The Demographic Transition, Social Network Change, and the Size of the West African Ebola Outbreak

Ashton Verdery, Nalyn Siripong, Jonathan Daw, James Moody

Abstract

Why is the current Ebola outbreak in West Africa orders of magnitude larger than any seen before? Here we examine the size and structure of family networks – the sorts of close personal relations needed to carry a disease transmissible only by close contact with bodily fluids. Using network simulation methods, we find that demographic changes experienced in West Africa have led to a rapid increase in numbers of living family members which now supports large, globally connected kinship networks. Prior Ebola outbreaks occurred in areas and times that corresponded to small fragmented networks. In this context, features of these countries' demographic transitions play a crucial role by generating a rapid growth in living family sizes that pushed community-level network structures past the connected component phase transition. This creates opportunities for a broad, fast moving outbreak that can spread far in the network solely through family contact.

Extended Abstract

Introduction

The first documented Ebola outbreaks occurred in the Democratic Republic of Congo (DRC) and Sudan in 1976, and, until this year, 26 other outbreaks had occurred in those countries and others in Central Africa including Gabon, Cote d'Ivoire, South Africa, Uganda, and Congo (WHO 2014). Most were small; only seven infected more than 100 people. Of these, four infected more than 300: the original outbreak in DRC (318 cases), another in DRC in 1995 (315 cases), and, in the largest outbreak prior to this year, in Uganda in 2004 (425 cases). In 2014, however, in addition to an unrelated outbreak of a different viral strain in DRC, five countries in West Africa - Guinea, Liberia, Sierra Leone, Nigeria, and Senegal – which have previously been unaffected began to experience the 2014 West African Ebola outbreak. As of September 14th, there had been 3,095 confirmed cases with an additional 2,262 suspected and probable cases in these countries which resulted in 2,630 deaths ("WHO | Situation Reports: Ebola Response Roadmap": 9/18/2014). The current outbreak is thus many orders of magnitude larger than any prior, and more have been infected and died from Ebola in the current West African outbreak than in all prior outbreaks combined. What accounts for this difference?

A combination of political and social factors have dealt a severe blow to the healthcare infrastructure of these countries – numerous press reports and policy briefs from global health organizations cite the lack of trained doctors and other healthcare workers as instrumental to the spread of the disease – but the underlying social structure necessary for a disease like Ebola to spread is a connected network of susceptible people linked to an infection source. Recent genomic surveillance indicates that the current outbreak in West Africa is likely to have come from a single human-animal point of contact, and been transmitted entirely between humans since then (Gire et al. 2014). While a substantial number of healthcare workers have contracted the disease (318 as of 9/14/2014) and several signs point to hospitals and funerals as key points of transmission in the current outbreak, there remains a substantial proportion of infections that are not accounted for by these sources. Since Ebola is only known to be transmissible from person to person through contact with blood or other bodily fluids, this network has to be intimate, which suggests families as an additional transmission substrate with potentially important consequences.

We examine trends in family size in the West African context by looking at changes in demographic rates governing reproduction and death. The historical trajectory of vital rates in these countries suggests that we can expect the average individual to have substantially more living family members at present than in the past. We use these insights to simulate the communal evolution of family networks in these countries, and examine the resulting network structure for epidemic carrying capacity.

Background

Nearly every contemporary society has experienced a “demographic transition” (Guest and Almgren 2003). Prior to the transition, birth and death rates are both high and in approximate equilibrium, but they go from high to low as the transition progresses. Improvements in infant mortality and health care associated with modernization are generally implemented more rapidly than social norms about family size, however, which means that death rates, especially those of infants and children, are lower than fertility rates during this stage. This combination leads to often substantial population growth and a younger shift in the age structure of the population, but it also facilitates dramatic increases in family size. When mortality declines prior to fertility it raises the average number of siblings in cohorts born during the transition because the shift from high to moderate mortality levels primarily affects the likelihood of infant and child survival (Coale 1972). Sets of siblings born after mortality has fallen but while fertility remains high will be larger than sets born in earlier (when more infants were likely to die) or later periods (when fewer children are born per couple). Finally, the transition is completed as education and modernization increase, birth rates tend to decline, and, after a considerable delay owing to the natural momentum embedded in the age structure which leads a large numbers of children to pass through the high fertility ages even after fertility rates have declined, population size stabilizes. There is evidence that the effects of the demographic transition on family size, however, can persist long after vital rates have ceased to change (Verdery 2010; Murphy 2011). Hammel (2005) shows that demographic shocks tend to affect family structure in small societies primarily by reducing the numbers of individuals without any living kin.

The processes of societal modernization have not only affected the demographic transition and family sizes, they have also tended to change living arrangements. In this regard, there are stark differences between the developed world – where modernization corresponded with large drops in amounts inter-generational co-residence – and the developing world – where little change has so far been evident except for slight drops in some African countries not currently experiencing Ebola (Ruggles and Heggenness 2008). Whether or not changes in co-residence have occurred, however, increases in numbers of living kin may still alter network structures because, in the family-oriented societies of the developing world, close kin tend to live near one another (Verdery et al. 2012) and have many opportunities for visitation.

Countries in Central and West Africa are in the growth stage of the demographic transition. Figure 1 shows population growth rates in the three countries currently experiencing substantial numbers of Ebola infections in the West African Ebola outbreak (Guinea, Liberia, and Sierra Leone) and five other countries which have had substantial outbreaks in the past (Congo, D.R. Congo, Gabon, Sudan and Uganda)[1]. There is a critical difference in growth rates between countries currently experiencing the West African Ebola outbreak and the other countries which have previously experienced Ebola. In the current outbreak countries, population growth rates surged above 4% for periods in the 1990s and 2000s. In Guinea, this included 5 years in the 1990s and 1 in the 2000s; Liberia experienced greater than 4% growth for 5 years in the 2000s and 1 year in the 2010s. Sierra Leone is a bit of an anomaly, as it only experienced 3 years of high growth in the 2000s. The only other country historically affected by Ebola with growth above 4% in this period is Sudan, which experienced four years of high growth in the early 1990s.

Figure 1. Trends in population growth in countries with large scale Ebola experience.

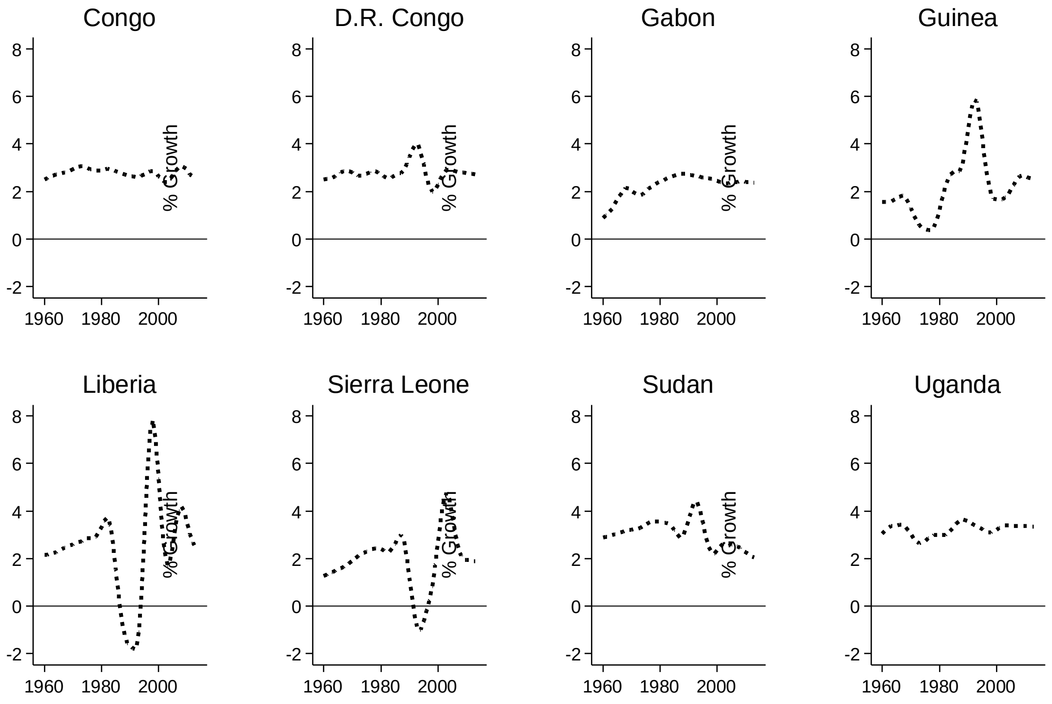

---

[1] Other countries which have experienced Ebola have all had very small outbreaks with fewer than 25 infections. These include Ivory Coast in 1994 (1 case), South Africa in 1996 (1 case), Senegal in 2014 (1 case), and Nigeria in 2014 (22 cases; 11 of whom were health workers).

We argue that these growth trends, a distinctive feature of the sub-Saharan demographic transition, may play a critical role in the size of the 2014 West African Ebola epidemic, because prior outbreaks are likely to have occurred in times and places where extant family sizes were comparatively smaller.

This argument draws not only on demographic transition theory, but also from mathematical theories of random graphs. Erdős and Rényi (1960) described the properties of graphs with *N* nodes and a *p* probability of an edge independently existing between any pair of nodes. As applied to social networks, nodes would be individuals and edges would be connections between them. Graphs and, by implication, social networks constructed in this way have remarkable properties with strong epidemiological relevance. At the most basic level, as *p* increases, the network will transition from almost surely having no large groups of individuals connected through chains of ties of arbitrary length, to a state where the vast majority of individuals can be reached through some chain of connection. More remarkably, however, is that this transition is highly non-linear and, indeed, there is a value of p (for any *N*) wherein the network almost immediately transitions from low levels of connectivity to high levels of connectivity. This shift is known in the literature as a phase transition. Figure 2 illustrates this for a network of *N*=100 and *N*=500 individuals. Its y-axis measures the proportion of the network in the largest connected component, that is, the proportion of the network reachable through chains of any length, where a chain from, e.g., A to C exists if A is tied to B and B is tied to C. The x-axis of Figure 2 measures values of *p* (independent probability of each edge existing) in steps of 0.0001. For each value of *N* and *p*, we simulated 10 Erdős-Rényi graphs, the connectivity measures for each of which are plotted. When the graph is large (*N*=1000), the network transitions rapidly from a state of low connectivity to high connectivity; by the time that each edge has a 1% probability of existing, the network is almost fully connected. In smaller graphs, the transition is slower (remember, smaller graphs have fewer potential edges), but still sharp. Between *p*=0.01 and *p*=0.02 the network shifts from almost completely disconnected to almost completely connected.

Figure 2. The connectivity phase transition in random graphs.

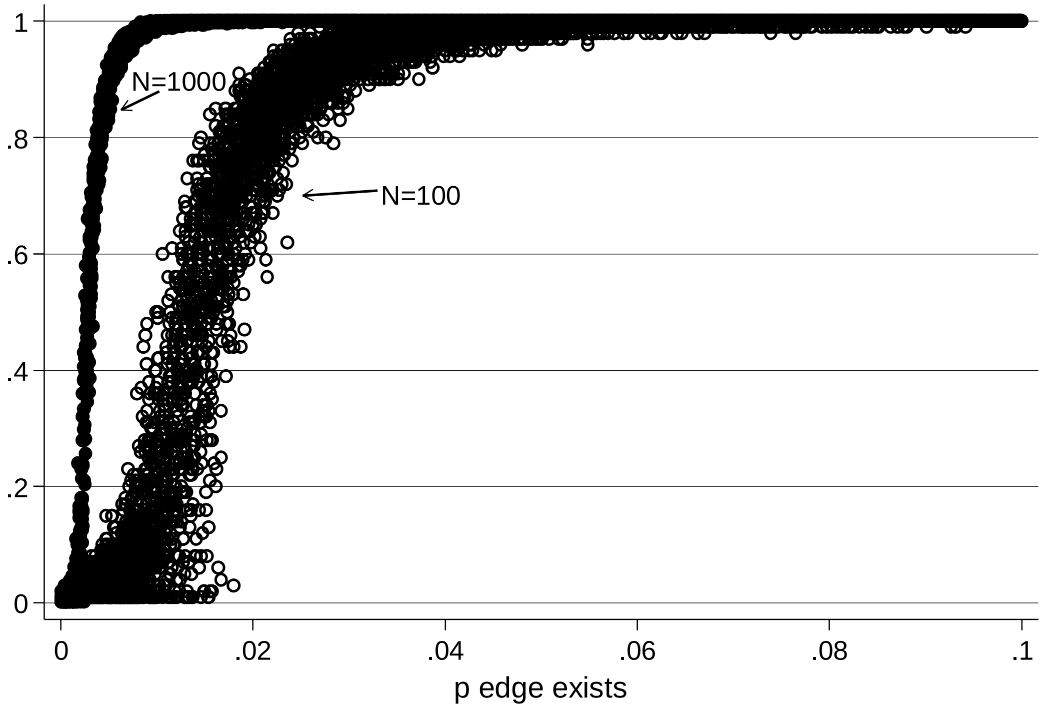

While it is known that random graphs exhibit phase transitions in global connectivity levels as the proportion of edges increase, whether such results generalize to human social networks – especially highly structured ones such as kinship networks – is unknown. The Erdős-Rényi model is built from a case where each edge is added to the network with independent probabilities, while a kinship network composed of parent-child ties, for example, would have structural prohibitions on parents being tied to their children's children. The potential for a connectivity phase transition in a kinship network is intriguing, however, and it may correspond to movement through the demographic transition. We explore this potential in the remainder of the paper.

Data and Methods

We use data from the Census's International Data base on total fertility rates in each of the countries which has experienced a large-scale Ebola outbreak. These are translated into age-specific fertility rates in each year using the ASFRPATT model (Arriaga 1992), which given a value of TFR generates a model age-specific fertility schedule. We translated lumped rates to single years of age with a fitted third-order polynomial. We also use data on each country's mortality rates as derived from life expectancies at birth. These come from the World Bank's World Development Indicators database, which covers the period 1960-2014, supplemented with additional data from Gapminder (Gapminder Documentation 2014) on early mortality, and knowledge of the mortality transition in each country offered in Riley (2005). To translate life expectancy at birth into age-specific mortality rates, we use the appropriate updated Coale-Demeny model life tables (found in UN 2012a) for each country as recommended by the United Nations (UN 2012b). Linear interpolation between the time-adjusted mid-points of age groups was used to obtain single years. We derive median age of marriage data for each country from Tabutin and Schoumaker (2004), and assume that in- and out-migration is constant at 1.5% of the village size per annum. We simulate village evolution from 1800-present.

Figure 3 shows the assumed parameters underlying the modeled fertility and mortality transitions in each country. As can be seen, there is substantial variation in the timing and tempo of both the fertility and mortality transitions. We examine what implications such variation has for kinship networks in each society, and whether it may help explain the size of the 2014 West African Ebola epidemic compared to the sizes of prior Ebola outbreaks.

Figure 3. Modeled fertility and mortality transitions in each country.

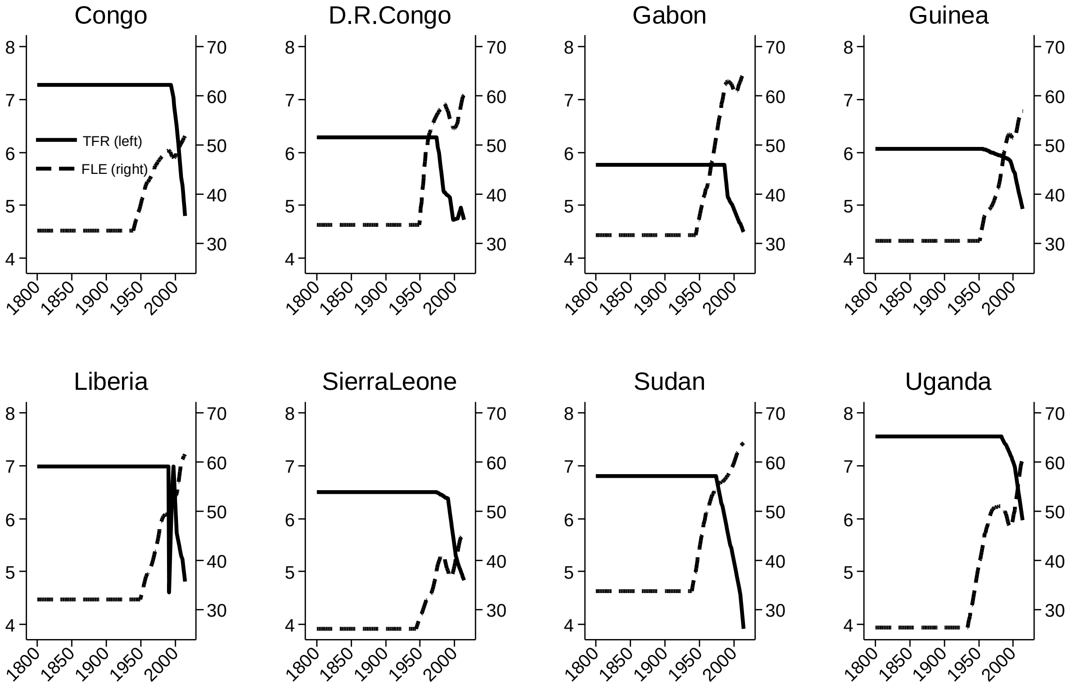

Figure 4 presents a simplified illustration of the evolution of within-village kinship networks captured by the demographic simulation. Four snapshots of the village at different points in time are presented. In year 0, the four individuals labeled A-D are unconnected recent in-migrants. Over the next two years, A and D both marry someone outside of the village while B and C marry each other; both types of marriage exist within the simulation. As the next 25 years pass, numerous other events occur: E in-migrates, D out-migrates, A and E bear two daughters (G and H), and B and C have a daughter (K) and two sons (I and J). In addition, B and C's son J died between years 2 and 25 and their daughter married L (who does not live in the village at the time). Finally, by year 60, L has in-migrated to join his spouse K with whom he has had a son, O. Other births have occurred in this interval as have some deaths. Over the period represented in this illustrative example, demographic processes have created complex kin structures. Our simulation model extends this sequence of events to capture the evolving network of kin relations that are conditional on such demographic processes. By modeling the historical evolution of vital rates across time in the countries experiencing the 2014 West African Ebola epidemic, we can simulate how a network of kin relations develops.

Figure 4. Schematic representation of the simulated evolution of kinship networks.

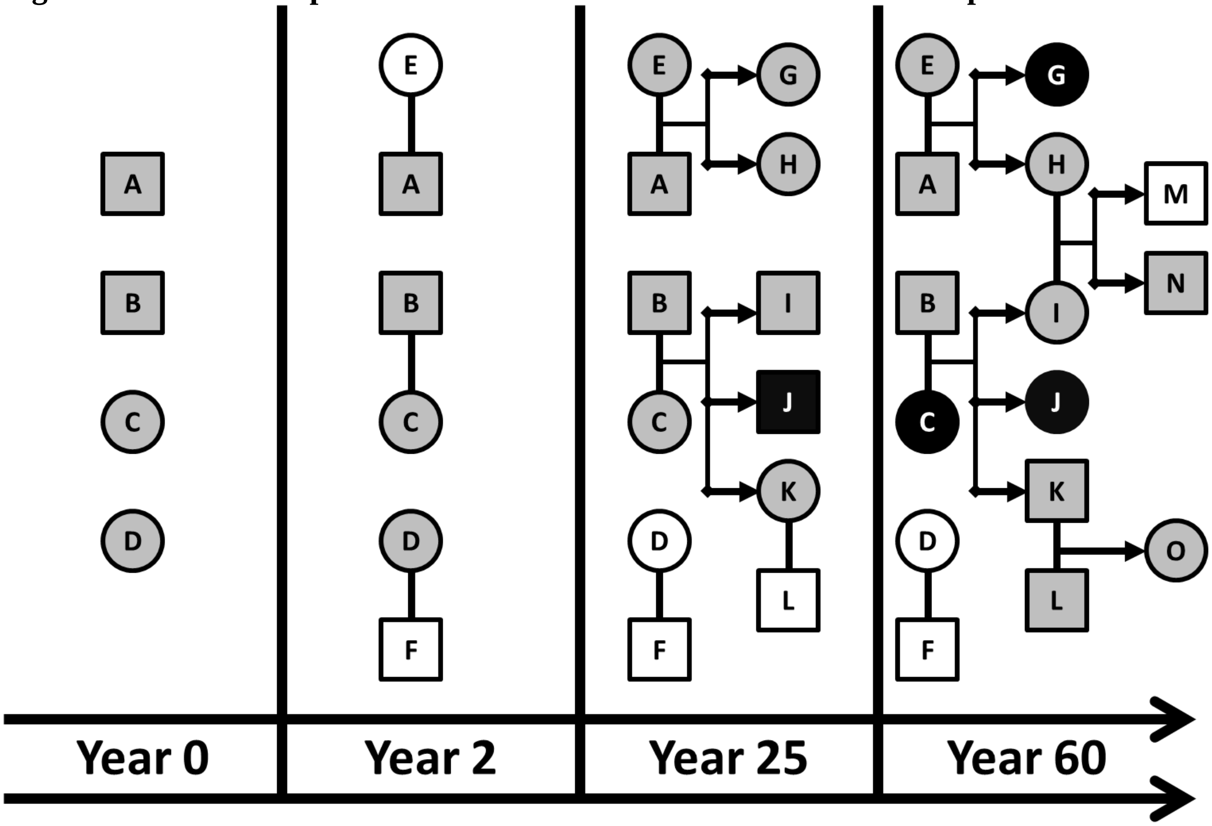

Notes: Time periods and generations within a period are represented as progressing from left to right. A letter denotes each person and can be used to track them across periods. Squares are males, circles are females. Gray nodes are living village residents, white nodes are those outside of the village (who may in-migrate [e.g., node E between years 2 and 25] or be out-migrants [e.g., node D in year 25]), and black nodes are dead individuals [e.g., node C in year 60]. Lines with arrows represent parent to child ties, lines without arrows represent marital ties.

After simulating kinship networks arising from these demographic transitions, we then simulate an Ebola epidemic on each network. To do this, we use a dynamic disease simulation tool. The tool is built and will be described in the full paper.

Preliminary Results

We first present estimates of the proportion of a simulated community connected to the largest connected component through kinship ties between spouses, parents, children and siblings. Figure 5 shows estimates for one of the current Ebola infected countries, Liberia, over time. The model was run 100 times for this country, which is why there are many estimates at each point in time. The steep rise in connectivity after 2000 seen here provides suggestive evidence that the specific timing of the demographic transition in this country impacted the network in a way that may correspond to the current Ebola epidemic. We further explore this possibility in the full paper for all eight countries.

Figure 5. Proportion connected for a simulated village with Liberia's vital rates.