated along with the accused in return for a lighter sentence, led to a multiplication of names. Second,

⁶Personal correspondence. The name of the prosecutor is withheld for confidentiality reasons.

⁷According to the 1991 census, Rwanda had 2,813,232 citizens aged 18 to 54, of which approximately 2,530,000 Hutu. Based on an annual average growth of 3%, the Hutu population in 1994 would have been close to 2,750,000 (Straus (2004), Verpoorten (2005)).

Longman (2009) claims that, over time, gacaca was undermined by government manipulation, aiming at a conviction of the largest possible number of Hutu in order to exclude much of the Hutu from holding public office. Third, several sources, including the Rwandan government, acknowledge that gacaca became a means of taking personal revenge on enemies, which contributed to the steep rise of the number of accused as gacaca proceeded. On the other hand, most sources evaluating gacaca also acknowledge that individuals may have escaped from accusation due to intimidation of witnesses, including murder or attempted murder of potential gacaca witnesses.

On the other hand, most sources evaluating gacaca also acknowledge that individuals may have escaped accusation due to intimidation of witnesses, including murder or attempted murder of potential gacaca witnesses. Moreover, the fifth report of the PRI research team (Penal Reform International) in Rwanda on the operation of gacaca makes mention of "very little participation of the population...Above all in towns or around churches where many people were killed...making it difficult to identify the culprits and to know exactly what happened". They find that "it is mainly the survivors who testify, while non-survivors hesitate and when they speak of killings, these usually concern the ones carried out by people who are either in prison or dead, or who have disappeared or fled" (Penal Reform International (2003)). Hence, it is likely that under-reporting of genocide suspects is especially problematic in those areas with few survivors.

Compared to the number of genocide suspects, the number of survivors that was recorded by the gacaca is expected to be more reliable since there are no clear motives for over- or under-reporting. However, the fact that a number of Hutu may also be counted among the survivors will have implications for the calculation of the genocide's death toll, a point on which we elaborate below.

The sum of male and female genocide survivors amounts to approximately 202,000. This is higher than the estimate of 150,000 survivors, based on counting in refugee camps immediately after the genocide (Prunier (1998)). In contrast, it is far lower than the reported 335,718 survivors in the census of survivors executed by the Rwandan government in 2006 (Government of Rwanda (2008)). However, apart from Tutsi living in the country at the time of the genocide, this census also includes Tutsi who escaped ethnic violence in neighboring countries, in particular Congo, as well as Tutsi who came back from living in exile abroad, especially Uganda.

This assessment of the quality of the gacaca data remains tentative, because the alternative data sources referred to are not flawless and comparison with the gacaca data is blurred

because different definitions are applied for identifying survivors and suspects. In any case, the above discussed reasons for over- and under-reporting of accused urge for a cautious interpretation of the gacaca data on genocide suspects.

3 Measuring Genocide

3.1 What are we (not) measuring?

The purpose of this section is to derive from the gacaca records a menu of meaningful genocide proxies, i.e. measures that capture (aspects of) genocide. The UN has recognized that genocide took place in Rwanda in the months April-June 1994, and several scholarly articles have estimated the death toll among Tutsi around 800,000, or approximately 75% of Rwanda's Tutsi population (e.g. Prunier (1998), Verpoorten (2005)).

Concurrently with the execution of genocide, other forms of violence took place in the same time span, including politicide (the killing of moderate Hutu by the genocidal regime), civil war between the RPF (Rwandan Patriotic Front) and the Habyarimana regime and reprisal killings on the part of the RPF directed against Hutu (Davenport and Stam (2009), Reyntjens (2009), the "Gersony report"⁸). Regarding the two latter forms of violence, Davenport and Stam (2009) estimate that, during April-June 1994, the sum of victims in zones under RPF control and the zones contested by the RPF and FAR+ amounts to respectively 80,000 and 90,000.

Also prior to and after 1994, different forms of violence took place in Rwanda, including the 1990-1992 civil war in the north of Rwanda and the 1995-1998 (counter-)insurgency in the northwest. Measuring the intensity of these forms of violence is out of the scope of this article, but is the topic of related work (Verpoorten (2010)).

3.2 A menu of measures

I define seven *genocide proxies* at the level of sector $i : 1...1484$. The first three genocide proxies measure *genocide participation*:

$$G1_i : \frac{(category_1_suspects_{i2005})(1 - d_{ns})^{-11}}{population_{i1994}},$$

⁸The "Gersony Report" is the name given to an unpublished report that identified a pattern of massacres by the RPF. The findings in the report were made by a team under Robert Gersony under contract to the United Nations High Commissioner for Refugees. Gersony's personal conclusion was that between April and August 1994, the RPF had killed "between 25,000 and 45,000 persons, between 5,000 and 10,000 persons each month from April through July and 5,000 for the month of August" (Des Forges (1999)).

$$G2_i : \frac{(category_2_suspects_{i2005})(1 - d_{ns})^{-11}}{population_{i1994}},$$

$$G3_i : \frac{(category_3_suspects_{i2005})(1 - d_{ns})^{-11}}{population_{i1994}},$$

with $category_1_suspects_{i2005}$, $category_2_suspects_{i2005}$ and $category_3_suspects_{i2005}$ as defined above, d_{ns} the annual natural death rate among genocide suspects between 1994-2005⁹, estimated at 10.9 per 1000, and $population_{i1994}$ the estimated 1994 sector level (pre-genocide) population projected forward from 1991 sector level population numbers on the basis of the 1978-1991 commune level population growth rate.

The next three genocide proxies measure *genocide survivorship* in the 1994 population defined as the proportion of surviving genocide victims, i.e. individuals who survived the genocide but are likely to have lost close relatives, be it widowed, orphaned or disabled genocide survivors:

$$G4_i : \frac{(widowed_survivors_{i2005})(1 - d_n)^{-11}}{population_{i1994}}$$

$$G5_i : \frac{(orphaned_survivors_{i2005})(1 - d_n)^{-11}}{population_{i1994}}$$

$$G6_i : \frac{(disabled_survivors_{i2005})(1 - d_n)^{-11}}{population_{i1994}}$$

with d_n a national level estimate of the natural death rate between 1994 and 2005 based on the 2000 Demographic and Health Survey - 9.2 per 1000 (Timaeus and Jasseh (2004)), and $population_{i1994}$ as defined above.

Finally, $G7$ measures *the genocide death toll* as the estimated number of Tutsi killed proportional to the 1994 Tutsi population:

$$G7_i = (1 - \frac{(genocide_survivors_{i2005})(1 - d_n)^{-11}}{Tutsi_population_{i1994}}) - d_n,$$

with $genocide_survivors_{i2005}$ the number of the male and female survivors reported in the 2005 gacaca records; and $Tutsi_population_{i1994}$ an estimate of the pre-genocide sector-level Tutsi population, based on the sector level total population and the commune level proportion of Tutsi as reported in the 1991 population census (Government of Rwanda, 1991; Minnesota Population Center, 2010). For details on the calculation of $G7$, I refer to the Appendix.

Which of these genocide proxies is to be used depends on the empirical question under

⁹The death rate is derived from information on the number of deceased genocide suspects between 1994 and 2005

study. It is rather straightforward that the genocide participation proxies, $G1$, $G2$ and $G3$, are useful for researchers studying the determinants of participation in the genocide, as is the case in Yanagizawa (2010) where slightly different definitions of $G1$ and $G2$ are used¹⁰. The genocide survivorship proxies, $G4$, $G5$ and $G6$ may be used to study the legacy of genocide, e.g. the impact of genocide on social capital in communities that are home to both surviving victims and perpetrators of the genocide. Finally, the genocide death toll, $G7$, may be used to analyze and explain the intensity of killings in the genocide.

In some cases, data analysis may benefit from aggregating a set of genocide proxies into one index or by transforming them into categorical variables. This is further explained and illustrated below.

3.3 Summary statistics

The first column in Table 1 gives the sector level mean for $G1 - G7$ across 1,390 sectors that could be matched with the population census. For the three measures of genocide participation ($G1$, $G2$ and $G3$) the means equal respectively 1.2%, 7.0% and 5.1%. For the different categories of genocide survivors ($G4$, $G5$ and $G6$), I obtain averages of respectively 0.5%, 1.2% and 0.2%, which are at the low hand, not only because entire families were killed during the genocide, but also because prior to the genocide Tutsi accounted only for approximately 12% of the 1994 population¹¹. The genocide death toll ($G7$) averages 69.0% (when censoring 69 anomalous values to zero and correcting for the under-reporting of Tutsi). When weighted by the share of Tutsi in the population of each of the sectors, we find a nationwide death toll of approximately 75.5%, which is very close to previous estimates by Prunier (1998) and Verpoorten (2005).

Table 2 shows the correlations across the different measures. The genocide participation and survivorship proxies exhibit a positive cross-correlation. This may be because both these categories of proxies increase with the pre-genocide Tutsi population share. On the other hand, it is a bit puzzling, since one would think that in areas with many perpetrators, there would be few survivors. Hence, the positive cross-correlation may suggest under-reporting of suspects in areas with few survivors, the more so because $GP1 - GP3$ correlate negatively

¹⁰In particular, Yanagizawa (2010) uses $\frac{category_1_suspects_{i2005}}{population_{i1991}}$, $\frac{category_2_suspects_{i2005}}{population_{i1991}}$.

¹¹The last population census prior to the genocide, conducted in 1991, reported 596,400 Tutsi living in Rwanda, representing 8.8% of the population. However, this figure is not reliable. Verpoorten (2005) provides evidence indicating that there was up to 40% under-reporting of Tutsi, either by the Habyarimana regime (in order to keep the school and public employment quotas of Tutsi low), either by Tutsi themselves (in order to avoid discrimination). Both in Verpoorten (2005) and Prunier (1998) it is argued that the proportion of Tutsi in 1991 was likely around 12%.

with the death toll $GP7$, which somewhat counter-intuitively suggests that the genocide death toll was lower in places with many perpetrators.

Table 3 gives the province level averages of the genocide proxies. To facilitate spotting provinces with high genocide intensity, for each of the proxies, the four highest values are put in bold. Butare stands out with top-4 values for all genocide proxies. Gikongoro, Kibuye, Gitarama and Cyangugu follow with each three or more genocide measures in the top-4. Kibungo and Kigali City feature twice in the top-4. The northern provinces Byumba, Gisenyi and Ruhengeri have the lowest values across the measures $G1 - G6$, which is in line with the historically low concentration of Tutsi in the north of the country (see Figure 1 for an administrative map of Rwanda).

4 Outlying and anomalous values (G1-G7)

The standard boxplots of the seven genocide intensity proxies are given in Figure 2. The genocide proxies $G1 - G6$ have a highly right-skewed distribution. This is in line with the fact that genocide intensity was very unequally distributed across sectors, mainly because the proportion of Tutsi across sectors in Rwanda was very uneven, but also because support for the genocide from the local administration and civilians varied across communes and provinces (Des Forges (1999)). In addition, from the boxplots, we detect a non-negligible number of outlying observations, and these outliers, whether stemming from real rare events, incidental or systematic error, amplify the skewness of the distribution.

It has been demonstrated that a high number of outlying observations can results in misleading statistics derived from the data, e.g. the sample mean and variance, making commonly used techniques such as OLS regression analysis and classical Principal Components Analysis (PCA) very sensitive to the presence of outliers (Barnett and Lewis (1993)). Detection of outliers as well as the use of outlier-robust techniques are often required to double-check results. To avoid arbitrariness in labelling extreme values as outliers, I turn to a procedure of outlier detection for skewed data (Hubert and der Veeken (2008)). For normal distributions, standard boxplot like those presented in Figure 2 can be used for detecting outliers. The whiskers of a standard boxplot are given by

$$[Q1 - 1.5IQR, Q3 + 1.5IQR],$$

with $Q1$ the first quartile, $Q3$ the third quartile and IQR the interquartile range for a

univariate continuous variable $X_n = \{x_1, x_2, \dots, x_n\}$. When the original variables are skewed, too many points tend to be flagged as outlying according to the standard boxplot whiskers. In order to identify outliers in skewed data, it is more appropriate to adjust the whiskers to

$$[Q1 - 1.5e^{-4MC} IQR, Q3 + 1.5e^{3MC} IQR],$$

with MC the medcouple defined as:

$$MC(X_n) = med_{xi < medn < xj} h(x_i, x_j),$$

$medn$ the sample median, and

$$h(x_i, x_j) = \frac{(x_j - medn) - (medn - x_i)}{x_j - x_i}$$

Using these definitions, I derive the skewness-adjusted whiskers of the genocide proxies $G1 - G6$. The values exceeding these whiskers are identified as outliers. On average, I find 9.6 outliers per genocide proxy. In total, 42 sectors have outlying observations for one or more of the six genocide proxies. When repeating Table 1 excluding these outliers, I find that the mean of the genocide proxies changes only marginally, whereas the standard deviations decrease considerable (not reported).

$G7$ includes 96 anomalous values for which the genocide's death toll is negative. The variable $G7$ is subject to several sources of measurement error. In the appendix, I give a detailed discussion of the possible causes and consequences of measurement error. In summary, the sources of error include the following: (i) d_n is a national level estimate of the natural death rate and may not be appropriate if there is large sub-national variation in post-genocide death rates; (ii) *genocide_survivors*₂₀₀₅ may include Hutu widows who were married to Tutsi; and (iii) the accuracy of *Tutsi_population*₁₉₉₄ hinges on the reliability of the 1991 population census and the extent of unobserved within commune variation in the population growth and the share of Tutsi in the population. These sources of error probably account for the anomalous cases. In the appendix, I discuss how these cases can be dealt with in an empirical application.

5 Construction of indices by PCA

In some cases, data analysis may benefit from aggregating a number of genocide proxies into one index¹². Several studies have persuasively argued for the use of principal component analysis as an aggregation method (PCA) (e.g. Filmer and Pritchett (2001)). PCA has the desirable property of reducing the dimensionality of a data set while retaining maximum variation in the data set. More precisely, from a set of variables, PCA extracts orthogonal linear combinations that capture the common information in the set most successfully. The first principal component (PC) identifies the linear combination of the variables with maximum variance, the second principal component yields a second linear combination of the variables, orthogonal to the first, with maximal remaining variance, and so on¹³. For our objective, i.e. defining an index of genocide intensity, we are interested in the first PC, which will be an appropriate summary of genocide intensity if it captures a relatively high percentage of the total variance present in the genocide proxies set and the "loadings" of that PC have roughly equal values.

A number of studies have used PCA for the purpose of summarizing conflict indicators by a conflict index. Pioneering work by Hibbs (1973) derives indices of "collective protest" and "internal war" from a 108-nation cross-sectional analysis of six event variables on mass political violence. Following Hibbs (1973) a large number of cross-country studies have used an index of sociopolitical instability as an explanatory variable in regressions in which the dependent variable is growth, savings or investment (e.g. Venieris and Gupta (1986), Barro (1991), Alesina and Perotti (1996)). To the best of our knowledge, only one micro-economic

¹²The rationale for aggregation is twofold. First, there may be complementarities between the measures, e.g. social capital (trust) may be more affected in a sector with a large share of survivors as well as a large number of suspects, than in a sector in which only one of those two groups is well represented. Second, assuming that the gacaca data include measurement error, this error can be attenuated by combining information across measures to reduce the effect of measurement error and outliers in each of the proxies separately.

¹³Formally, suppose that x is a vector of p random variables and x^* is a vector of the standardized p variables, having zero mean and unit variance, then the first principal component $PC1$ is the linear function $\alpha_1' x^*$ having maximum variance, where α_1 is a vector of p constants $\alpha_{11}, \alpha_{12}, \dots, \alpha_{1p}$ and $'$ denotes transpose.

$$PC1 = \alpha_1' x^* = \alpha_{11}x_1^* + \alpha_{12}x_2^* + \dots + \alpha_{1p}x_p^*,$$

Mathematically, the vector α_1 maximizes $var[\alpha_1' x^*] = \alpha_1' \Sigma \alpha_1$, with Σ the covariance matrix of x^* , which corresponds to the correlation matrix of the vector x of the original, unstandardized variables. For the purpose of finding a closed form solution for this maximization problem, a normalization constraint, $\alpha_1' \alpha_1 = 1$, is imposed. To maximize $\alpha_1' \Sigma \alpha_1$ subject to $\alpha_1' \alpha_1 = 1$, the standard approach is to use the technique of Lagrange multipliers. It can be shown that this maximization problem leads to choosing α_1 as the eigenvector of Σ corresponding to the largest eigenvalue of Σ , λ_1 and $var[\alpha_1' x^*] = \alpha_1' \Sigma \alpha_1 = \lambda_1$.

To interpret the PC in terms of the original variables, each coefficient α_{1l} must be divided by the standard deviation, s_l , of the corresponding variable x_l . For example, a one unit increase in x_l , leads to a change in the 1st PC equal to α_{1l}/s_l .

For a detailed exposition of principal component analysis we refer to Jolliffe (2002) and Duntzman (2001).

study, González and Lopez (2007), uses PCA to summarize variables into a micro level index of violent conflict. This study looks at the effect of political violence in Columbia on farm household efficiency. Five indicators of violence are defined: homicides, the number of attacks by FARC guerrillas, the number of attacks by ELN guerrillas, kidnappings, and displaced population. The first PC accounts for 43% of the joint variance of the five indicators, and is retained as an index of political violence.

As an illustration I subject the first six genocide proxies to PCA. The first principal components corresponds to the following linear combination:

$$GI = 0.40 \times G1 + 0.45 \times G2 + 0.40 \times G3 + 0.45 \times G4 + 0.43 \times G5 + 0.28 \times G6,$$

which explains 58% of the total variation in the underlying set of variables.

PCA relies on maximizing the classical sample variance. Therefore, it is sensitive to outliers. Since the data is highly skewed data, I also subject the set of variables to a recently proposed PCA that is robust to outliers in skewed distributions (Hubert, Rousseeuw, and Verdonck (2009)), referred to as ROBPCA. ROBPCA reduces the effect of outliers by replacing the classical sample covariance matrix used in classical PCA with a robust covariance matrix that is calculated for a subset of data points for which outlyingness is below a predefined threshold value. ROBPCA for skewed data uses the skewness-adjusted whiskers as a benchmark for defining outlyingness (see above). I find that PCA and ROBPCA yield roughly the same result: the correlation coefficient between the first principal components is 0.97.

Finally, I make a correction for possible survival bias. GI may be biased downward in communes where many families were entirely exterminated. In order to attenuate the effect of survival bias, I increase the weight of communes that are close to sites of large-scale massacres. The proximity to a large-scale massacre is taken into account by adding the natural logarithm of the commune level distance to the nearest mass grave to the set of variables subjected to PCA. This distance is calculated in km by overlaying a geo-referenced administrative map with the location of 71 mass graves in Rwanda taken from the Yale Genocide Studies website. The resulting GI is given by the following linear combination:

$$GI' = 0.39 \times G1 + 0.44 \times G2 + 0.38 \times G3 + 0.44 \times G4 + 0.42 \times G5 + 0.27 \times G6 - 0.27 \times G7,$$

with mg_j "log(distance to mass grave)". The correlation coefficient between GI and GI' is as large as 0.99.

Figure 3 displays the quintiles of the GI' on an administrative map, showing rather large within-province variation in GI' , with a large number of top quintile sectors in Butare, the eastern part of Gikongoro province and Kibuye province, as well as in the northwestern corner of Kibungo. In addition, smaller local clusters can be spotted in and around Kigali City and in the western province Cyangugu.

6 The construction of categorical variables using LISA

Categorical measures may be preferred to continuous variables for some purposes, e.g. summary statistics across conflict intensity or interaction effects in a regression analysis. Defining the categories based on percentiles (e.g. assigning "1" to the top 10% or top 20% values), or a cut-off value (e.g. assigning "1" to values of the standardized indices that exceed 0.5) involves a degree of arbitrariness. In addition, these methods run the risk of wrongly classifying erroneous outlier values. LISA avoids these caveats by identifying significant high-high spatial clusters, i.e. areas with high values of a variable that are surrounded by high values on the neighboring areas. Concomitantly, the low-low clusters are also identified from this analysis (Anselin (1995)).

More formally, LISA provides a measure of the extent to which the arrangement of values around a specific location deviates from spatial randomness. A general expression of a LISA statistic for a variable y_i , observed at location i , is:

$$L_i = f(y_i, y_{J_i}),$$

where f is a function expressing the correlation between y_i and y_{J_i} , and the y_{J_i} are the values observed in the neighborhood J_i of location i . The LISA statistic I look at is the local Moran statistic for an observation i :

$$I_i = (y_i - \bar{y}) \sum_{j=1}^n w_{ij} (y_j - \bar{y}),$$

with w_{ij} a spatial weighting matrix indicating the relevant neighbors for the LISA analysis. The weighting matrix w_{ij} can be defined in different ways, although contiguity-based definitions are by far mostly used. I use a first order rook-contiguity based weighting matrix for neighbors, where w_{ij} equals 1 for sectors with a common boundary.

By looking explicitly at areas instead of individual sectors, one can to a large extent avoid wrong classification of erroneous outliers. Arbitrariness in identifying "high" is avoided by assessing the significance of high-high clusters. The procedure employed to assess statistical significance relies on a Monte Carlo simulation of different arrangements of the data and the construction of an empirical distribution of simulated statistics. Afterwards the value obtained originally is compared to the distribution of simulated values and, if the value exceeds the 95th percentile, it is said that the relation found is significant at 5%.

LISA has been used in Anselin (1995) for analyzing spatial patterns of conflict in Africa. In addition, a number of micro-level studies have used LISA for detecting hot spots in crime (e.g. Murray, McGuffog, Western, and Mullins (2001)). Several other cluster detection methods have been proposed and used for analyzing the location of armed conflict across countries (e.g. Ward and Gleditsch (2002)). A recent micro-level application uses the SaTScan program for detecting space-time clusters in DR Congo (Raleigh, Witmer, and Loughlin (2009)).

By means of illustration, Figure 4 shows the locations with significant high-high (dark grey) and significant low-low clusters (light grey) for GI' . Note the very large low-low cluster in the North, corresponding to low shares of Tutsi in the northern provinces. The significant high-high clusters confirm the pattern detected before: Butare clearly stands out with almost half of its territory belonging to a high-high cluster, while Kibuye comes in second with several high-high clusters on a relatively small area; Gikongoro, Kibungo and Gitarama follow closely. Finally, a few small high-high clusters turn up in Rural Kigali and Cyangugu.

7 Conclusion

This article describes the data released by the Rwandan transitional justice system charged with judging genocide suspects. After discussing the general reliability of the data, I presented a menu of genocide proxies that capture genocide participation, survivorship of close relatives of genocide victims, and the genocide's death toll. The summary statistics of the different measures across provinces correspond to the intensity of genocide described in event data.

Subjecting a number of these proxies to PCA yields an index of genocide intensity, which can be corrected for survival bias, using the distance to the nearest massgrave. Finally, I used Local Indicators of Spatial Association (LISA) to transform the continuous indices into categorical variables in a non-arbitrary way that is robust to spatial outliers.

The gacaca data can be matched with several existing nationally representative Rwandan household surveys in which sectors are used as sample units. Hence, the scope for using the proposed genocide intensity measures in empirical applications is large. However, given the issues of reliability - especially surrounding the numbers of suspects - the proposed measures have to be used with caution.

Appendix

The legal definition of the genocide suspect categories

Originally, four categories of genocide suspects were created in 1996 by the Act on the Organization and Pursuits of Crimes against Humanity. However, the Organic Law 16/2004 of 19.06.2004 reduces the categories to three: the former categories 2 and 3 were combined into category 2 and the 4th category became the 3rd one.

- Category 1:

- (a) The person whose criminal acts or criminal participation place him or her among the planners, organizers, incitators, supervisors and ringleaders of the genocide or crimes against humanity, together with his or her accomplices;
- (b) The person who, at that time, was in the organs of leadership, at the national level, at the level of Prefecture, Sub-prefecture, Commune, in political parties, army, gendarmerie, communal police, religious denominations or in militia, has committed these offences or encouraged other people to commit them, together with his or her accomplices;
- (c) The well known murderer who distinguished himself or herself in the location where he or she lived or wherever he or she passed, because of the zeal which characterized him or her in killings or excessive wickedness with which they were carried out, together with his or her accomplices;
- (d) The person who committed acts of torture against others, even though they did not result into death, together with his or her accomplices;
- (e) The person who committed acts of rape or acts of torture against sexual organs, together with his or her accomplices;
- (f) The person who committed dehumanizing acts on the dead body, together with his or her accomplices.

- Category 2:

- (a) The person whose criminal acts or criminal participation place him or her among the killers or who committed acts of serious attacks against others, causing death, together with his or her accomplices;

- (b) The person who injured or committed other acts of serious attacks with the intention to kill, but who did not attain his or her objective, together with his or her accomplices;
- (c) The person who committed or aided to commit other offences against persons, without the intention to kill, together with his or her accomplices.

- Category 3:

- (a) The person who only committed offences against property.

The calculation of the genocide's death toll

The sector-level number of Tutsi killed proportional to the 1994 Tutsi population can be calculated as follows:

$$genocide_toll_{1994,i} = (1 - \frac{(genocide_survivors_{i2005})(1 - d_n)^{-11}}{Tutsi_population_{i1994}}) - d_n, \quad (1)$$

with $genocide_survivors_{i2005}$ the sector-level number of male and female survivors, d_n the annual post-genocide natural death rate, and $Tutsi_population_{i1994}$ the 1994 pre-genocide sector-level Tutsi population.

In what follows, I explain each of the three components at the right hand side of this equation in detail and discuss the possible causes and consequences of measurement error.

The number of genocide survivors

As mentioned in the article, the survivors recorded by the Gacaca information round are Tutsi who were living in the sector at the time of the genocide and survived but can also include Hutu widows (or widowers) who were married to Tutsi. If, besides Tutsi, a number of Hutu are counted among the survivors, this has implications for the estimation of the genocide's death toll. More precisely, $genocide_toll_{1994,i}$ will be biased downward if Hutu relatives are included as survivors, which is likely to occur more in localities in which Tutsi were well integrated. In the empirical application of the article, the possible bias stemming from this form of measurement error is attenuated by including controls for the share of Tutsi in the population and the proportion of inter-ethnic marriages

The pre-genocide Tutsi population

$Tutsi_population_{i1994}$ is an estimate of the 1994 sector-level Tutsi population, which is calculated on the basis of 1991 sector-level total population (pop_{i1991}), the 1978-1991 commune

level annual population growth rate ($pop_growth_{c1978-1991}$) and the 1991 commune-level proportion of Tutsi ($share_Tutsi_{c1991}$)¹⁴. More precisely:

$$Tutsi_population_{i1994} = (pop_{i1991} * (pop_growth_{c1978-1991})^3) * share_Tutsi_{c1991} \quad (2)$$

Within commune variation in $pop_growth_{c1978-1991}$ and $share_Tutsi_{c1991}$ will cause measurement error in $Tutsi_population_{i1994}$. The accuracy of $Tutsi_population_{i1994}$ also depends on the appropriateness of projecting pop_{i1991} forward to 1994 using information on population growth between 1978-1991. Finally, the reliability of $Tutsi_population_{i1994}$ hinges on the reliability of the population census data. Verpoorten (2005) provides evidence indicating that the 1991 population census is highly reliable for total sector-level population numbers, but not for ethnicity-specific numbers. In particular, using data for Gikongoro Province, it is shown that there was on average 40% under-reporting of Tutsi, either by the Habyarimana regime, or by Tutsi themselves¹⁵. In order to obtain a more accurate figure for the overall death toll, I repeat the calculation of $genocide_toll_{1994,i}$ using an adjusted measure of $share_Tutsi_{c1991}$, i.e. $share_Tutsi_{c1991} * 1.4$. However, sub-national variation in the extent of under-reporting may still cause bias. For example, in areas with relatively large under-reporting of Tutsi, the calculated death toll will be biased downward.

The natural death rate between 1994 and 2005

There are no estimates of post-genocide natural death rates among genocide survivors. Therefore, I proxy for d_n using information on adult mortality rates from the general population estimated from the 2000 Rwandan DHS survey (Timaeus & Jasseh, 2004). The estimates equal 8.1 and 11.8 for women and men aged 15 to 60, respectively. Taking a weighted average with the weight for women reflecting their approximate share among the survivors recorded in the Gacaca (2/3), I obtain a proxy of 9.2 for d_n .

This proxy is at the national level. Hence, $genocide_toll_{1994,i}$ will be biased upwards (downwards) in localities with a relatively high (low) post-genocide natural death rate. This bias is partly controlled for in the empirical application by including commune-level infant mortality as a control variable.

¹⁴ $Population_{i1991}$ and $population_growth_{c1978-1991}$ were obtained from the 1991 and 1978 population census data provided by the Rwandan National Census Service, while $share_Tutsi_{c1991}$ was downloaded from IPUMS International website (Minnesota Population Center).

¹⁵ An undetermined number of Tutsi registered as Hutu in order to avoid discrimination. In addition, the Habyarimana regime is said to have deliberately under-reported the number of Tutsi in order to keep their school enrolment and public employment quotas low.

Anomalous values and implications for empirical applications

The genocide's death toll $genocide_toll_{1994,i}$ includes a number of anomalous values. In particular, it is negative for 96 administrative units, which is likely to be due to one or more of the measurement errors discussed above.

Hence, when used in empirical applications, it is useful to perform a series of robustness checks in order to verify if measurement error interferes with the identification of the research question. For example, one could verify if the distribution of anomalous values is unrelated to the explanatory variables of interest. In addition, the analysis can be performed both including and excluding the observations for which $genocide_toll_{1994,i}$ is negative. Finally, several control variables may act to attenuate the impact of possible measurement error in the remaining observations, i.e. the commune level share of Tutsi in the population, the 1991 commune level proportion of inter-ethnic marriages, and the commune level infant mortality rate (see Verpoorten (2011) for an example).

References

- ALESINA, A., AND R. PEROTTI (1996): “Income distribution, political instability, and investment,” European Economic Review, 40(6), 1203–1228.
- ANSELIN, L. (1995): “Local indicators of spatial association,” Geographical Analysis, 27, 93–115.
- BARNETT, V., AND T. LEWIS (1993): Outliers in Statistical Data (3rd edn). John Wiley & Sons Ltd.: Chichester, U.K.
- BARRO, R. J. (1991): “Economic growth in a cross-section of countries,” Quarterly Journal of Economics, 106, 407–444.
- DAVENPORT, C., AND A. STAM (2009): “Rwandan Political Violence in Space and Time,” Unpublished manuscript.
- DES FORGES, A. (1999): Leave None to Tell the Story: Genocide in Rwanda. New York: Human Rights Watch.
- DUNTEMAN, G. H. (2001): Principal Components Analysis, Sage university papers, Quantitative applications in the social sciences. Sage.
- FILMER, D., AND L. H. PRITCHETT (2001): “Estimating wealth effects without expenditure data—or tears: an application to educational enrollments in states of India,” Demography, 38(1), 115 – 132.
- GONZÁLEZ, M. A., AND R. A. LOPEZ (2007): “Political Violence and Farm Household Efficiency in Colombia,” Economic Development and Cultural Change, 55(2), 367 – 392.
- GOVERNMENT OF RWANDA, . (2005): “Report on Data Collection: Annexes,” Discussion paper, National service of Gacaca Jurisdiction.
- (2008): “Recensement des Rescapés du Génocide de 1994: rapport final,” Discussion paper, Service National de Recensement.
- HIBBS, D. (1973): Mass Political Violence: A Cross-national Causal Analysis. John Wiley & Sons.

- HUBERT, M., AND S. V. DER VEEKEN (2008): “Outlier detection for skewed data,” Journal of Chemometrics, 22, 235–246.
- HUBERT, M., P. ROUSSEEUW, AND T. VERDONCK (2009): “Robust PCA for skewed data and its outlier map,” Computational Statistics and Data Analysis, 53, 2264–2274.
- INGELAERE, B. (2009): “Does the truth pass across the fire without burning? Locating the short circuit in Rwanda’s Gacaca courts,” Journal of modern African studies, 47(4), 50–528.
- JOLLIFFE, I. (2002): Principal Component Analysis, no. XXIX in Springer Series in Statistics. Springer, NY, 2nd edn.
- LONGMAN, T. (2009): “An Assessment of Rwanda’s Gacaca Courts,” Peace Review, 21(3), 304–312.
- MOLENAAR, A. (2005): “Gacaca: grassroots justice after genocide,” Research report 77/2005, African Studies Centre.
- MURRAY, A., I. MCGUFFOG, J. WESTERN, AND P. MULLINS (2001): “Exploratory Spatial Data Analysis Techniques for Examining Urban Crime,” The British Journal of Criminology, 41, 309–329.
- OFFICE OF THE PROSECUTOR, . (2002): “Abantu Bafungiyе mu Magereza Kasho na Buri-gade,” Discussion paper, Rwandan Ministry of Justice.
- PENAL REFORM INTERNATIONAL, . (2003): “Research on the Gacaca - Report V,” Discussion paper, PRI.
- PITSCH, A. (2002): “The gacaca law of Rwanda: possibilities and problems in adjudicating genocide suspects,” Discussion paper, Working Paper NUR-UMD Partnership.
- PRUNIER, G. (1998): The Rwanda Crisis: History of a Genocide. London: Hurst & Company.
- RALEIGH, C., A. LINKE, H. HEGRE, AND J. KARISEN (2010): “Introducing ACLED: An Armed Conflict Location and Event Dataset,” Journal of Peace Research, 47(5), 651–660.
- RALEIGH, C., F. WITMER, AND J. O. LOUGHLIN (2009): A Review and Assessment of Spatial Analysis and Conflictchap. in "The Geography of War in Geographic Contributions to International Studies" (ed. C. Flint). Oxford: Basil Blackwell.

- RESTREPO, J., M. SPAGAT, AND J. VARGAS (2006): “The Severity of the Colombian Conflict: Cross-Country Datasets Versus New Micro-Data,” Journal of Peace Research, 43(1), 99–115.
- REYNTJENS, F. (2009): The great African war: Congo and regional geopolitics, 1996-2006. New York, N.Y.: Cambridge University Press.
- STRAUS, S. (2004): “How many perpetrators were there in the Rwandan genocide? An estimate,” Journal of Genocide Research, 6(1), 85–98.
- TIMAEUS, I., AND M. JASSEH (2004): “Adult Mortality in Sub-Saharan Africa: Evidence from Demographic and Health Surveys,” Demography, 41, 757–72.
- VENIERIS, Y., AND D. GUPTA (1986): “Income distribution and socio-political instability as determinants of savings: A cross-sectional model,” Journal of Political Economy, 96, 873–883.
- VERPOORTEN, M. (2005): “The death toll of the Rwandan genocide: a detailed analysis for Gikongoro Province,” Population, 60(4), 331–368.
- (2010): “Detecting Hidden Violence: The Spatial Distribution of Excess Mortality in Rwanda,” Discussion Paper 254, LICOS - Centre for Institutions and Economic Performance, K.U.Leuven.
- (2011): “Leave none to claim the land. A Malthusian catastrophe in Rwanda?,” Discussion Paper 291, Licos Discussion Paper.
- WARD, M. D., AND K. S. GLEDITSCH (2002): “Location, Location, Location: An MCMC Approach to Modeling the Spatial Context of War and Peace,” Political Analysis, 10, 244–260.
- WOLTERS, S. (2005): “The Gacaca Process: Eradicating the culture of impunity in Rwanda,” Situation report, Institute for Security Studies.
- YANAGIZAWA, D. (2010): “Propaganda and Conflict: Theory and Evidence from the Rwandan Genocide,” Unpublished manuscript.

Figure 1: administrative map of Rwanda

Figure 2: boxplots of G1-G7, standardized to have mean 0 and st. dev. 1

Figure 3. Quintiles of G_I' (darkest grey = upper quintile)

Note: G_I' is the first principal component of the genocide proxies G_1 - G_6 and the logged distance to nearest mass grave; the checkered areas are areas with missing data, including the national park, forest areas and lakes

Figure 4. significant high-high (dark grey) and low-low clusters (light grey) of GI'

Note: GI' is the first principal component of the genocide proxies G1-G6 and the logged distance to nearest mass grave; the checkered areas are areas with missing data, including the national park, forest areas and lakes

Table 1. Measuring genocide at the sector level

(N=1390)	Mean	St. Dev.	Mean ^a
(a) Genocide participation			
(G1) Category 1 suspects (% 1994 population)	1.2%	1.7%	
(G2) Category 2 suspects (% 1994 population)	7.0%	7.3%	
(G3) Category 3 suspects (% 1994 population)	5.1%	6.5%	
(b) Genocide survivorship			
(G4) Widowed genocide survivors (% 1994 population)	0.5%	0.6%	
(G5) Orphaned genocide survivors (% 1994 population)	1.2%	1.7%	
(G6) Disabled genocide survivors (% 1994 population)	0.2%	0.5%	
(c) Genocide death toll			
(G7) Genocide death toll as a % of 1994 Tutsi population	55.5%	42.7%	66.0%
- excl. anomalous values	63.1%	27.7%	66.6%
- anomalous values censored to zero	58.7%	31.2%	66.6%
- corrected for under-reporting of Tutsi	68.0%	30.5%	75.5%
- corrected for under-reporting of Tutsi & excl. anomalous values	70.6%	23.0%	75.5%
- corrected for under-reporting of Tutsi & anomalous values censored to zero	69.0%	25.1%	75.5%

Notes: the sector level 1994 population is projected forward from the 1991 population census; details on the calculation of GP7 are given in the Appendix; ^aWeighted with the size of the pre-genocide sector-level Tutsi population

Table 2: Correlation matrix (N=1390)

	G1	G2	G3	G4	G5	G6	G7
(a) Genocide participation							
(G1) Category 1 suspects (% 1994 population)	1						
(G2) Category 2 suspects (% 1994 population)	0.612***	1					
(G3) Category 3 suspects (% 1994 population)	0.551***	0.740***	1				
(b) Genocide survivorship							
(G4) Widowed genocide survivors (% 1994 population)	0.503***	0.590***	0.436***	1			
(G5) Orphaned genocide survivors (% 1994 population)	0.468***	0.566***	0.422***	0.782***	1		
(G6) Disabled genocide survivors (% 1994 population)	0.285***	0.260***	0.189***	0.516***	0.388***	1	
(c) Genocide death toll							
(G7) Genocide death toll as a % of 1994 Tutsi population	-0.184***	-0.223***	-0.225***	-0.284***	-0.255***	-0.192***	1

Notes: G7 is corrected for under-reporting of Tutsi & anomalous values censored to zero

Table 3: Province level averages of G1 - G7

Province	G1	G2	G3	G4	G5	G6	G7
Butare	0.021	0.106	0.079	0.009	0.023	0.004	0.750
Byumba	0.002	0.016	0.018	0.001	0.003	0.001	0.608
Cyangugu	0.011	0.074	0.032	0.007	0.019	0.002	0.673
Gikongoro	0.015	0.085	0.088	0.006	0.014	0.002	0.772
Gisenyi	0.005	0.037	0.033	0.001	0.005	0.001	0.633
Gitarama	0.017	0.091	0.075	0.007	0.016	0.003	0.580
Kibungo	0.017	0.093	0.052	0.005	0.014	0.002	0.665
Kibuye	0.020	0.103	0.079	0.004	0.012	0.002	0.833
Rural Kigali	0.012	0.085	0.048	0.004	0.011	0.002	0.651
Kigali City	0.011	0.079	0.033	0.004	0.015	0.001	0.810
Ruhengeri	0.002	0.015	0.015	0.001	0.002	0.000	0.738
Umutara	0.007	0.029	0.021	0.002	0.006	0.001	0.712

Notes: G7, corrected for under-reporting of Tutsi & anomalous values censored to zero