Aldo Benini
A note for ACAPS

Persistence of severity patterns over time

A case study of two estimates in Nepal

---

16 February 2016

Contents

Tables and figures

Summary

After sudden-onset disasters, needs assessments may produce severity estimates repeatedly, apace with updated information. We expect updates to add new information; at the same time, if impacts last, we also expect that earlier severity estimates resonate in later ones.

The extent to which severity persists is not well known. To our knowledge, nowhere have severity measures from subsequent needs assessments been correlated. There is a practical interest in gauging this correlation. If needs change rapidly in nature, degree and direction, severity estimates should be repeated frequently. If the severity pattern is more persistent, estimates have a longer shelf-life. Subsequent assessment may then focus more on indicators with direct operational value.

Two measurements in Nepal, following the earthquakes in 2015, provide an opportunity to study the degree and structure of persistence. UNOCHA produced severity measures for low-level administrative units – Village Development Committees (VDCs) - in April and June. Over 600 VDCs were evaluated at both points in time. Although the indices used for the purpose employed different indicators, they can be correlated.

Figure 1: Scatterplots of the April vs. June severity measures, by district

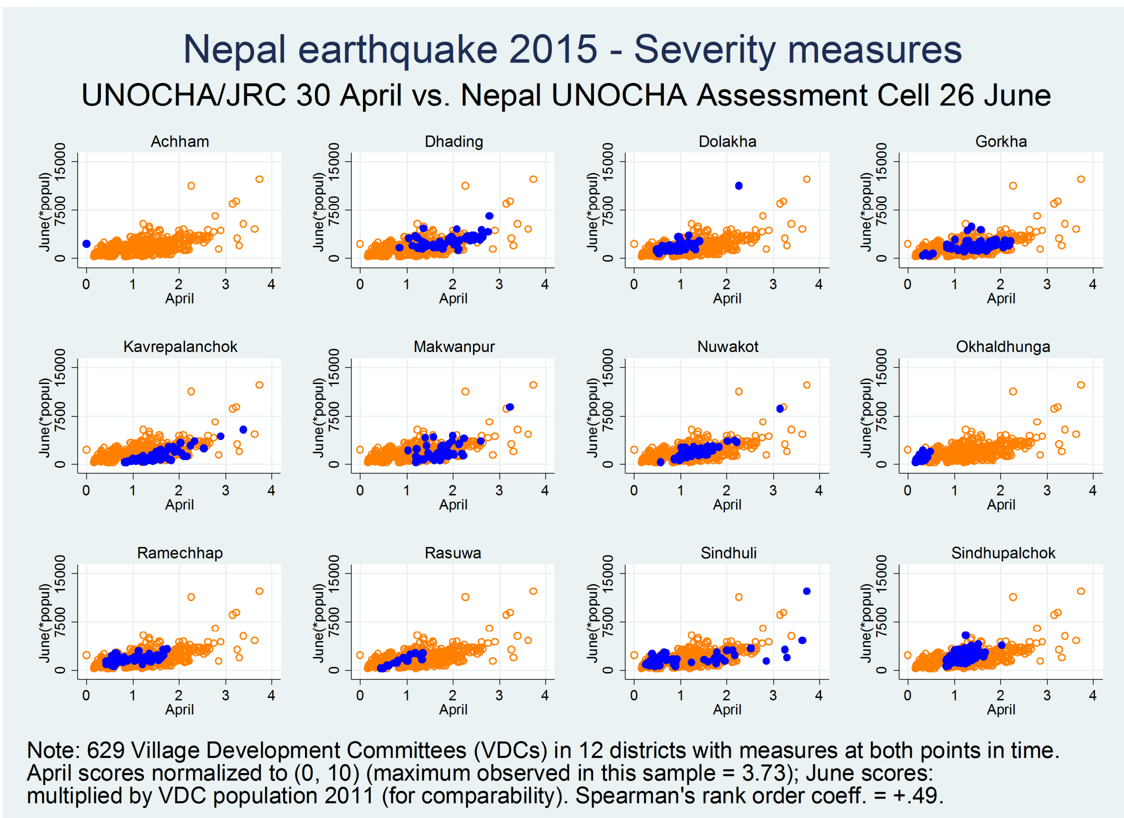

At face value, the correlation between the two estimates is moderate (+.49). It suggests that the impact on the affected communities changed considerably within two months. This is to be expected because a second earthquake struck Nepal during this period, with an epicenter in a different district from the first. However, some of the variability is a statistical artifact. Both indices incorporated indicators with fixed weights. When these are released, the connection between the indices shows up much stronger (+.76).

To do so, we rely on an established statistical technique – canonical correlation analysis – that tests how closely two sets of indicators measure the same concept. The same technique can be employed to express the degree to which a composite measure – severity in our case – has changed between two points in time. The indicators collected at the two points need not be the same – and, in fact, in Nepal they were largely different. For example, the strength of the earthquake was at first measured geophysically, as the Modified Mercalli Index; in the second assessment a combination of casualty, building damage, and food security indicators took its place.

The much higher correlation coefficient that we obtain by this method leads to a paradoxical conclusion. On one side the severity pattern shows persistence. On the other, UNOCHA and its partners collected several new indicators during the two months. These they combined in two sub-indices – physical and socio-economic vulnerability – that are scarcely correlated with the components of the first severity index. In other words, they contribute genuinely new information.

Still the question can be asked whether this second severity measure added much value, the way it was done. Both times, the assessment produced one global severity index value for each VDC. The second index, in June, included indicators with primary information about three sectors only – one each about shelter, health and food security. It was a considerable feat to collect them in over 600 VDCs. But it is hard to imagine that the humanitarian community was able to fine-tune its response, leaning on one combined severity index, to this much variability among small local communities.

The operational value of the second assessment might have been greater if instead more effort had been expended to estimate persons in need. These estimates might have been limited to the district level, or perhaps to district headquarter towns, plus and small samples of outlying communities. There were only twelve affected districts. Estimates, in sectors critical to survival and recovery, of persons in acute need as well as of those in moderate need would not only provide a basis for severity estimates, but they would also be of direct value to response planners and implementers.

For Nepal, that question is now academic. We cannot generalize from this one study to severity patterns in other places, disaster types or time intervals. But whenever we believe that the severity pattern is stable, then the initial estimates may buy us enough time to refine subsequent assessments to produce greater operational information value.

Information across successive assessments

Current needs assessment doctrine aims at a rapid, coarse situation overview after the onset of a crisis or disaster. This is to be followed by successively finer-grained and elaborate severity estimates as the humanitarian response community grows more knowledgeable with the affected areas, sectors and communities (NATF 2015). The amount of assessment resources fielded for the purpose is scaled up steeply during the initial days, weeks, and even months. It will peak, then decrease, then perhaps plateau at a level deemed adequate for continuing surveillance, only to produce intermittent spikes such as when funding cycles call for updated information.

The sequence of assessments can be looked at in the perspective of information value. The value of humanitarian information is impossible to state in absolute terms. It is impossible to know what decisions the responders would have made had they not received detailed needs assessments. Moreover, even if we knew the uses of this information from close observation, the value of different sectoral interventions could not be combined in a unified metric. In relative terms, however, some legitimate questions can be asked about the added value of sequential assessments, or at least about the degree to which subsequent needs estimates depart from the early ones. The answers will be relative, not only in the sense that we do not have an absolute basis for information values, but even more so because the criteria for relevant information may change as the response unfolds. Notably, the requirements for very specific information in terms of sectors and affected groups increase.

The meaning of persistence

From an assessment policy viewpoint, a general question of interest is to what extent the information collected at a later stage (e.g., two months into the response) is already implicitly contained in the information that was available soon after the sudden onset (e.g., in the first week). If we may presume that generally this is the case to a low degree only, then the situation must have changed a lot, and a fresh, detailed assessment is justified.

Conversely, it may happen that the initial information persists to a high degree. "Persistence" means that for any two units of interest – e.g., local communities X and Y – if X ranked higher than Y on the initial metric A (e.g., severity measured in the first week), there is a high probability that later on metric B (e.g., severity measured after two months, in part using different questions), X again ranks higher than Y. If so, this may have different practical consequences. Conceivably, one would want to continue using the initial assessment and concentrate the assessment resources on narrower domains, such as on a selection of the most highly impacted units, or on technical aspects of the response plan.

It is important to see the finer points in this logic. First, it is unaffected by the change of metric as long as the underlying concept remains the same – severity, impact, etc. What matters is the repeated comparison between units of interest, and the ability of an earlier ranking to predict a later one. Second, it is affected only to a degree by the refinement of measurements that the assessment effort achieves over time. Initially, estimates such as of severity may be limited to areas as large as provinces, to be refined subsequently down to the district and sub-district levels. However, the measure of interest will likely remain locally correlated, notably in space or for the same social groups. Thus the probability that after two months community X in province X' outranks community Y in province Y' may still be fairly high if on the initial (province-level) measure X' outranked Y'. To exemplify by an earthquake disaster: If X', and therefore X, is close to the epicenter, and Y' and Y rather far from it, knowing that X' is more impacted than Y' lets us predict that X is more impacted than Y. Assessment teams that two months later go to both X and Y will likely confirm this.

[Sidebar:] Mutual information and correlation

When we talk about information in the course of time, meanings often differ. They depend on everyday language or on particular scientific disciplines. A sentence like “Much of the previous information is still valid” will likely be understood as “The things that we are looking at are still the same; only the attributes of some have changed.” For example, a patient visiting his doctor’s office again may confirm that the administrative information has remained the same – name, address, insurance policy -, but his mobile phone number has changed. The phrase “the information persists” is not taken from everyday language; yet it still signals identity of objects, but change only in their attributes. For example, for psychologists, “information persistence” is a measure of visual memory (Wikipedia 2015).

In an information-theoretic perspective, however, the attributes that we compare at different points in time are not necessarily the same. To stay with our example of patients visiting doctor’s offices, it would be possible to compare patients’ cities and postal codes collected at time X to their insurance companies, member cards for which were issued and verified at a later time Y. The relationship between the two is not random. Translated to statistical language, the meanings of variables are less interesting than the relationship between their distributions. Information scientists look for a measure that expresses the non-randomness in the relationship between two variables in abstract and flexible manner.

One such measure is known as “mutual information” (Wikipedia 2016b). Knowing the value of attribute A may tell us something about the likely values of attribute B, and vice versa. The form of the relationship may be highly variable. The Wikipedia article comes with a chart of examples of bivariate distributions. In terms of their mutual information scores, some very different distributions have strikingly similar information values.

Figure 2: Mutual information - Examples of distributions

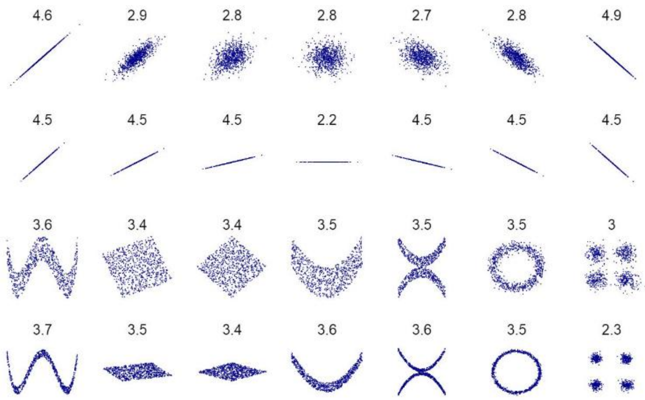

The mutual information statistic is most appropriate when the attributes compared are categorical and unordered. Patients' residences and health insurance is a case in point. In the social sciences, the mutual information concept has given rise to a number of measures expressing the strength of such relationships, such as in residential segregation or in the occupational structure of industries. For comparisons on ordered and continuous variables, however, mutual information is not a satisfying approach. Information persistence here needs measures that express the extent to which order is preserved over time, or even, how far the measurements deviate from the mean in direction and distance. This requirement takes us to the domain of correlations. This diagram, taken from the Wikipedia article on correlation (Wikipedia 2016a), depicts some of the same distributions; yet the relative strength of the correlations differs greatly from that of the mutual information measures.

Figure 3: Correlations - Examples of distributions

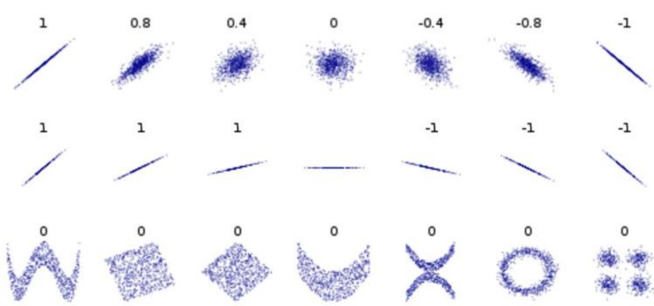

In particular, all the distributions in the bottom row display zero correlation. All the strictly linear ones in the middle row return a coefficient of 1 or -1 or, if flat, are undefined.

Our topic – severity – is measured *at least* at an ordinal level, in which case comparisons over time call for rank-order correlations. Severity measurements in Nepal and elsewhere have produced indices on interval or ratio levels. Therefore, the information persistence can be expressed in correlation coefficients for continuous variables.

This does not mean we have to compare apples with apples, and oranges with oranges. We can compare the number of apples at time A with the number of oranges at time B. From the information-theoretic viewpoint, it is perfectly legitimate to study the correlation in between, say, the Mercalli Index values at the time of an earthquake and the cumulative number of fatalities two months later. What matters is the joint distribution of the two. The content of the indicators is critical for the validity of the severity measures at each point in time; for the relationship between them, only the non-randomness is of interest.

An opportunity in Nepal

To return to the question of information value, it may therefore be of interest for the humanitarian community to develop a notion of the persistence of information over time, and hence a better grasp of the need for repeated updates of the same types of information versus the judicious continued use of older information. By a strike of good fortune, two severity assessments carried out in the same disaster zone within a short period in 2015 permit us to gauge the persistence of information in one such situation where the added value of assessments was of concern. This note develops a measure of persistence in this one particular case. It is not more than an isolated, if first-of-its-kind, case study; we cannot generalize to other crisis situations, let alone to a general rule.

This opportunity arises from two published severity assessments in Nepal, after the earthquakes that struck the country in 2015, on April 25 and again on May 12. UNOCHA, supported by the EU Joint Research Center (JRC), produced a first assessment remotely, within days after the first quake (the final version was released on April 30) (JRC 2015). After the second tremor in May, the Nepal UNOCHA Assessment Cell produced a field assessment-based severity assessment at the district level (the socio-economic indicators, such as the poverty rate, used pre-disaster data). In late June, more than one month after the second tremor, the Assessment Cell refined this down to the Village Development Committee (VDC) level. The Cell documented rationale and composition of this measure in a conceptual note (Liew 2015).[1]

The metrics of the April and June assessments differ considerably. Yet both captured severity at a relatively fine grain, i.e. for low-level administrative units, the VDCs (the April assessment imputed a number of indicators that were then available only for districts). For over 600 VDCs, estimates are available at both points in time. Both metrics are on the ratio level, which enables statistical techniques to gauge the extent to which information from the first conveyed to the second.

The structure of the severity metrics

As mentioned, the metrics of the April and June assessments differ. The severity index formed, within days after the first quake, by remote sensing and from secondary datasets combines measures of hazard, exposure and vulnerability. The figure (JRC, op.cit.) is nearly self-explanatory.

[1] We thank Anthony Liew, United Nations, and Leonie Tax, ACAPS, for various additional comments and clarifications.

Figure 4: Conceptual model of severity measurement in April 2015

	NEPAL EARTHQUAKE SEVERITY INDEX				
Dimension	**Hazard (1)**	**Exposure (1)**	**Vulnerability (1)**		
Component	Earthquake Intensity	Total Population (1)	Housing (1)		Poverty (1)
Indicator			Wall type (1)	Roof type (1)	Human Poverty Index
Scale	Admin Level 4 (Village)	Admin Level 4 (Village)	Admin Level 3 (District)		Admin Level 3 (District)
Data source	*USGS*	*Nepal Census 2011*	*Nepal Census 2011*		*Nepal Human Development Report 2014*

Figure 1: Conceptual model of the index with weights shown by scale and in brackets. i.e. Hazard, Exposure and Vulnerability have same weight. Wall type and Roof type have same weight.

The 2011 population census counts proxy for exposure; as a result, the index is a population-weighted measure. The aggregation is multiplicative, which obviated the need for indicator weights. However, the components are exponentially dampened, as in

$$Severity = Hazard^{1/3} \times Exposure^{1/3} \times Vulnerability^{1/3}$$

(op.cit, page 3).

From that setup, the June index departed on several aspects. It replaced the multiplicative hazard-exposure-vulnerability trifecta with an additive one grouping the lower-level indicators in earthquake impact, physical vulnerability and socio-economic vulnerability sub-indices. These require weights, which are detailed in this schematic. Population-weighting is external to the index formula. This is deliberate, in order to demonstrate the drastic differences in the geographic severity distributions between the unweighted and weighted indices. Liew (op.cit.) presents maps of both flavors.

Figure 5: Additive index of severity measurement in June 2015

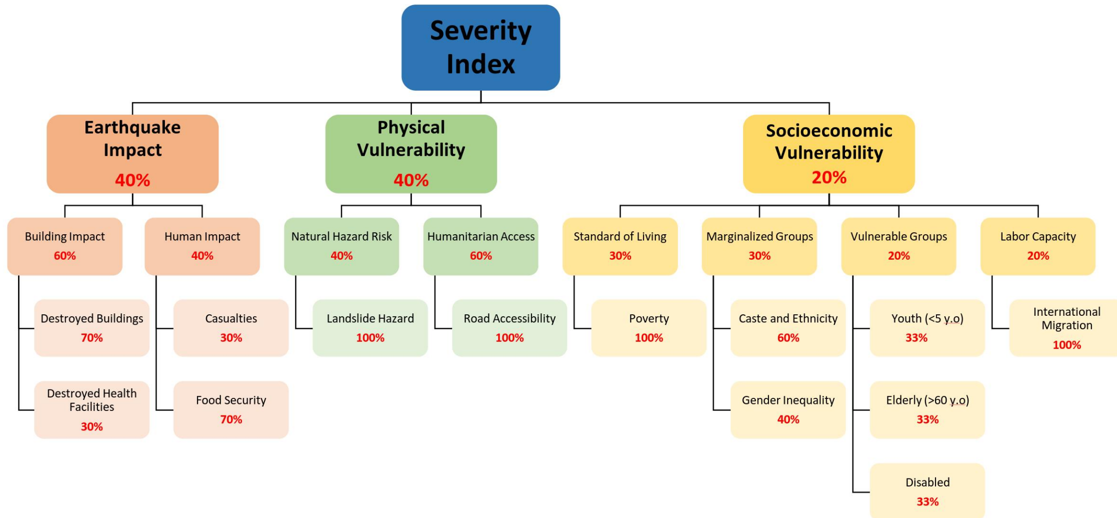

From April to June - Correlations

At face value

We first compare the indices at their face values. For comparability, we either have to divide the April index by its population component or to multiply the June index by the population. We do the latter, but without the exponential dampening that was applied in April. Thus our June measure is proportionate to the population-weighted index that the Nepal Assessment Cell used.

We visualize the correlations between the two measures, district by district, for the 629 VDCs for which they were taken at both points in time.

For technical reasons – deviations from normality -, we express the strength of the overall correlation through a statistic that looks at the ranks, rather than at the absolute values. The correlation coefficient is a significant, if not very strong +0.49. The scatterplots make it obvious that the agreement is particularly weak at the higher extreme. Nevertheless, a considerable part of the severity distribution in April conveys to the one observed in June.

Yet, the moderate +0.49 correlation value is not the last word. This face-value comparison underestimates the degree of agreement between April and June. The measures, as they are, are captive to functional forms and to fixed dampening powers, respectively fixed weights that fuse the indicators into the two indices. The real degree of information agreement is hidden; to unveil it, we must release those fixed parameters. We do this in two steps:

Hidden in the underlying distributions

First we establish identical aggregation modes for both indices. We make the aggregation mode of the April index additive by setting the components to their logarithms, as in:

$$log(Severity) = 1/3*log(Hazard) + 1/3*log(Exposure) + 1/3*log(Vulnerability)$$

Second, as we have seen, each indicator is formed of the three components; we let their weights float freely to the point where we maximize the correlation between the resulting indices. In other words, we estimate weights a_1, a_2, a_3, and b_1, b_2, b_3, such that the correlation between

$a_1*log(Hazard) + a_2*log(Exposure) + a_3*log(Vulnerability)$ [i.e., the April sub-indices]

and

b_1* *(Earthquake impact)* + b_2* *(Physical vulnerability)* + b_3* *(Socio-economic vulnerability)* [i.e., the June sub-indices]

is maximized. Note that we do not include population in the June index formulation because it would cancel out with the April exposure (both use census 2011 values). Since the June index is not population-weighted, we expect a_2 to be zero. If it is far from zero, then we must assume that the June index compensated for aspects of severity that the April index did not capture, and which are correlated with VDC populations. This has nothing to do with whether the index in *June* ought to be population-weighted or not, but whether the measures in *April* for hazard and vulnerability (Mercalli Index, pre-crisis poverty and building quality) by themselves anticipated the severity pattern found in June.

This maximum correlation is the true estimate of information conveyance in the statistical sense. We estimate its value through an established statistical procedure known as canonical correlation[2]. We show the optimized weights as well as the canonical correlation coefficient in this diagram.

[2] We document results in more detail in the appendix. There is no introductory article on canonical correlation in Wikipedia. However, a STATA manual chapter presents both formulas and an empirical example: *"Consider two scientists trying to describe how 'big' a car is. The first scientist takes physical measurements—the length, weight, headroom, and trunk space—whereas the second takes mechanical measurements—the engine displacement, mileage rating, gear ratio, and turning circle. Can they agree on a conceptual framework?"*. Their common concept is "size", but their indicators are entirely different. See http://www.stata.com/manuals13/mvcanon.pdf . In humanitarian information management, the procedure has rarely been used; Hoja, D. et al. (2013) for Haiti is an exception, albeit one involving far more demanding additional algorithms than in our basic usage here.

Figure 6: Canonical correlation between the major components of the April and June measures

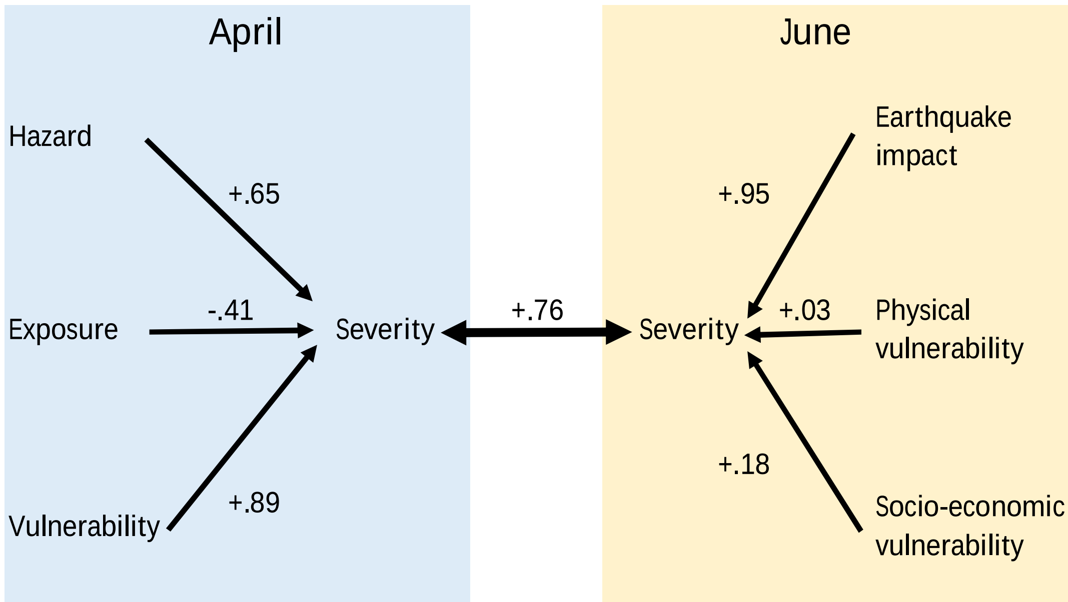

The correlation between these two severity constructs is now strong, much stronger than in their original, independent formulations. As long as we accept that the weights can float freely, and thus can be optimized for maximum conveyance, the result is clear and encouraging: Yes, a lot of the April information is preserved in the measures taken in June.

But the pattern of contributions from the sub-indices raises difficult questions of interpretation[3]. From the first assessment, in April, both hazard and vulnerability contribute strongly. The effect of exposure, measured solely by the VDC population, is negative – in other words, the April severity information conveys optimally only if larger communities are substantially penalized. Most likely, this is due to the fact that the vulnerability indicators in April were available only at the district level. In June, UNOCHA had VDC level data for them. Poverty and the quality of buildings probably – this remains to be tested – are more precarious outside the district headquarter towns (the urban districts of Kathmandu and Lalitpur are not covered in the June severity estimates). In the absence of detailed data, the VDC population size variable in part corrected for that.

From the second assessment, in June, the contributions all have the expected positive sign, but their strengths vary greatly. The high weight of the earthquake impact means that the information available in June contributed almost nothing that was not yet known in April. This is so despite the change in indicators - from physical (Mercalli) in April to societal (buildings, loss of life, food) in June.

[3] To clarify for readers conversant with canonical correlation: The coefficients are standardized; the contributions are thus comparable although the metrics of the sub-indices are different.

By contrast, the weight on the physical vulnerability is almost zero. This construct is new; almost nothing from April informs it. The weight on socio-economic vulnerability is also quite low – not surprisingly so because the poverty now shares its influence with several newly added indicators that had no precedent in April. In other words, both of these sub-indices contribute a lot of new information.

Back to (almost) face value

In the real life of severity estimates, weights do not float freely. The formula used in June fixed them precisely in the hierarchical additive aggregation scheme. These weights may not be optimal – determining weights via the Betti-Verma algorithm might have been preferable (Benini and Chataigner 2014), but this possibility is academic by now. What matters is what the Nepal Assessment Cell did, then and there.

In this perspective, it is interesting to describe the alignment between the three components of the April severity index – hazard, exposure, vulnerability – with the (population-unweighted) June index as calculated by the Assessment Cell. If we only think in linear relationships, the total correlation will be lower than in the canonical model (0.63, to be precise - because we do not optimize the sub-index weights[4]). However, we can now look out for non-linear influences, and these may be of interest particularly because a second earthquake struck between the two severity estimates, modifying the hazard distribution. In this local-regression graph on the next page, the April sub-indices have first been set to their logarithms (for additivity, as mentioned above), and then have been normalized in order to make the slopes comparable[5].

[4] Regression output is shown in the appendix.

[5] The local regression was done using Royston and Cox's *mrunning* multi-variate scatterplot smoother (Royston and Cox 2005). We only have graphic output (we did not save the prediction variables to the dataset); therefore we do not document this model further in the appendix. The R2 value was displayed at the end of execution and was copied by us into the edited graph.

Figure 7: Non-linear effects of the April sub-indices on the June index

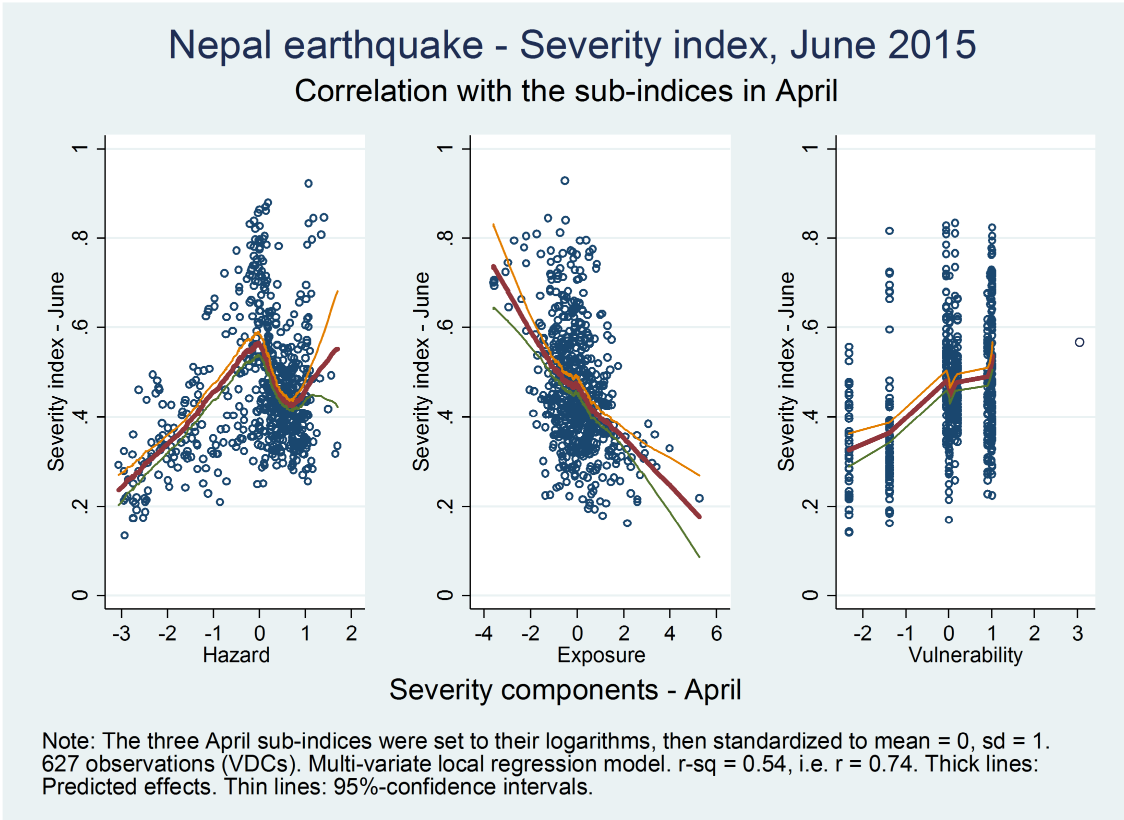

Admitting non-linear relationships, the overall correlation strength is again close to that of the canonical correlation. Practically, this finding does not help very much – since the forms of the relationships are not known in advance, in April we could not have predicted the further course anyway. However, the presence or not of strong non-linearities adds some ex-post insight. The effect of the hazard, as measured in April, obviously is strongly non-linear; i.e., VDCs shaken by tremors around the center of the transformed Mercalli Index range tended to be rated, by late June, as more severely affected. Those higher on Mercalli would up somewhat lower down than they would have on a straight line. This pattern may incorporate the impacts of the second earthquake, in May, which the index in June must have picked up through some of its indicators (more buildings destroyed).

Similarly striking is the linear relationship in the exposure (it is linear because a ruler can be stuck straight through the confidence interval). Apparently, the June index captures a heightened level of severity in smaller communities that the hazard and vulnerability sub-indices in April did not yet sense (again, we suspect that this is so because, within given districts, poverty and its measured correlates, such as disability, are higher in the remote and rural VDCs – a variation that was not measured in April).

Because the April vulnerability values are bunched, not much can be said about its linear of non-linear effect.

Something old, something new

Taken together, these findings are somewhat paradoxical:

- The severity index in June contains a lot of the information established in April already. "Information" here does not mean the use of the same indicators, but rather the extent to which the indices and their components correlate, and thus the extent to which the distribution of severity in June was anticipated by that already figured out in April.

- Nevertheless, the Assessment Cell incorporated substantial amounts of new information in the June version of the index. Physical vulnerability – landslides, impassable roads – was not implied in any of the April components. The socio-economic vulnerability model incorporated more facets than before; the distribution of its scores in June was thus predictable to a negligible degree only.

However, this updated severity information has substantial value only if we make some further assumptions. The weights on indicators and sub-indices must be meaningful. And, having one global severity score per VDC has significant operational value. With the latter assumption, we step outside the box of the strictly statistical discussion.

Severity measurement evolving

Reallocating the effort

A question that needs to be asked at this point is whether the effort to collect severity-relevant data after April could have been applied differently. The outside observer is immediately struck by the leap in granularity – from just 12 districts to 600+ VDCs, a ratio of 1 : 50. It is hard to imagine that the humanitarian community was able to fine-tune its response, leaning on one combined severity index, to this much variability among small local communities, even in the months after June.

This level of detail seems to have been purchased at the cost of comparable sectoral information across the affected region. On the earthquake impact side, the June index admits one indicator each from shelter, health and food security. Other sectors are not present; and for those three sectors the proportions of the persons in need are not visible – if they had been estimated, they were not brought into the severity format.

In future evolutions of severity measurement over the first few months after a sudden onset, one might want to consider a different distribution of the effort. There would be substantive – what to measure – as well as sampling considerations – which units, and how many, and where. For the first, early assessment, the district level (or its rough equivalent in the administrative hierarchy) and a small number of rapidly collected indicators, all relevant for the chosen process model, would do. In this sense, the UNOCHA/JRC estimates in April were adequate.

Estimating persons in need

Later on, the process model would have to be revised, as, in fact, the Assessment Cell did by anticipating the changed needs in the monsoon and following recovery periods. The indicators should then reflect more specific sectoral concerns. In particular, they should include estimates of *persons in need* in sectors that have been identified as critical to survival and recovery. Ideally, particularly in areas of endemic poverty, the persons-in-need estimates should distinguish between acute and moderate need (Benini 2015).In addition, local experts should be asked to rate the needs of their communities in various sectors on ordinal scales of severity that are clearly linked to long-term damage and excess mortality. Evolving to persons in need will overcome the duality between population-weighted and unweighted models that bedevils the operational value of the June estimates.

Adaptive sampling

But the more consequential changes may happen on the sampling side. It does not seem rational to collect primary data on 600 small local communities, but more rational to use the extant secondary data as well as extrapolations from a sample of freshly assessed communities. One might want to stratify, such as between district headquarter communities and outlying ones. All headquarters might be visited, plus a fraction of the outlying ones. The sampling rate for them might vary, based on the first district-based severity estimates. Alternatively, we might sample adaptively (Thompson and Seber 1996), snowballing to neighboring communities whenever we detect a particularly severely affected one.

These reflections are extremely sketchy. They are encouraged by the main finding from these two severity measures in Nepal. A considerable degree of the early information conveyed to the subsequent estimate even though the indicators differed between April and June, and a second major event (earthquake) intervened. This inspires confidence that early estimates, though necessarily rather coarse, will be good enough for long enough to permit elaborating the instruments for higher-resolution and more operational measures.

Appendix

Calculations were performed in the statistical application STATA, version 14.

Descriptive statistics

For 629 VDCs in 12 Districts, UNOCHA calculated severity scores in both April (remotely) and June (with on-the-ground primary data). 627 observations are used in the canonical correlation model (excluding two that had zero scores in the April dataset and were thus lost when the sub-indices were set to their logarithms).

April data (UNOCHA / JRC)

As in the UNOCHA / JRC dataset

variable name	storage type	display format	value label	variable label
intensnorm	float	%8.0g		April Hazard (normalized) [Modified Mercalli]
popnorm	float	%8.0g		April Exposure (normalized) [VDC population]
vulnerabscore	float	%8.0g		April Vulnerability (normalized)
severnorm	float	%8.0g		April Severity Score (normalized)

[normalized by UNOCHA]

Variable	Obs	Mean	Std. Dev.	Min	Max
intensnorm	629	1.31261	1.226784	0	10
popnorm	629	.0456784	.0482026	.0042544	.8680172
vulnerabsc~e	629	5.911452	.7419622	4.317401	8.832682
severnorm	629	1.236229	.6054248	0	3.730715

Log-transformed

April_log10_i~m	float	%9.0g	Hazard
April_log10_p~m	float	%9.0g	Exposure
April_log10_v~e	float	%9.0g	Vulnerability
April_lo~rscore	float	%9.0g	Severity score

Variable	Obs	Mean	Std. Dev.	Min	Max
April_log1..	627	-.1887024	.697609	-2.318761	1
Ap~0_popnorm	629	-1.431617	.2602934	-2.371158	-.0614717
April_log1..	629	.7680374	.057462	.6352224	.9460926
April~rscore	627	-.2843514	.2650593	-1.092906	.2598432

Standardized log-transforms

Ap~d_intensnorm	float	%9.0g	Hazard
April~d_popnorm	float	%9.0g	Exposure
April_log10_s..	float	%9.0g	Vulnerability

Variable	Obs	Mean	Std. Dev.	Min	Max
April_log1..	627	5.32e-09	1	-3.05337	1.703967
Ap~d_popnorm	629	1.69e-07	1	-3.609547	5.263849
April_log1..	629	-1.96e-07	1	-2.311354	3.098659

June data (Nepal Assessment Cell)

variable name	storage type	display format	value label	variable label
population	long	%8.0g		Population 2011 (VDC)
eqimpact	float	%8.0g		Earthquake impact score
physical	float	%8.0g		Physical vulnerability score
socioeconomical	float	%8.0g		Socio-economical vulnerability score
severityindex	float	%8.0g		Severity index - June (not population-weighted)
sevIndexXP~2011	float	%9.0g		Severity index - June - multiplied by VDC population size

Variable	Obs	Mean	Std. Dev.	Min	Max
population	629	4458.397	4701.332	415	84671
eqimpact	629	.6110371	.1947429	.1979476	.952288
physical	629	.2713831	.2917204	0	1
socioecono~l	629	.5356935	.1614404	.0532	.9868
severityin~x	629	.4601176	.1778059	.105	.923
sevInde~2011	629	1806.454	1143.757	313.962	12257.44

Canonical correlations

All three correlations were calculated, but only the first was used and reported in the main part of the note. The two severity measurements are closely related if their best linear combinations are highly correlated. The figure indicates that this indeed is the case.

Figure 8: Scatterplot of the first linear combinations, by district

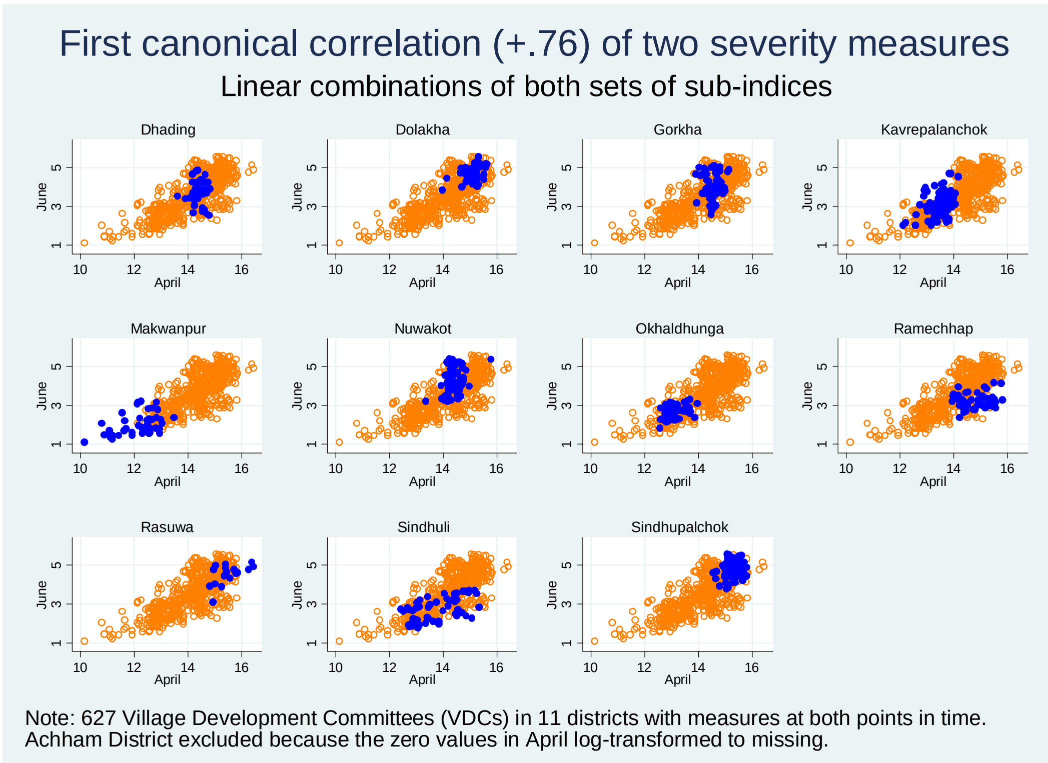

Command

```
. canon ( April_log10_intensnorm April_log10_popnorm April_log10_vulnerabscore) ( eqimpact
physical socioeconomical), stdcoef first(3)
```

Estimates

```
Canonical correlation analysis                      Number of obs =        627
```

Standardized coefficients for the first variable set

	1	2	3
April_log1..	0.6468	0.5830	-0.6624
Ap~0_popnorm	-0.4093	0.7198	0.6555
April_log1..	0.8939	0.2168	0.5488

Standardized coefficients for the second variable set

	1	2	3
eqimpact	0.9479	-0.5163	-0.1822
physical	0.0259	0.6872	0.9652
socioecono~l	0.1838	0.5593	-0.9201

Canonical correlations:

0.7619 0.3268 0.0277

Tests of significance of all canonical correlations

	Statistic	df1	df2	F	Prob>F	
Wilks' lambda	.374452	9	1511.5	83.5066	0.0000	a
Pillai's trace	.688013	9	1869	61.7986	0.0000	a
Lawley-Hotelling trace	1.50388	9	1859	103.5449	0.0000	a
Roy's largest root	1.38356	3	623	287.3190	0.0000	u

e = exact, a = approximate, u = upper bound on F

Post-estimation output:

Correlations

Correlations for variable list 1

	April_..	April_..	April_..
April_log1..	1.0000		
Ap~0_popnorm	0.2956	1.0000	
April_log1..	-0.3359	-0.2185	1.0000

Correlations for variable list 2

	eqimpact	physical	socioe~l
eqimpact	1.0000		
physical	0.4046	1.0000	
socioecono~l	0.1248	0.4002	1.0000

Correlations between variable lists 1 and 2

	April_..	April_..	April_..
eqimpact	0.2067	-0.2617	0.5669
physical	-0.0923	-0.3410	0.3223
socioecono~l	-0.1181	-0.3162	0.2069

Loadings

Canonical loadings for variable list 1

	1	2	3
April_log1..	0.2256	-0.7230	-0.6530
Ap~0_popnorm	-0.4135	-0.8447	0.3398
April_log1..	0.7661	0.1363	0.6281

Canonical loadings for variable list 2

	1	2	3
eqimpact	0.9813	-0.1684	0.0936
physical	0.4830	0.7021	0.5233
socioecono~l	0.3124	0.7699	-0.5565

```
Correlation between variable list 1 and canonical variates from list 2

             |        1         2         3
-------------+------------------------------
April_log1.. |   0.1718   -0.2362   -0.0181
Ap~0_popnorm |  -0.3150   -0.2760    0.0094
April_log1.. |   0.5837    0.0445    0.0174
--------------------------------------------

Correlation between variable list 2 and canonical variates from list 1

             |        1         2         3
-------------+------------------------------
    eqimpact |   0.7476   -0.0550    0.0026
    physical |   0.3680    0.2294    0.0145
socioecono~l |   0.2380    0.2516   -0.0154
--------------------------------------------
```

Linear regression of June scores on log April sub-indices

On page 13, we mention a lower overall correlation value when these relationships are estimated in a linear regression model.

Description of variables

```
variable name   type    format     label      variable label
-----------------------------------------------------------------------------
severityindex    float    %8.0g                   Severity index - June (not population-
weighted)
Ap~d_intensnorm float   %9.0g                 Hazard
April~d_popnorm float   %9.0g                 Exposure
April_log10_s.. float   %9.0g                 Vulnerability

    Variable |        Obs        Mean    Std. Dev.       Min        Max
-------------+---------------------------------------------------------
severityin~x |        629    .4601176    .1778059       .105       .923
April_log1.. |        627    5.32e-09           1   -3.05337   1.703967
Ap~d_popnorm |        629    1.69e-07           1  -3.609547   5.263849
April_log1.. |        629   -1.96e-07           1  -2.311354   3.098659
```

Model (severityindex = June!)

```
.     regress     severityindex     April_log10_std_intensnorm     April_log10_std_popnorm
April_log10_std_vulnerabscore

      Source |       SS           df       MS      Number of obs   =       627
-------------+----------------------------------   F(3, 623)       =    140.43
       Model |  7.99824606         3  2.66608202   Prob > F        =    0.0000
    Residual |   11.827569       623  .018984862   R-squared       =    0.4034
-------------+----------------------------------   Adj R-squared   =    0.4006
       Total |  19.8258151       626  .031670631   Root MSE        =    .13779

-----------------------------------------------------------------------------------------------
                severityindex |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
------------------------------+----------------------------------------------------------------
   April_log10_std_intensnorm |    .051643   .0060254     8.57   0.000     .0398105    .0634756
      April_log10_std_popnorm |  -.0655139   .0058072   -11.28   0.000    -.0769179   -.0541098
April_log10_std_vulnerabscore |   .0927126   .0059818    15.50   0.000     .0809657    .1044595
                        _cons |   .4606034    .005503    83.70   0.000     .4497968      .47141
-----------------------------------------------------------------------------------------------
```

Overall correlation

The square root of the Adj R2 is 0.4006^0.5 = .6329297.

References

Benini, A. (2015). Moderate need, acute need. Valid categories for humanitarian needs assessment? Evidence from a recent needs assessment in Syria [26 March 2015]. Geneva and Beirut, ACAPS and MapAction.

Benini, A. and P. Chataigner (2014). Composite measures of local disaster impact - Lessons from Typhoon Yolanda, Philippines. Geneva, Assessment Capacity Project (ACAPS).

Hoja, D., T. Krauss and P. Reinartz (2013). Detailed damage assessment after the Haiti earthquake. Earth Observation of Global Changes (EOGC). J. M. Krisp and E. al. Heidelberg, New York, Dordrecht, London, Springer**:** 193-204.

JRC (2015). NEPAL EARTHQUAKE SEVERITY INDEX – VERSION 4, 30 April 2015. Ispra, Italy, Joint Research Center of the European Commission (JRC).

Liew, A. (2015). "Concept note: Severity Index [16 July 2015]." Kathmandu, Nepal Earthquake Assessment Unit. Retrieved 28 July 2015, from https://www.humanitarianresponse.info/en/operations/nepal/document/150716-severity-index.

NATF (2015). Multi-Sector Initial Rapid Assessment (MIRA) Guidance [Revision July 2015]. Geneva, United Nations Inter-Agency Standing Committee Needs Assessment Task Force (IASC/NATF)

Royston, P. and N. J. Cox (2005). "A multivariable scatterplot smoother." Stata Journal **5**(3): 405-412.

Thompson, S. K. and G. A. F. Seber (1996). Adaptive Sampling. New York, Wiley-Interscience.

Wikipedia. (2015). "Iconic memory." Retrieved 3 February 2016, from https://en.wikipedia.org/wiki/Iconic_memory.

Wikipedia. (2016a). "Correlation and dependence." Retrieved 3 February 2016, from https://en.wikipedia.org/wiki/Correlation_and_dependence.

Wikipedia. (2016b). "Mutual information." Retrieved 3 February 2016, from https://en.wikipedia.org/wiki/Mutual_information.