Story Builders

Jacob Glazer

The Coller Faculty of Management, Tel Aviv University
and Department of Economics, Warwick University
glazer@tauex.tau.ac.il

Ariel Rubinstein

School of Economics, Tel Aviv University
and Department of Economics, New York University
rariel@tauex.tau.ac.il

March 14, 2020

ABSTRACT: We study methods for constructing a story from partial evidence where a story is defined as a path along a finite directed graph from the origin to a terminal node. Each node in the graph represents a possible event. A *story builder* receives evidence, i.e. a subset of events consistent with at least one story, and expands it into a coherent story. The analysis focuses on a consistency property whereby if the story builder believes in a particular story, given a certain set of facts, then he believes in it given a broader set of facts consistent with the story.

AEA Classification: D01.

KEYWORDS: Story builder, non-Bayesian beliefs.

1. Introduction

It is a common occupance that an individual receives partial evidence about some chain of events and develops it into a complete and coherent "story". In constructing the story, the individual is aware of the constraints on how the story can develop. We are interested in methods of constructing a full story from partial evidence. Given that people often fail to apply Bayesian reasoning even in very simple situations, the approach taken will be non-Bayesian.

In our setting, a story is a sequence of events, i.e., a path along a finite directed graph (without cycles), which starts from the origin and continues along the graph until it reaches a terminal node. Each node in the graph represents a possible event. A *story builder* receives some evidence in the form of a subset of events consistent with at least one story and expands it into a full and coherent story that he believes to be true. Thus, the story builder holds a point belief rather than a probabilistic belief about the real story.

The formalization of the concept of a "story builder" makes it possible to define and analyze a variety of procedures. Prominent among them is the "rational" story builder who has in mind an ordering over the set of possible stories and chooses the story that maximizes that ordering, given the evidence. The ordering might embody a belief about the likelihood of the stories, in which case he chooses the most likely one. Alternatively, the ordering might reflect wishful thinking, and then the story builder selects the best story that does not conflict with the evidence.

Whereas the rational story builder approaches the situation holistically, other types of story builders construct the story in steps, sequentially adding events according to some rule. For example, the story builder might have in mind a probability measure over the set of possible stories and will advance from the origin by selecting the most likely event at each stage, given the evidence and the path he has chosen so far.

Much of the analysis centers around what we call Story Consistency which states that if, after receiving a particular evidence set, the story builder believes in some story, then he continues to do so if he receives the same evidence set together with an additional piece of evidence consistent with the story.

If the graph is a tree, then Story Consistency is satisfied only by a story builder who can be described as rational. If not, then there are story builders that satisfy the property but are not "rational". In most of the paper we assume that the story builder is naive and does not take into account the source of the evidence. However, we also comment on a story builder who believes that the evidence is presented by a party with a vested interest.

Finally, we present an example of a matching game in a large population consisting of various types of story builders that are matched into pairs. Each of them needs to come up with a story identical to that of his partner, given a particular graph. We show that miscoordination may take place in equilibrium, and (unlike the case of mixed strategies equilibria in a standard game) this equilibrium is robust to small changes in the distribution of types in the population.

The most closely model is that of Sadler (2018), in which an agent encounters information in the form of propositions that are either true or false. A proposition is identified according to the subset of states in which it is true while a belief is a set of propositions held by the individual. An individual is modeled as an updating function that determines a belief as a function of a previous belief and a new proposition. This formalization provides a language for specifying a number of updating rules that are not necessarily consistent with Bayesian reasoning.

In Bjorke (2019), the set of "states" is a product set. The individual receives information about some components of the vector and possesses a "focal state" function (analogous to our story builder) that completes any subset of characteristics so as to become a full vector. An individual chooses a "focal state" for every subset of values. Bjorke (2019) focuses on two functions, "most likely" and "most distinctive", and investigates the following implementability question: What are the states that can be evoked given the true state?

Eliaz and Spiegler (2018) focus on a decision maker who interprets objective data about the realizations of variables in terms of a causal model (a Bayesian network). The decision maker in their framework receives statistical data and builds a model (a narrative) that organizes the data.

2. The model

Let $G = \langle A, \rightarrow \rangle$ be a finite directed graph. The set A is a finite set of possible *events,* each of which may or may not have occurred. The existence of an arch $a \rightarrow b$ in the graph means that the event b may succeed the event a. The event a is *terminal* if there is no event b such that $a \rightarrow b$. A *path* in G is a sequence of events $(a_0, a_1, ..., a_K)$ such that $a_k \rightarrow a_{k+1}$ is an arch in G, for every $0 \leq k \leq K-1$.

We restrict our attention to graphs satisfying two properties:

(i) There is no cycle. That is, G does not have a path of the form $(a_0, a_1, \ldots, a_K = a_0)$ with $K \geq 1$. It follows that a terminal event exists.

(ii) The graph contains an event, denoted by O (the starting point), such that for every $a \in A$ there is at least one path from O to a.

A *story* in G is a path $x = (x_0, .., x_{l(x)})$ where $x_0 = O$ and $x_{l(x)}$ is terminal. The integer $l(x)$ is the length of the story x. The set of all stories is denoted by S. A *partial story* is a path that starts at O and does not necessarily end at a terminal event.

For any two stories x and y, denote by $j(x, y)$ the maximal j for which $(x_0, \ldots, x_j) = (y_0, \ldots, y_j)$. That is, $j(x, y)$ is the length of the longest partial story joint to x and y. Define $a_{xy} = x_{j(x,y)} = y_{j(x,y)}$. That is, a_{xy} is the event at which the two stories split.

We say that the event a *appears (weakly) before* the event b if $a = b$ or there is a path from a to b. By the assumptions on the graph, this relation is anti-symmetric (and may be incomplete) and a appears before b if and only if there is a story in which a appears before b.

We have in mind that one and only one of the stories is the *true story*. An individual gets to see evidence in the form of a set of events that have occurred. We do not allow the individual to receive direct information about an event that has not occurred, though he might infer it from the graph and the evidence he receives. Formally, an *evidence set* is a subset $E \subseteq A$ such that there is at least one story $s \in S$ that is *consistent* with E, in the sense that it passes through all the events in E. The model does not specify what the source of the evidence is; in particular the individual may come across the information himself, or an interested party might provide it to him.

The individual is familiar with the graph G, observes an evidence set and builds a story that is consistent with the evidence. A *story builder* is a method of developing any evidence set into a story. Let $\mathbb{E}$ be the set of all evidence sets. For every $E \in \mathbb{E}$, let S_E be the (non-empty) set of stories consistent with E. A *story builder* (for the graph G) is a function that assigns a unique story in S_E to every $E \in \mathbb{E}$.

The "rational" story builder. A prime example of a story builder is based on a strict ordering $\succsim$ on the set S. For each evidence set E, the story builder $F_{\succsim}$ selects the story in S_E that is $\succsim$ -maximal according to the ordering. The ordering $\succsim$ can represent a variety of psychological phenomena. In particular, it might represent a likelihood relation between the stories. Under this interpretation, the rational story builder can be thought of as a Bayesian agent who uses only point-wise beliefs. Other interpretations include *wishful thinking* (the ordering describes the story builder's preferences regarding what he wishes the truth to be) and *simplicity seeking* (the story builder seeks a simple story consistent with the evidence). We say that a story builder F is *rationalizable* if there is an ordering $\succsim$ such that $F = F_{\succsim}$.

3. The story builder as a choice function

A story builder who observes an evidence set E chooses a story in S_E, the set of stories consistent with E. Our notion of a story builder allows for the individual to believe in two different stories after receiving two different evidence sets that share the same set of consistent stories. In other words, if we think about the graph as a description of the logical constraints on a story, then the concept of a story builder allows for the story builder to believe in two different stories given two evidence sets that logically lead to the same conclusion (i.e. the same set of consistent stories). The following property excludes such a possibility.

Invariance: A story builder F satisfies the *invariance* property if $F(E) = F(E')$ whenever $S_E = S_{E'}$.

When a story builder satisfies the invariance property it can be thought of as a choice function with the (restricted) domain $D = \{S_E \mid E \text{ is an evidence set}\}$. The graph G imposes restrictions on the story builder's domain. For example, if G is a tree and E is an evidence set, then S_E is the set of all stories that include the event in E that is the furthest away from O. Thus, if the graph is a tree, then the number of elements in the domain D is at most $|A|$.

The choice from the sets of size 2 is often a basic ingredient of a choice function. However, and regardless of the graph, there is only a limited set of subsets of size 2 in D.

Claim 1: For no graph G are there three stories $p, q, r \in S$ and three evidence sets E_{pq}, E_{qr} and E_{pr}, such that $S_{E_{pq}} = \{p, q\}$, $S_{E_{qr}} = \{q, r\}$ and $S_{E_{pr}} = \{p, r\}$.

Proof: Assume to the contrary that such a graph G exists. Recall that the relation "the event a appears before the event b" is anti-symmetric.

Let x and y stand for any two stories in $\{p, q, r\}$. If the stories x and y share an event after a_{xy} (the event after which x and y split), then denote by b_{xy} the first event after a_{xy} that is common to both x and y. The stories x and y must coincide after b_{xy}, since otherwise the evidence E_{xy} would be consistent with at least four stories (there are at least two possibilities for the path between a_{xy} and b_{xy} and two for the continuations from b_{xy}).

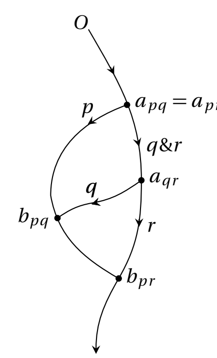

Figure 1

Figure 1 illustrates the rest of the proof. At the earliest event where the three stories p, q and r do not coincide, one or them, say p, splits from the other two, which may split as well. Thus, $a_{pq} = a_{pr}$ and a_{qr} appears (weakly) after a_{pq}. The set E_{pq} must contain an event after a_{pq} since otherwise $r \in S_{pq}$. Thus, b_{pq} exists. Similarly, b_{pr} exists and without loss of generality, b_{pq} appears (weakly) before b_{pr}. Since b_{pr} exists, all events in E_{pr} are either between O and a_{pr} or (weakly) after b$_{pr}$. Since p and q coincide after b_{pq} and since q and r coincide after b_{qr}, all events in E_{pr} are either between O and $a_{pq} = a_{pr}$ or (weakly) after b_{pq}, which implies that $q \in S_{pr}$, a contradiction. ■

Obviously, given a graph G, any choice function with the restricted domain D induces a story builder. If the choice function does not satisfy Sen's property α, then the story builder is not rationalizable. Following is an example: The story builder is a judge. The stories in S are the possible chains of events in a particular case. The set S is partitioned by the judge into two subsets: P ("punish") and N ("don't punish"). The judge has in mind a prior p over S. Given an evidence set E, he compares $p(P \cap S_E)$ to $p(N \cap S_E)$. If the former is larger, he finds the defendant guilty and justifies the decision based on his belief in the most likely story in $P \cap S_E$; otherwise, he finds the defendant not guilty and justifies the verdict based on his belief in the most likely story in $N \cap S_E$. That is, the judge first makes up his mind whether the defendant is guilty or innocent, given the evidence that has been brought before him, and only then does he adopt the most likely story given the evidence and given his verdict. Such a judge may not be rationalizable. For example, given the tree and the probability measure on S presented in Figure 2, the judge believes (O, d, b) if he does not see any evidence and (O, d, c) if he sees the evidence set $\{d\}$ even though d is consistent with (O, d, b). Using the concept defined in the next section, the judge does not satisfy "story consistency". A similar example can be constructed for the case in which the judge finds the defendant guilty only if he believes in the verdict "beyond a reasonable doubt" (i.e. above a certain probability).

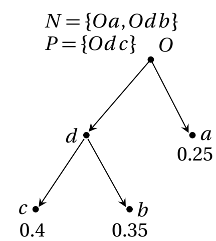

Figure 2

4. Story Consistency

We focus on the property of a story builder which states that if, after receiving a particular evidence set, the story builder believes a certain story, then he holds the same belief if the evidence set contains one additional event consistent with the story.

Story Consistency. A story builder F satisfies Story Consistency if $F(E \cup \{a\}) = F(E)$ for every evidence set E and any event a in the story $F(E)$.

Claim 2: If F satisfies Story Consistency, then F satisfies the invariance property.
Proof: Assume that $S_E = S_{E'}$. By definition, $F(E) \in S_E$ and therefore $F(E) \in S_{E'}$ and as a result the events in $E' - E$ appear in $F(E)$. Applying Story Consistency successively to the events in $E' - E$ yields $F(E \cup E') = F(E)$. Similarly, $F(E \cup E') = F(E')$. Thus, $F(E) = F(E')$. ■

The story builder in our model receives evidence only once and then forms his belief. The model does not specify the order in which the evidence arrives. Consider a story builder with a function F who first receives E_1 and then E_2. He might ignore the order they arrived in and adopt the belief $F(E_1 \cup E_2)$. Or he might first form the belief $F(E_1)$ and then stick with it if E_2 is consistent with $F(E_1)$. If F satisfies Story Consistency, then the story builder reaches the same belief under both possibilities. If F does not satisfy Story Consistency and if the story builder uses the second approach, then a speaker, who possesses the evidence set $E_1 \cup E_2$ and is obliged to present all the evidence, might manipulate the order in which the evidence is presented in order to manipulate the story builder's conclusion.

Obviously, a rational story builder satisfies Story Consistency. We now show (Claim 3) that if the graph is a tree then Story Consistency implies that the story builder is rationalizable. Somewhat surprisingly, we provide an example (Claim 4) of a graph and a story builder who satisfies Story Consistency but is not rationalizable.

Claim 3. Assume that G is a directed tree. If F satisfies Story Consistency, then F is rationalizable.

Proof: Given F, we say that the story x is revealed to be better than the story y, denoted by $x \succsim y$, if there exists an evidence set E such that $F(E) = x$ and $y \in S_E$. We show that the relation $\succsim$ is anti-symmetric and transitive and thus can be extended to a preference relation. Then, by definition, $F(E)$ is the $\succsim$-maximal element in E.

Anti-symmetry (implied by Story Consistency even if G is not a tree). Assume that $x \succsim y$ and $y \succsim x$, that is, there are evidence sets E and E' such that $x, y \in S_E \cap S_{E'}$, $F(E) = x$ and $F(E') = y$. Both x and y are in $S_{E \cup E'}$. Since $F(E) = x$ and $E' - E$ contains only events that appear in x, it follows from Story Consistency that $F(E \cup E') = x$. Similarly, $F(E \cup E') = y$, which implies that $x = y$.

Transitivity. Assume that $x \succsim y$ and $y \succsim z$. Let E be an evidence set such that $x, y \in S_E$ and $F(E) = x$, and let E' be an evidence set such that $y, z \in S_{E'}$ and $F(E') = y$. Both a_{xy} (the split event of the stories x and y) and a_{yz} are in y. The event a_{xy} appears (weakly) before a_{yz}. Otherwise, and since G is a tree, E' consists of events in y that appear (weakly) before a_{yz} and therefore $x \in S_{E'}$, implying $y \succsim x$ and contradicting the anti-symmetry of $\succsim$. Since G is a tree, E consists only of events in y that appear (weakly) before a_{xy} and since z splits from y not before a_{xy} the story z is also consistent with E, implying that $x \succsim z$. ■

Claim 4: There exists a graph and a story builder that satisfies Story Consistency but is not rationalizable.

Proof: Let $G = \langle A, \rightarrow \rangle$ be the graph where $A = \{1,2,..,T\} \times \{0,1\}$ with $(t,\delta) \rightarrow (t+1,\delta')$ for any t,δ,δ'. Assume $T \geq 4$. An event (t,δ) can be thought of as "δ happens at date t". Any event with index t can be followed by any event with index $t+1$. Obviously, G is not a tree. A story $(O,(1,\delta_1),...,(T,\delta_T))$ can be thought of as a sequence of zeroes and ones of length T and thus can simply be denoted by $(\delta_1,\ldots,\delta_T)$.

Define $F(\emptyset) = (0,\ldots,0)$ and F($\{(t^1,\delta^1),...,(t^K,\delta^K)\}$) where $1 \leq t^1 < t^2 < .... < t^K \leq T$ as the story $(a_1,..,a_T)$ where $a_t = \delta^l$ for any $t^l \leq t < t^{l+1}$ and $a_t = \delta^K$ for $t < t^1$ or $t^K \leq t \leq T$. In other words, if we think about $1,\ldots,T$ as a cycle (1 comes after period T), then any evidence that δ occurs at t is taken as a proof that δ persists from t on, until contradictory evidence is received. Thus, for example, when $T = 7$ we have $F(\{(2,1),(4,0),(6,0)\}) = (0,1,1,0,0,0,0)$.

This story builder satisfies Story Consistency but cannot be rationalized. For example, if (for $T = 4$) F was rationalized by $\succsim$ then:

$F(\{(1,1),\ (3,0)\}) = (1,1,0,0) \succ (1,0,0,1)$

$F(\{(2,0),(4,1)\}) = (1,0,0,1) \succ (0,0,1,1)$

$F(\{(1,0),(3,1)\}) = (0,0,1,1) \succ (0,1,1,0)$

$F(\{(2,1),(4,0)\}) = (0,1,1,0) \succ (1,1,0,0)$, a contradiction. ■

A "slight" modification of the story builder described in the proof can be rationalized. Define $F(\{(t^1,\delta^1),...,(t^K,\delta^K)\})$ where $1 \leq t^1 < t^2 < .... < t^K \leq T$ as the story $(a_1,..,a_T)$ where $a_t = \delta^l$ $(l = 1,...,K-1)$ for every $t^l \leq t < t^{l+1}$, $a_t = \delta^1$ for all $t < t^1$ and $a_t = \delta^K$ for all $t \geq t^K$. In other words, the story builder initially adopts $a^1 = \delta^1$. He proceeds to build the story so that $a_t = a_{t-1}$ unless $t = t^k$ for some k and $a_{t-1} \neq \delta^k$ in which case he switches to $a_t = \delta^k$. In addition, $F(\emptyset) = (0,\ldots,0)$. Thus, for example, $F(\{(2,1),(4,0),(6,0)\}) = (1,1,1,0,0,0,0)$.

This modified story builder is rationalized by a preference relation defined by $x \succsim y$ if $P(x) \geq P(y)$ where $P(z_1,z_2,\ldots,z_T) = \sum_{t=2,\ldots,T} 1_{\{z_t = z_{t-1}\}} \lambda^t$ for $\lambda < 1$ close to 1. That is, P

evaluates a story $(z_1,\ldots,z_T)$ by adding the weight λ^t (almost equal to 1) for every period t in which $z_t = z_{t-1}$.

5. Step-by-step story builders

The story builders in the following examples construct a story in steps, starting from O and proceeding from one event to one of its immediate successors, by applying some "local" rule.

a. Recursive Construction I: Progress along the most likely story line as long as you don't encounter a contradiction.

The story builder has in mind a likelihood relation $\trianglerighteq$ (a strict ordering) over all stories (that is, $s \trianglerighteq s'$ means that s is more likely than s'). Given an evidence set E, let $C(E)$ be the set of all partial stories that do not contradict E, that is, those that can be developed into a story consistent with E. Formally, $C(E) = \{(O, a_1,\ldots,a_k) \mid \text{there is } s \in S_E \text{ that starts with } (O, a_1,\ldots,a_k)\}$. By the definition of an evidence set, $C(E)$ contains (O). The story is built recursively so that at each stage the story builder has a full story in mind, one that is not necessarily consistent with the evidence. He starts the construction having in mind the most likely story in S. He arrives at stage $k+1$ with a story s such that $(s_0,\ldots,s_k) \in C(E)$. He keeps the story if $(s_0,\ldots,s_k,s_{k+1}) \in C(E)$. Otherwise, he replaces s with a story t, the most likely story from those that start with $t_0 = s_0,\ldots,t_k = s_k$ and for which $(t_0,\ldots,t_k,t_{k+1}) \in C(E)$. He stops when he reaches a terminal event.

In other words, the story builder starts with the most likely story without taking into account the evidence. He advances along this story until he reaches a terminal event or until the partial story, say of length $k+1$, contradicts the evidence he possesses. At this point, he does not abandon the partial story of length k, which did not contradict the evidence. Rather, he considers all stories that start with the partial story of length k and for which the partial story of length $k+1$ does not contradict the evidence. *Note, this set might contain stories that are not consistent with the evidence, although the inconsistency is not apparent from their first $k+1$ events.* He chooses the most likely story in this set and continues with this "revised story" in mind until he reaches a terminal event or until he needs to revise the story again by applying the same procedure.

To demonstrate the psychology behind this procedure, consider a person who is known to have the following daily routine: (home, cafe, office, gym, return home). One day, the person is observed stopping at a bank and it is known that he could only have gotten there from either the cafe or the gym. According to the above procedure, the story builder now believes that the routine is: (home, cafe, office, gym, **bank**, return home) rather than (home, cafe, **bank**, office, gym, return home) since he revises the story he has in mind only when he realizes that it does not have a consistent continuation.

Consider Figure 3 and assume $(O,b,c,d) \triangleright (O,a,c,e) \triangleright s$ for any other $s \in S$. In the absence of any evidence, the story builder chooses (O,b,c,d). Given the evidence set $\{e\}$, the story builder starts with the most likely story, i.e. (O,b,c,d). He keeps it in mind at steps 1 and 2 since $(O,b) \in C\{e\})$ and $(O,b,c) \in C(\{e\})$. In the third step, he realizes that $(O,b,c,d) \notin C(\{e\})$ and appends the event e to (O,b,c), thus ending up with the story (o,b,c,e). Notice that $\trianglerighteq$ does not rationalize this story builder since $(O,a,c,d) \triangleright (O,b,c,e)$ and both are in $S_{\{e\}}$.

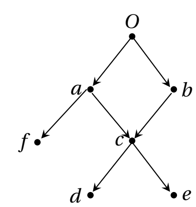

Figure 3

The story builder who follows this procedure clearly satisfies Story Consistency. Although he is not necessarily rationalized by $\trianglerighteq$, the following claim states that the story builder can be rationalized, regardless of the graph.

Claim 5: Let G be a graph. If F is a story builder that follows "Recursive Construction I" with the likelihood relation $\trianglerighteq$, then F is rationalizable.

Proof: Recall that $l(x)$ is the length of the story x and $j(x,y)$ is the length of the path from O to a_{xy} (the event where x and y first split). Thus, $(O,x_1,...,x_{j(x,y)}) = (O,y_1,...,y_{j(x,y)})$. Define $x = (O,x_1,...,x_{l(x)}) \succsim y = (O,y_1,...,y_{l(y)})$ if there is a story s starting with the partial story $(O,x_1,...,x_{j(x,y)+1})$ such that $s \trianglerighteq s'$ for all s' starting with $(O,y_1,...,y_{j(x,y)+1})$. In other words, the comparison between any two stories x and y is carried out using a criterion that involves the immediate successors following the split at a_{xy}. Each successive event b receives an evaluation $v(b)$, which is the maximum likelihood of a story that proceeds from O to a_{xy} and from a_{xy} to b. The story x, in which b_x follows a_{xy}, is determined to be $\succsim$-superior to the story y, in which b_y follows a_{xy}, if $v(b_x) \geq v(b_y)$.

This relation is complete and antisymmetric. To verify transitivity, assume that $x \succsim y$ and $y \succsim z$. Let s be a story starting with $(O,x_1,\ldots,x_{j(x,y)+1})$ and is $\trianglerighteq$-preferred to any story starting with $(O, y_1, \ldots, y_{j(x,y)+1})$ and let t be a story starting with $(O, y_1, \ldots, y_{j(y,z)+1})$ and is $\trianglerighteq$-preferred to any story starting with $(O, z_1, \ldots, z_{j(y,z)+1})$.

By definition, it is impossible that $j(x,y) = j(y,z) > j(x,z)$. It is also impossible that $j(x,y) = j(y,z) < j(x,z)$ since then s is also consistent with $(O, z_1, \ldots, z_{j(x,y)+1}) = (O, x_1, \ldots, x_{j(x,y)+1})$ and therefore $z \succsim y$. We are left with three cases:

Case (i): $j(x,y) < j(y,z)$. Then, $j(x,y) = j(x,z) < j(y,z)$ and $(O, y_1, \ldots, y_{j(x,y)+1}) = (O, z_1, \ldots, z_{j(x,y)+1})$ and s is also $\trianglerighteq$-preferred to any story starting with $(O, z_1, \ldots, z_{j(y,z)+1})$. Therefore, $x \succsim z$.

Case (ii): $j(x,y) > j(y,z)$. Then, $j(x,z) = j(y,z) < j(x,y)$ and the story t is also consistent with $(O, x_1, \ldots, x_{j(x,y)+1})$ and thus $x \succsim z$.

Case (iii): $j(x,y) = j(y,z) = j(x,z)$. The story s is $\trianglerighteq$-preferred to any story starting with $(O, y_1, \ldots, y_{j(x,y)+1})$ and, in particular, $s \trianglerighteq t$. Therefore, s is $\trianglerighteq$-preferred to any story starting with $(O, z_1, \ldots, z_{j(y,z)+1})$ and therefore $x \succsim z$.

Finally, we verify that $F = F_{\succsim}$. Let $s = F(E)$ and $t = F_{\succsim}(E)$. By the definition of $j(s,t)$, $(O, t_1, \ldots, t_{j(s,t)}) = (O, s_1, \ldots, s_{j(s,t)})$ and $s_{j(s,t)+1} \neq t_{j(s,t)+1}$. Then $t \succsim x$ for any story x consistent with E and, in particular, $t \succsim s$. By the definition of $\succsim$, there is a story starting with $(O, t_1, \ldots, t_{j(s,t)+1})$ that is $\trianglerighteq$-preferred to any story starting with $(O, t_1, \ldots, t_k, s_{j(s,t)+1})$, a contradiction of the definition of $s_{j(s,t)+1}$. ■

b. Recursive Construction II: Advance to the most likely event given the evidence and the partial story. The story builder has in mind a probability measure p on S. For an evidence set E, let p_E be the conditional of p on E. The story builder constructs a story recursively. At stage 0, he sets $s_0 = O$. He arrives at stage $k+1$ after constructing a path $(s_0, s_1, ..., s_k)$. He then appends to the partial story the event that, given the evidence and given $(s_0, s_1, ..., s_k)$, is the most likely follower of s_k. In other words, he chooses s_{k+1} to be a maximizer of the function $\phi(a) = p_E(\{t \in S \mid s \text{ starts with } (s_0, s_1, ..., s_k, a)\})$.

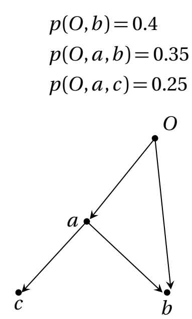

Figure 4

Such a story builder may not satisfy Story Consistency. For example, consider the graph and probability measure p depicted in Figure 4. In that case $F(\emptyset) = (O, a, b)$, that is, in the absence of any evidence, the story builder advances from O to a, which is a more likely successor than b, and then proceeds to b

since the path (O,a,b) is more likely than (O,a,c). However, if he receives the evidence $\{b\}$ the story builder advances from O to b since (O,b) is more likely than (O,a,b) and thus $F(\{b\})=(O,b)$.

c. Recursive Construction III: Advance to the most likely event given the evidence but independently of the partial story. The story builder has in mind a probability measure p on S. He builds the story recursively. At stage 0, he sets $s_0 = O$. He arrives at stage $k+1$ after constructing a partial path $(s_0, s_1, ..., s_k)$. He proceeds by adding the most likely event following s_k conditional on the evidence and ignoring the partial path he has constructed so far. In other words, he chooses s_{k+1} to be a maximizer of $\psi(a) = p_E(\{t \in S \mid a \text{ is in } t\})$. This procedure is similar to the previous one with except that when constructing the story recursively, the appended event s_{k+1} is selected independently of what the story builder believes to be the path that led to s_k. Note that the test for choosing s_{k+1} is based on assessing $p_E(\{t \in S \mid a \text{ is in } t\})$ rather than $p_E(\{t \in S \mid t \text{ contains } s_k \to a\})$.

Such a story builder F may not satisfy Story Consistency. Consider Figure 3 and assume that $p(O,a,f) = 0.2$, $p(O,a,c,e) = 0.35$ and $p(O,b,c,d) = 0.45$. In this case, $F(\emptyset) = (O,a,c,d)$. The story builder proceeds from the partial path (O,a,c) to d since it is a more frequent event than e, although the path (O,a,c,d) is a zero probability story. Given the evidence $\{a\}$, he reaches the conclusion that $F(\{a\}) = (O,a,c,e)$.

6. The story builder considers the source of the evidence

Story builders differ in how they relate to the way in which evidence has reached them. We distinguish between naive and non-naive approaches.

A Naive story builder: A naive story builder does not consider the reason that evidence has reached him. In other words, he does not make any conjectures that connect the true story to the evidence he receives.

A Non-naive story builder: A non-naive story builder makes a conjecture that connects the evidence he receives to and the true story. For example, he might have in mind a function μ that assigns to each story s a belief over the evidence sets that are consistent with s, where $\mu(s)(E)$ is the probability he assigns to receiving the evidence set E when the truth is s. Given an evidence set E, the story builder applies Bayesian reasoning to update his beliefs about s and then adopts the most likely story.

Story Consistency makes sense in the naive approach but is not intuitive when the story in the non-naive approach as the following example demonstrates. Consider, the boss of two employees A and B who has just arrived at work. In the absence of any information to the contrary he believes that employees A and B are already at work. If his secretary tells him that A is already at work, then he assigns significance to the fact that the secretary did not mention B and concludes that B has not arrived yet. In contrast, if the boss coincidentally passes by A's office and sees him there, then he probably will not change his mind as to whether B is already at work. Thus, if the evidence appears coincidentally, the boss' reasoning satisfies Story Consistency. But if the evidence is given in the course of conversation, the boss attributes importance not only to the evidence but also to the evidence that he does not have, thus violating Story Consistency.

A story builder who adopts a non-naive approach may not even satisfy the invariance property. Consider the graph in Figure 5. Assume that the story builder believes that the evidence set includes the event b if and only if the true story ends at z_2. Then, for any $E \subseteq \{a,b\}$, $F(E) = (O,a,b,z_2)$ if $b \in E$ and $F(E) = (O,a,b,z_1)$ if $b \notin E$. This F does not satisfy the invariance property since $S_{\{a\}} = S_{\{b\}}$ while $F(\{a\}) \neq F(\{b\})$.

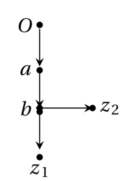

Figure 5

A strategic non-naive story builder: A story builder might be strategic, meaning that he views the evidence as originating from a source with a vested interest in the conclusion that the story builder will come to. We say that the story builder is *strategic* if:
(i) He believes that a source provided the evidence and that the source holds some preferences with respect to the story that the story builder will construct.
(ii) The source is not necessarily able to present the full true story. If he can present a particular evidence set then he is also able to present any subset of it.
(iii) Given the evidence set presented by the source, there is no subset of the evidence set that leads the story builder to believe in a story that is preferred by the source to the one that the story builder constructs.

Formally, F is strategic if there is a preference relation $\succsim_s$ on S such that $F(E) \succsim_s F(E')$ for any $E' \subset E$. The following claim states that being strategic is equivalent to being rational.

Claim 6: F is rational if and only if F is strategic.

Proof: If F is rational, then there is an ordering $\succsim_l$ such that if $F(E) = s$ then $s \succsim_l s'$ for all $s' \in S_E - \{s\}$. Define $\succsim_s$ by $x \succsim_s y$ iff $y \succsim_l x$. If $E' \subset E$, then $F(E') \succsim_l F(E)$, and $F(E) \succsim_s F(E')$.

In the other direction: Assume F is strategic and that it is backed by the preference relation $\succsim_s$. Let $\succsim_l$ be the inverse preferences. Let E be an evidence set and $F(E) = s$, and $s' \in S_E$. Let $E_{s'}$ be the set of all events in the story s'. Then, $F(E_{s'}) = s'$. The evidence set E is a subset of $E_{s'}$ and $s' \succ_s s$ since F is strategic, which implies that $s \succ_l s'$. Thus, F is rational. ■

We use the term "rational story builder" since it is consistent with the definition of a story builder as a rational choice function. However, the term is somewhat misleading. A strategic story builder is required to hold beliefs about the source's preferences which are independent of the true story. A story builder who would like to know the truth and believes that the source of the evidence wants him to construct the true story is not rational by this definition. This is analogous to a choice function not being rational if the decision maker holds preferences that depend on the choice set from which he has to choose.

7. A Robust Failure of coordination between different types of story builders

We conclude the paper by demonstrating the potential use of the story builder concept using the following simple example of an interaction between different types of story builders: A "large" number of individuals are randomly pair-wise matched and play a coordination game in which each tells a story represented by a path in the graph shown in Figure 6.

Before telling a story, each individual builds a point-wise belief (a story) about his partner's choice. Assume that no individual has any evidence of his partner's actual choice. The population is partitioned into three types of story builders as described below. All individuals use the same distribution of stories p as a parameter in their belief formation procedure. In equilibrium, this distribution is the actual distribution of stories in the population.

Following are the three types of story builders:

Type 1 chooses the most likely story given p.

Type 2 uses the Recursive II procedure (building the story by advancing along the graph in steps, each time moving to the most likely next event, given p and given the path he has chosen so far).

Type 3 uses the Recursive III procedure (i.e. choosing the most likely next event independently of the path he has chosen so far).

Denote the frequencies of these three types in the population by α, β and γ, respectively. Assume that $\beta + \gamma > \alpha > \delta > \gamma$.

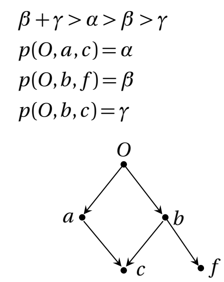

Figure 6

A candidate for equilibrium is a profile $(s_1,\ s_2,\ s_3)$ where the story s_i is chosen by the individuals of type i. Since each individual wishes to coordinate with his matched partner, we require that in equilibrium s_i be identical to the story built by a type i story builder about his partner, given p. In addition, we require that p be the distribution of stories in the population derived from $(s_1,\ s_2,\ s_3)$ and the distribution of types (α, β, γ).

Under these assumptions, $s_1 = (O, a, c)$, $s_2 = (O, b, f)$, $s_3 = (O, b, c)$ is an equilibrium. Type 1 finds (O, a, c) to be the most likely story. Type 2 chooses s_2 since the event b is the most likely event that follows O ($\beta + \gamma > \alpha$) and f is the most likely event given the path (O, b) ($\beta > \gamma$). Type 3 chooses s_3 since the event b is more likely than a ($\beta + \gamma > \alpha$) and the event c is more likely than f ($\alpha + \gamma > \beta$).

In this equilibrium, when two individuals of different types are matched they will fail to coordinate. This failure differs from that in a non-degenerate mixed strategy Nash equilibrium in a standard coordination game. In such an equilibrium, each player is indifferent between several alternatives. The mixed strategy equilibrium is not robust in the sense that a small change in the action distribution will push the players away from equilibrium. In our case, the behavior of each type is robust to "slight changes" in the distribution of stories in the population.

Note that such a robust equilibrium with coordination failure would not exist if all individuals were of the same type or if only two of the three types (with nonidentical proportions) existed in the population.

One can view the story builder in this example as a problem solver in the sense of Glazer and Rubinstein (2018) whose task is to come up with a story that will match his partner's choice. In this respect, the model is one of interactions between different problem solvers.

References

Bjorke Isaac (2019). "Misleading Information", mimeo.

Eliaz, Kfir and Ran Spiegler (2018). "A Model of Competing Narratives", mimeo.

Glazer, Jacob and Ariel Rubinstein (2019). "Coordinating with a Problem Solver". *Management Science 65* (2019), 2813-2819.

Sadler Evan (2018). "False Information and Disagreement in Social Networks", mimeo.