Behavioral Economists, Human-Computer Interactions and Research Transparency*

Kweku A. Opoku-Agyemang†

August 11, 2017

Click Here For The Latest Version

Abstract

I model economists as being themselves behavioral, and not fully rational in how they implement their research. I characterize research results as a network graph of data analyses at the level of human-computer interactions that contain graph copies whenever the economist manifests behavioral shortcomings. Ethical and transparent research emerges as a unique graph. I allow researchers to hold themselves accountable with commitment devices that eliminate behavioral copies of research graphs. In so doing, I introduce a solution to the problem of the so-called garden of forking paths. Applications are based in economic theory and qualitative research for program evaluation design. Transparency may scale if researchers become more aware of software commitment devices that enable creativity via preanalysis plans, open data, dynamic documents and time-stamped analyses. Human-computer interaction has a fundamental role to play in research transparency. (*JEL* A14)

Keywords: Sociology of economics, psychology and economics, research transparency, human-computer interaction.

*This research was funded by the National Science Foundation under Grant Number IIS-1054332. I thank Temina Madon, Ted Miguel, Kuranda Morgan, Kelsey Mulcahy and Alex Wais for various helpful conversations. The perspective and implications of the paper are solely and fully mine and not to be thought of as the responsibility of, or as being approved by any entity or person other than the author. The usual disclaimer applies.

†Cornell Tech, Human-Computer Interaction and Social Computing Research Area, New York, NY 10011, United States. Email: `kao75@cornell.edu`. Web: `www.kwekuopokuagyemang.com`.

1 Introduction

Research tansparency and ethics processes aimed at improving reliability and credibility are gaining momentum in the economics profession. In this paper, I model research objectivity under the constraint that most economists are *behavioral* in pursuing research. The goal of the paper is to help direct the diverse implementations of economics practice towards actions that are transparent, by focusing on the interfaces between economists and their computers[1].

For example, an empirical economist doing a natural experiment on observational data continuously decides how analyze the data when new insights emerge with her software tools. She pivots her theoretical model and refines her research question as her priors co-evolve. An applied theorist must chip away at her theorems and propositions while her knowledge of the topic and its context is evolving. An theoretical econometrician may occasionally worry that her propositions and lemmas exhibit too little overlap with empirical data and its limits, and can use her intuition to control his research path in near-real-time. Various professional incentives interact with research choices and may significantly complicate research results to become non-replicable (Merton, 1973). As such, research may begin with the best of intentions but end up being non-transparent. In practice, such factors collectively mean that even if the trajectory of many research decisions begin objectively, they may later mimic elements of randomness as behavioral problems manifest.

A common theme in all the above applications is that in the current regimes, much "rough" work goes unreported, and such efforts may help transparency efforts. As such, any commitments which impose a transparent structure on research must necessarily be coupled with random movements by researchers that ultimately harm the integrity of research. In the paper, I argue that *all* economists, regardless of field, are "behavioral" in this sense. For

[1]The human-computer interaction field is at the nexus of computer and information science on the one hand, and the social and behavioral sciences and related fields such as design on the other. Using a computer (for research in my case, although this is obviously unrepresentative) is considered to take place as an open-ended dialogue between the user and the computer. Human-computer interaction is viewed through the lens of human-to-human interaction, an analogy which is convenient for economic thoretical modeling. The term was mainstreamed by the Card, Newell, and Moran (1983) classic text, *The Psychology of Human-Computer Interaction*. The term can be traced, however, to slightly earlier works (see Carlisle, 1975; Card, Newell and Moran, 1980). A recent insightful overview for an economist audience is contained in Aker, Burrell and Ghosh (2009).

example, if the political economist can be self-committed to ethics in a way that credibly allows the researcher to explore the data in a transparent way, then in the absence of such a commitment, the reader may infer that the research is significantly diminished in integrity. Such an self-commitment approach must build on preanalysis plans. Preanalysis plans add a credilibity dimension to a dichotomy between preregistered results and exploratory findings: it makes preregistered results credible but also negatively alters the credibility of non-registered reported results. There is a trade-off between using a preanalysis plan for transparency and the credibility attached to reported exploratory work, which has a unique role to play in economics. This belief may be one reason why so few scholars use traditional preanalysis plans.

I develop a general model to motivate dynamic preanalysis plans that can accomodate observational data, while recording all human-computer interactions associated with a paper. I argue that such recordings should include code that would normally be unreported for the sake of brevity or claimed irrelevance, as they might provide important clues to improving research transparency. Formally, the underlying model is an extension of the time-inconsistent network model of Kleinberg and Oren (2014) to incorporate random graphs. A researcher generates an authentic graph with several random graph copies. The agent is behavioral, first taking research actions to generate an ethical graph and then taking actions to create copies of the graph that are not transparent when her behavioral shortcomings kick in. Relative to the time-inconsistent node-level model of Kleinberg and Oren (2014), graph-level analyses add several interesting dimensions to understanding the sociology of economics research for publishers and other actors. The main implication is the generalization of preanalysis plans to timestamp every data analysis decison taken, the use of using dynamic documents and sharing all recorded human-computer interactions on publication. Specifically, even after the time a paper is written and published, with a updatable and time-stamped preanalysis plan that shows the universe of research decisions, an objective and transparent account of the research findings can be backed out. If the transparent graph differs from the publication, it may be presented as a substitutive or complementary narrative. By lessening the stigma associated with non-transparency, such an approach may encourage more economists to engage with research transparency.

To illustrate these ideas, consider the following example. An economist is concerned about being transparent while doing research on publically-available observational data, and is considering how to configure a preanalysis plan to integrate exploratory research in a way that minimizes unethical behavior. Currently, the preanalysis plan is configured to contain a predictable research graph, giving the researcher an irresistible commitment to the preregistered plan.

The researcher considers using the preanalysis plan throughout for the observational data analysis. However, the information available to the researcher is not perfect: she comes up with newer ideas to test in the data. Although the preanalysis plan allows her to report these additional findings as exploratory, and not preregistered as is appropriate, this trade-off between exploration and credibility motivates the economist to decide to throw out the preanalysis plan entirely. This idea holds because it will be sufficiently likely that when *all* results are exploratory, they will be given equal weight by the audience. The problem is to incentivize her to use a preanalysis plan that enables exploratory research to be transparent, so as to minimize the current trade-off.

I now consider each decision as independent graph copies. Let $G_{i\dots k}$ be independent copies of $\mathbb{G}_{v,e}$, where by each copy graph is scaled by a factor $\beta\left(G_i, \dots G_k\right)$ that discounts how credible we should consider the output of each independent graph. I do not impose any particular structure on the discount factor. If the researcher is rational and transparent, there are no independent copies of $\mathbb{G}$ or discount factor at all; only the graph $\mathbb{G}$. By contrast, when the researcher is behavioral, there exist independent copies of $\mathbb{G}$, in addition to the transparent research graph. All copies reproduce the graph structure. There are four kinds of copies that are relevant from an empirical graph programming perspective, and all are accounted for in the approach:

1. scenarios where the graph structure as well as all data attributes and any objects they might contain are copied. The entire graph object is new so that changes in the copy do not affect the original object.
2. scenarios where the graph structure is copied but the edge, node and graph attributes

are references to those in the original graph. This saves time but could cause confusion if the researcher changes an attribute in one graph and it changes the attribute in the other.

3. scenarios where any attributes that are containers are shared between the new graph and the original
4. scenarios where the resulting graph is independent of the original

The approach of isolating the transparent graph from non-transparent graphs is without loss of generality, although it is simplest to think of the graphs as being independent of the original. I relax this assumption in Section 3.

To drive home the simple approach, consider a simple graph where triangles represent non-transparent behavioral research graphs. These may be generated as follows: from two edges starting from the same node, with one edge being rational and the other behavioral. Both edges meet in a single node to form a triangle. Assume that the total number of triangles in the finite graph (i.e. the time-stamped code of the paper) are defined. Then we may define a graph *without* triangles by removing triangles, and that such a graph represents a way to identify credible parts of a research project. For an illustration, let V_1, V_2, V_3 refer to sets of vertices. Let $G = (V_1, V_2, V_3, E_{12}, E_{23}, E_{31})$be a triangle on these sets of vertices. Suppose that G contains at most $\delta\,|V_1|\,|V_2|\,|V_3|$ triangles. Then there exists a graph

$$G' = (V_1, V_2, V_3, E'_{12}, E'_{23}, E'_{31})$$

which contains no triangles whatsoever, and such that $\left|E_{ij} \setminus E'_{ij}\right| = o_{\to 0}\left(|V_i \times V_j|\right)$ for $ij = 12, 23, 31$.

In terms of the model used in the paper, the economist is the decision maker, and the nodes are subsets of the research. There is an uncertain state of the world (since the researcher can make progress on any part of the research at a point in time, even by just thinking about it) that is stochastically changing over time and observed only by the researcher. The problem of the researcher is how to maintain an objective research trajectory and drop the behavioral research trajectory, here, how to keep the objective research graph and drop the behavioral graph. Among the feasible policies for managing research, using a preanalysis plan to preregister and direct experimental research, and doing strictly exploratory research on theoretical or observational data are the two extremes.

Between these are a relevant range of research policies. For example, consider research on nonexperimental analyses (such as regression discontinuity designs or natural experiments), or simulations or theory. Such research often relies on exploratory research decisions as knowledge accumulates in a random process. Observe that if the research graph process is overly predictable, the associated results might be credibly identified as p-hacking, specification fishing or other antisocial research behavior. I consider commitment to objective reesearch graphs to be denoted by the dropping of exploratory graphs by a commitment device. I think of the commitment device manifesting in a Spencian signaling manner, so that the device responds to every exploratory graph.

Characterizing the qualitative decisions that economists make during the long gestation research periods common in the profession, and how they can be made more transparent is complicated. In this paper, by extending the computing discussion of Kleinberg and Oren (2014) to discussions on research transparency in economics (summarized in Christiansen and Miguel 2017), I develop a tractable characterization of research transparency problems in economics research and a method for maintaining research integrity using a wide class of available software and other solutions which retain the benefit of exploratory research by making it more credible.

When I apply these techniques to the special case of microeconomic theoretical modeling (broadly based on Varian (2016) tutorial essay, *How To Build An Economic Model in Your Spare Time*. I find a simple qualitative prescription for research management in economic modeling. First, theorists may also use dynamic documents to preregister their conceptual frameworks, and adapt them according to their evolving preferences. The elements which do not appear in the final work may be dropped via a commitment device. In this setting, homeomorphisms that reflect behavioral problems are removed, in a stronger version of the graph-based main argument. Again, the commitment device acts as a signaling tool to researchers. The shortcoming of this approach is that it inherently equates an absence of commitment with a disinterest in transparency. I exploit the fact that dynamic document tools are free downloads, but costly in terms of the effort required to create them, so that some researchers use them and others do not. A separating equilibrium explains which homeomorphisms are dropped.

Qualitative research can also benefit from research transparency. Such tools are often used in applied fields such as development economics or comparative political economy while evaluating policy programs. Although such research is often unreported in development economics research, their presentation in supplementary documentation can enable such research to be more transparent.

1.1 Related Literature

The model studied in this paper is a network-level extension of the time-inconsistency graph edge-based model of Kleinberg and Oren (2014). It adapts results from that paper to characterize behavioral research network graphs and isolate objective elements of research from research that contains non-transparent parts and involve large decision-making data streams. The key insight of Kleinberg and Oren (2014), that some research edges are time-inconsistent and others are rational, were in turn, adapted from quasi-hyperbolic discounted utility modeling of Laibson (1997).

Concurrently and independently of this work (Fafchamps and Labonne, 2017) and (Anderson and Macgruder, 2017) evaluate preanalysis plans, but with a different focus. Both papers are based on splitting samples into a training dataset and a test dataset. Fafchamps and Labonne (2017) are interested in scholars incorporating feedback prior to publication. Once the paper is accepted for publication the method is applied to the testing sample so that the latter results are exclusively published. Anderson and Macgruder (2017) propose conducting exploratory analysis in the fraction of the data not withheld as well then registering a simple analysis plan to assess the approach. Instead, I argue for removing the dichotomy between exploratory and confirmatory research, since most research and software innovations enable researchers to use dynamic documents and timestamp all research decisions. As such, I take a more "micro" look at research decisions made by economists and other researchers. Another computer science field, human-computer interaction, is relatively relevant for this discussion on how researchers interact with software by giving commands and queries, and this paper focuses closer to this level of analysis than on individual economists.

Network formation models have counterparts in the rich literature, which the paper benefits from and builds upon. For example, Chandrasekhar and Jackson (2014) show when random graph model parameters are consistent, while reinterpreting network formation as a distribution over the space of sufficient statistics to enable empirical work. In a spinoff paper, Chandrasekhar and Jackson, (2016) then generate subgraphs such as links and triangles and estimate the rates at which they are formed. Mele (2010) discuss how agent types are segregated in networks. Bedev (2013) presents a solution concept that subsumes Nash to analyze agent choices within the context of endogenous graphs and discrete games. Graham (2015) provides approaches for identification in networks. The paper departs from this literature by exploiting triangle and graph removal, an important approach in mathematical graph theory (Szemerédi, (1975, 1978), Fox, 2011). This approach has played a significant role in computer science theories and applications such as graph property testing (Rubinfield and Sudan 1996; Goldreich, Goldwasser and Ron, 1998). I model the economic decision-maker as an agent that generates graphs that define her research. In doing so, she may signal demand for transparency by using a commitment device that responds by eliminating behavioral graphs and converts exploratory research to confirmatory investigation.

There is an fast-paced but mostly relatively recent literature on research transparency in economics. Much of the attention in the profession on preanalysis plans stems from Casey et al (2012), which showed how their absence could yield wildly conflicting results from the same dataset. Overviews of the need for transparency are contained in several recent articles(e.g. Miguel et al, 2014; Olken, 2015; Chang and Li (2017); see Christiansen and Miguel, 2017 for the state of the literature). Another policy approach has focused on replications, with an investigation performed by Camerer et al (2016), with a growing number of ovewviews[2]. Coffman and Niedderle (2015) consider preanalysis plans and replications as substitutes. This paper connects the entire research workflow from prior to preanalysis plans through

[2]See Berry et al (2017), Sukhtankar (2017), Hamermesh (2017), Coffman et al (2017), Duvendack et al (2017), Höffler (2017), Anderson and Kichka (2017).

replications, considering the two as complements to provides an avenue for the two tools to be integrated into a single graph. In so doing, the methods in this paper may help foster the adoption of preanalysis plans and replications alike to minimize transparency problems.

My discussion merges transparency with improvisation. A seminal text of scientific improvisation is a classic text by Flexner (1939), titled *The Usefulness of Useless Knowledge*. The article argues for enabling spontaneity in research. The proposal of Flexner was to allow researchers to pursue science without worrying about empirical application, given the difficulties of predicting the social utility of research ex-ante. Although it was mainly written for an audience of physical scientists, I find the incentive structures it proposes to be relevant for my discussion research transparency. As such, I think of physical scientists as being "behavioral economists" in how their research graphs are generated. Like economists, scientists face a variety of incentives that stymie this approach, and modifications to commitment devices may help such goals to be more feasible.

Selection for statistical significance is not only reflected in the issue of not reporting unappealing results (called the "file drawer" problem) and mining datasets for acceptable p-values (called p-hacking). The researcher *herself* is dealing with degrees of freedom and navigating a so-called garden of forking paths while doing research (Simmons, Nelson, and Simonsohn, 2011; Gelman and Loken, 2013), so that undisclosed flexibility can lead to exaggerations of statistical significance. It is less clear how to approach this problem, although reporting data collection, annotation and analysis may help in the sciences (Plant et al, 2014). Other complementary approaches for the sciences may use grants to assess "pathways to reproducibility" along with "pathways to impact" (e.g. Sené, Gilmore and Janssen, 2017) after a paper is completed. Software solutions have the potential to make such steps relatively efficient. Although measurement has been important for answering economic questions (Deaton, 2016), the gap in the discussion stems at least partly from inadequate economist self-modeling of measurement processes in the transparency space.

In this paper, I attempt to study literal economist decision-making, using a model to understanding the scale and scope of behavioral problems. The paper relies on human-computer interactions for research transparency. Human-computer interaction (HCI) is traditionally based in information schools and computer science and has developed a science around how people interact with computer interfaces (see Dix, 2009 for an overview). This allows the paper to model how research is done and how it may be pursued transparently. Given the prevalence of computers in both theoretical and applied research in economics, it is surprising that such an engagement is absent in the literature (to the best of my knowledge). Exceptions that focus on how economic agents interact with computers include Aker, Burrell and Ghosh (2009), Opoku-Agyemang, Shah and Parikh (2017), and Opoku-Agyemang (2017). Modeling how economists *themselves* interact with computers is a contribution of the paper.

2 Model

The basic notation follows Kleinberg and Oren (2014), applied to the context of research transparency and hence extended to focus on graph copies. The results in a research paper is summarized by a graph G, with each node v presenting output generated from some command or action taken by a researcher on their computer. I consider the movement from action to action to be seperated by an edge e. The main idea is that by the time a paper is written with all of its robustness checks, multiple analyses and sub-analyses motivated by seminar audiences, conference discussants, journal reviewers and editors, it consists of multiple research trajectories. I assume that only one of these trajectories is objectively authentic to the research question at hand (I use the words objective, transparent and authentic interchangeably). Each research trajectory consists of multiple graph "copies" of an objective research graph. I assume that all research has results that can be objectively identified: there is no such thing as a research question whose answer is undefined[3]. In this section, I discuss how on identifying the objective graph. I also use the following convenient standard $o\left(\right)$ and $O\left(\right)$notation conventions from the computer science graph literature. Let $o_{x\to 0;y_1\ldots,y_n}\left(X\right)$ describe any quantity bounded in magnitude by $Xf\left(x,y_1,\ldots,y_n\right)$so that

[3]This is only a harmless simplifying assumption, as an undefined answer may be considered an objective one, without loss of generality.

f is a function and $f \to 0$ as $x \to 0$ for each fixed choice of $y_i, \ldots y_h$. For a function $F()$of $y_1, \ldots y_n$, I let $O_{y_1, \ldots y_n}(X)$refer to a quantity bounded by $XF(y_1, \ldots y_n)$. I also let F be finite so that $|F|$refers to its cardinality.

Behavioral Economist without Commitment Device. In the model, the economist has behavioral problems in navigating research path R. The economist's research trajectory is summarized as:

$$d(u, v) = \min_{R \in R(u,v)} \sum_{e \in R} c(e)$$

whereby the research path is defined by R, the start node is a, and the end node is b. Edges are $e_i \in E$, and cost is c. The economist starts with edge e_1, but the time-discounting sets in at the rate of $\beta = [0, 1]$. In each step, the agent at node u chooses an out-neighbor v that minimizes the time-inconsistency equation $c(u, v) + \beta d(v)$. This means that the economist always starts from node u in one direction in the short term and can end up going in a discounted direction in the long term due to behavioral problems. The economist cares less about the future for smaller β. If R is the $a \to b$ path chosen by the researcher, the cost ratio is given as the actual cost divided by the optimal cost, $\sum_{e \in R} \frac{c(e)}{d(a)}$.

Commitment device: A commitment device removes parts of the graph, so as to reduce the set of options available to the researcher. The researcher sends a signal associated with specific behavioral problems to the commitment device, which provides the best response by dropping those parts of a graph generated by behavioral problems. The graph that remains after this process is,by definition, the objective research that is most credible. Recall that such graphs have no discounting. If this graph is found to fall short of answering the research question, the researcher or research replicator may build on the available information to pursue the objective research trajectory.

Dropping Behavioral Graph Elements with a Commitment Device. Let the researcher take a mix of objective and behavioral paths in the research. Behavioral paths emerge when $i > i^*$. Theeconomist research decisions is summarized y $c(e_{i<i^*}(R)) + \beta \sum_{i \geq i^*} c(e_i(R))$.

Behavioral research graphs.

A simple way to identify non-transparency in research is through a sub-graph pattern such as triangles, as done in the introduction. Consider two edges starting from the same node, with one edge being rational and the other behavioral. Since a triangle has three sides and three edges, it is straightforward to see that an objective edge and two behavioral edges (or two objective edges and one behavioral path) start from the same one node and end in another node for both rational and behavioral edges. In the simple case, such triangles represent a *clustering* in the graph that denote behavioral shortcomings. These may be quantified with standard clustering coefficients.

The global clustering coefficient based on the number of triangles in a graph G are given as $C = \frac{\sum_j triangle_j}{\sum_j \binom{degree_j}{2}}$, where the degree of node j is $degree_j$ and $triangle_j$ refers to the number of triangles incident to node j. If a wedge is a research path of length 2, then the number of wedges centered at node j is $\binom{degree_j}{2}$. The clustering coefficient based on the number of triangles in a graph are $\frac{triangle_j}{\binom{degree_j}{2}}$ so that the global coefficient sums both the numerator and denominator[4].

In cases where the graph is non-binary, the weighted version of the global clustering coefficient is $C_w = \frac{\sum_{\tau_\triangle} \omega}{\sum_{\tau} \omega}$ where the numerator refers to the total number of closed triplets and the denominator refers to the total number of triplets[5].

In the model, a commitment device always drops all triangles from a graph. This is known as becoming triangle-free.

Triangle-free: A graph G is triangle-free if it contains no triangles. This occurs through triangle-removal, or commitment.

Network density: The density tracks the relative fraction of links present.

Commitment: commitment is the mechanism of triangle-removal, which occurs as a signal from the researcher to software, via adoption $\theta = \{0, 1\}$. There is a one-to-one mapping between the nodes, edges, graphs of the researcher to software functions. The commitment device responds in a Spencian manner with a best-response function $\Gamma(\theta)$. If $\theta = 1$, all triangles are removed. Both the signal and the best response are subject to the

[4]See Durak et al (2012) for a treatment.

[5]Opsahl and Panzarasa (2009).

network structure, so that signals are not sent when there are no triangles, for example. I first have to count the number of triangles and graphs available.

Counting Lemma. The counting lemma (Lee, 2015) denotes the number of graph copies with the following lemma:

Lemma: For every graph H and $\delta > 0$, there exists c and ε_0 such that for all $\varepsilon \leq \varepsilon_0$:

1. Let G be a graph and let $\prod$ be a vertex partition of G. If there exists a copy of H in $R(\prod)$ then there exists a copy of H in G.
2. If the copy of H is over the vertices $1 \ldots r$,

Then there are at least $c \prod_{i=1}^{r} |V_i|$copies of H in G. I now discuss the removal of triangles from a graph.

Triangle removal commitment lemma. Lemma: Let $\theta = 1$. For every $\varepsilon > 0$, there exists $\delta > 0$ such that for any graph G, on n vertices with at most δn^3triangles, G is made triangle-free by a commitment device responding to εn^2 commitment signals.

The lemma eliminates triangles from the research graph, so that only the rational choices remain in the research graph. The commitment device employs best responses to Spencian signals exerted by the researcher, subject to the underlying network structure. The commitment device drops the behavioral triangles.

Proof: Let $X_1 \cup \ldots \cup X_M$ be a $\frac{\varepsilon}{4}$ partition of G. Recall that $\theta = 1$, so the signal is sent by the researcher for edges to be removed. Both actions are subject to the constraint of the network structure. Edges are removed when

1. $(x, y) \in X_i \times X_j$ whereby (X_i, X_j) is not a $\frac{\varepsilon}{4}$-regular pair,
2. $(x, y) \in X_i \times X_j$ whereby $d(X_i, X_j)$ is $< \frac{\varepsilon}{2}$
3. $x \in X$ whereby $|X_i| \leq \frac{\varepsilon}{4M} n$.

From condition 1, the number of edges removed is at most,

$$\sum_{(i,j)} |X_i|\,|X_j| \leq \frac{\varepsilon}{4} n^2.$$

From condition 2, the number of edges removed is at most, $\frac{\varepsilon}{2} n^2$.

From condition 3, the number of edges removed is at most, $Mn\frac{\varepsilon}{4M} n = \frac{\varepsilon}{4} n^2$.

At most, εn^2 edges are removed.

I now end the proof by contradicting the existence of a triangle in the graph after these steps are taken. Assume that some triangle xyz remains in the network, such that $x \in X_i$ $y \in Y_i$ and $z \in Z_i$ This implies that the pairs $(X_i, X_j), (X_j, X_k)$ and (X_k, X_i) are all $\frac{\varepsilon}{4}$-regular with density of at least $\frac{\varepsilon}{2}$.

Since $\sum_{(i,j)} |X_i|\,|X_j|\,|X_k| \geq \frac{\varepsilon}{4M} n$, from the counting lemma" the number of triangles is at least $\left(1 - \frac{\varepsilon}{2}\right)\left(\frac{\varepsilon}{4}\right)^3 \left(\frac{\varepsilon}{4M}\right)^3 n^3$. Taking $\delta = \frac{\varepsilon^6}{2^{20} M^3}$ provides a contradiction of a remaining triangle. □

Graph-level discussion. Obviously, the above discussion does not give researchers enough credit: nearly any researcher would take more just than two actions before behavioral shortcomings set in. I aggregate the discussion from triangles to graphs, so that in the simple case of one graph and one copy, there is a full graph $G'$for when the researcher is not time-inconsistent (at time $t_{i<i*}$) and a separate graph G'' for when she is time-inconsistent (at time $t_{i \geq i*}$), we have the following format for the full research network $\mathbb{G}$, whereby

$$\mathbb{G}_{t_{i=1 \ldots i* \ldots}} = \left(G'_{v',e'}\right)_{t_{i<i*}} \cup \left(G''_{v'',e''}\right)_{t_{i \geq i*}}$$

This overall approach is assumed to be applicable in some research cases, where the research question context is sufficiently understood at u, so that the graph actually exists in advance. However, this approach would not be sufficient to explain much economics research. Perhaps most critically, it assumes that the graph is mapped out in an identical way for all researchers answering the same question. In some cases, however, the assumption that all economists would be anticipated to approach an identical research question in the same way would be too strong to be empirically founded. An approach which is a better fit should also allow that in practice, researchers may rarely approach the same question in the same way in the first place. For example, some may generate a theory out of empirical results, or take the reverse approach under the same data circumstances. The inherent subjectivity of certain decisions may not be adequately represented in the basic discussion.

I therefore introduce randomness into the research graph formation, while keeping the setting of time-discounting. The research graph is based on randomly-generated graphs that also involve discounting across nodes and/or networks, although I suppress the discount rate at the graph-level discussion for ease of exposition. Another benefit of this approach is that I can relatively cleanly discuss how multiple behavioral networks may emerge in more detail (the previous discussion only assumed one point splitting rational from behavioral (i.e. at the start)). Notwithstanding, the main ideas in the paper apply to both discussions without loss of generality and the distinction outlined above merely serves an expository purpose. I still assume that there are fixed number of ways of generating a paper.

Let$\mathscr{G}_{v,e}$ refer to the family of all labeled graphs such that we have a set of vertices $V = [v] = \{1, 2 \ldots v\}$ and e edges, so that $0 \leq e \leq \binom{e}{2}$. To every graph $G \in \mathscr{G}_{v,e}$, I assign a probability $C(G) = \binom{\binom{v}{2}}{e}^{-1}$. This is akin to starting with an empty graph of research on the set $[v]$ so that the researcher inserts e edges in such a way that all possible choices are equally likely, due to behavioral research problems that occur in practice. We denote the graph by

$$\mathbb{G}_{v,e} = ([v], S_{v,e})$$

Independent copies of graphs (definition).

Monotone increasing

Having a copy is a monotone increasing graph property of a research graph.

Having copies are monotone increasing p. 6 (definition). Lessening transparency.

The property of having copies, $\mathscr{C}$, is monotone increasing if

$$G \in \mathscr{C} \Longrightarrow G + e \in \mathscr{C}$$

. This means that adding an edge e to a graph G does not destroy the property. This is defined as lessening transparency.

Having copies are monotone decreasing (definition). Increasing transparency. A graph is monotone increasing if and only if its complement is monotone decreasing.

Monotone decreasing. The property of not having copies, $\mathscr{C}'$, is monotone decreasing if

$$G \in \mathscr{C}' \Longrightarrow G - e \in \mathscr{C}'$$

This means that removing an edge e from a graph G does not destroy the property. This is defined as *raising transparency.*

Proposition. If $\mathscr{C}$ is monotone increasing then, whenever $c < c'$,

$$\mathbb{C}\left(\mathbb{G}_{v,c} \notin \mathscr{C}\right) \leq \mathbb{C}\left(\mathbb{G}_{v,c'} \notin \mathscr{C}\right)$$

A Threshold as an "Equilibrium"

Thresholds in the research graph generate graph copies while the research is being done by the economist. Dropping graph copies also occurs via thresholds by the use of the commitment device.

Identifying behavioral graphs: Behavioral copies of graphs from quasi-exogenous thresholds.

I discuss threshold functions for the monotone graph property of having copies $\mathscr{C}$.

Threshold: A function $f^* = f^*(v)$ is a threshold for $\mathscr{C}$ in the graph $\mathbb{G}_{v,e}$ if, as $v \to \infty$,

$$\lim_{v \to \infty} \mathbb{C}\left(G_{v,e} \in \mathscr{C}\right) = \begin{cases} 1 & \text{if } \frac{f}{f^*} \to 0 \\ 0 & \text{if } \frac{f}{f^*} \to \infty \end{cases}$$

Definition. I say that some function of vertices d grows asymptotically at least as quickly as a constant times a different function of vertices b, denoted $d(v) \ll b(v)$ or $d(v) = o(g(v))$ if $\frac{d(v)}{b(v)} \longrightarrow 0$ as $v \longrightarrow \infty$.

Proposition: $\mathscr{C}$ has a threshold, c^*. (p. 10)

I start from the characteristic that $\mathscr{C}$ is a monotone increasing graph property. Assuming that $0 < \varepsilon < 1$, I define $c(\varepsilon)$ as $C\left(G_{v,c(\varepsilon)} \in C\right) = \varepsilon$. Note that $c(\varepsilon)$ exists from the fact that $\mathscr{C}$ is monotone increasing and increasing the likelihood that $\mathbb{G}_{v,e} \in \mathscr{C}$.

I now show that $c^* = c\left(\frac{1}{2}\right)$ is a threshold for $\mathscr{C}$. Let $G_{1,\ldots,k}$ be independent copies of $\mathbb{G}$. The graph defined by the union of the copies,

$$G_1 \cup \ldots \cup G_k$$

is distributed as $G_{v,1-(1-c)^k}$. Now $1-(1-c)^k \leq kc$. From the coupling argument[6]

$$\mathbb{G}_{v,1-(1-c)^k} \subseteq \mathbb{G}_{v,kc}$$

thus, $\mathbb{G}_{v,kc} \in \mathscr{C} \Longrightarrow G_{1\ldots k} \notin \mathscr{C}$. Hence

$$\mathbb{C}\left(\mathbb{G}_{v,kc} \notin \mathscr{C}\right) \leq \left[\mathbb{C}\left(\mathbb{G}_{v,kc} \notin \mathscr{C}\right)\right]^k$$

Let ω be a function of v such that $\omega \longrightarrow \infty$ arbitrarily slowly as $v \to \infty, \omega \ll \mathrm{loglog}v$. (Recall that $d(v) \ll b(v)$ or $d(v) = o(g(v))$ if $\frac{d(v)}{b(v)} \longrightarrow 0$ as $v \longrightarrow \infty$). Suppose that $c = c^* = c\left(\frac{1}{2}\right)$ and $k = \omega$. Then $\mathbb{C}\left(\mathbb{G}_{v,\omega c} \notin \mathscr{C}\right) \leq 2^{-\omega} = o(1)$.

On the other hand, for $c = \frac{c^*}{\omega}$,

$$\frac{1}{2} = \mathbb{C}\left(\mathbb{G}_{v,c^*} \notin \mathscr{C}\right) \leq \left[\mathbb{C}\left(\mathbb{G}_{v,\frac{c^*}{\omega}} \notin \mathscr{C}\right)\right]^{\omega}$$

As such,

$$\mathbb{C}\left(\mathbb{G}_{v,\frac{c^*}{\omega}} \notin \mathscr{C}\right) \geq 2^{\frac{-1}{\omega}} = 1 - o(1)$$

$\square$.

Utility function

A "seriousness" graph discount factor $\mathscr{B}$ is associated with each graph copy. The transparent network is generated from $t = i = 1 \ldots i^* - 1$. The threshold is activated at $t \geq i^*$ and

[6] In probability theory, coupling is a proof technique that allows one to compare two unrelated random variables (distributions) X and Y by creating a random vector whose marginal distributions correspond to X and Y respectively. The choice of W is generally not unique, and the whole idea of "coupling" is about making such a choice so that X and Y can be related in a certain way we desire. See Lindvall (1992).

non-transparent graph copies are generated. The utility function of the researcher is summarized as:

$$(\mathbb{G})_{t_{i=1\ldots i*\ldots}} = (G_{v,e})_{t_{i<i*}} - \mathscr{B}\,(\mathscr{G}_{\mathscr{V},\mathscr{E}})_{t_{i\geq i*}}$$

where $\mathscr{B} = [0,1]$. By focusing on the network, I am able to integrate all behavioral problems into one simple representation. In the next section, I outline the general model, which relaxes the assumption that behavioral graph copies are independent.

3 General Model

In the general model of multiple graphs, the research graph and its copies are not necessarily independent and easily seperated in data as represented in the above discussion. Identifying and dropping specific graph copies is trivial in the simple case. On the other hand, part of a research path may be in both the rational and behavioral parts of the research graph. In this section, I show that the approach of eliminating graph copies is feasible in such complex cases as well.

Graph partitions: A partition of the vertex set V of G is

$$V:\ G \to V_1 \cup \ldots \cup V_k$$

on n vertices into parts which differ in cardinality by at most 1.

Mean square density: The mean square density of the partition is given as

$$\sum_{1\leq i,,j\leq k} p_i p_j d\,(V_i, V_j)^2$$

whereby $d\,(V_i, V_j)$ is the edge density between two subsets of vertices of a graph G or the fraction of pairs $(v_i, v_j) \in V_i \times V_j$ that are edges of G.

Refinement of partition: A refinement $\mathcal{A}$ of a partition refers to another partition $\mathcal{B}$ such that each member of $\mathcal{B}$ is a subset of some member of $\mathcal{A}$.

Regularity lemma: Any graph on n vertices with $o\left(n^3\right)$ triangles can be made triangle-free by removing $o\left(n^2\right)$edges.

If the graph partition does not satisfy the regularity lemma, the partition can be refined so that the mean square density increases by $\Omega\left(\varepsilon^5\right)$. The number of parts is exponential at most exponential in k. The process must stop after $O\left(\varepsilon^{-5}\right)$ because the mean square density cannot exceed 1.

$\mathscr{G}$-free: Being "free" of a graph means eliminating each copy of graph $\mathscr{G}$.

Szemerédi's regularity lemma: Every large graph can be partitioned into a small number of parts such that the bipartite subgraph between almost every pair of parts is random-like.

Mean entropy density: The mean entropy density with respect to the partition is

$$\sum_{1\leq i,,j\leq k} p_i p_j f\left(d\left(V_i, V_j\right)\right)$$

where $f\left(x\right)=x\text{log}x$ for $0<x\leq 1$ and $f\left(0\right)=0$. Since $f\left(x\right)$ is nonpositive for $0<x\leq 1$ implies that the mean entropy desnity is always nonpositive as well.

Shattering of sets: An important tool is based on "shattering" sets with a few copies of G.

Shattering definition: Label the graph G with vertex set $[g] := \{1,\dots g\}$. If $V_1,\dots,V_g$ are vertex subsets of a graph such that there are few copies of G with the copy of vertex i in V_i for $i\in[g]$, then there is an edge $(i,\ j)$ of G such that the pair (V_i,V_j)can be shattered as follows:

An $(\alpha,m,t)-$ *shattering* of a pair (A,B) of vertex subjects in a graph G is a pair of partitions $A=A_1\cup\dots\cup A_r$ and $B=B_1\cup\dots\cup B_s$ such that $r,s\leq t$ and the sum of $|A_i|\,|B_j|$over all pairs (A_i,B_j) with $d\left(A_i,B_j\right)<\alpha$ is at least $m\,|A_i|\,|B_j|$.

Jensen defect inequality. The Jensen defect inequality is implemented with the set shattering. The Jensen defect inequality states that if f is a convex function, and $\varepsilon_{i=1\dots s}$ are nonnegative reals which sum up to 1, and $x_1,\dots x_s\in\mathbb{R}$ then

$$\varepsilon_i\sum_{i=1}^{s} f\left(x\right)\geq f\sum_{i=1}^{s}\left(\varepsilon_i x_i\right)$$

$$\Rightarrow \varepsilon_1 f(x_1) + \ldots + \varepsilon_s f(x_s) \geq f(\varepsilon_1 x_1 + \ldots + \varepsilon_s x_s)$$

Theorem 1: Let H refer to behavioral research graphs. For each graph H on h vertices, if δ^{-1} is a tower of twos of height $5h^4 \log\varepsilon^{-1}$, then every graph G on n vertices with at most δn^h copies of H can be made H-free by removing εn^2 edges.

I sketch the proof here with the full argument relegated to the Appendix. Let H be a fixed graph with h vertices. The proof proceeds by contradiction. Suppose that $G = (V, E)$ is a graph on n vertices for which εn^2 edges need to be removed to make it H-free, but that G contains fewer than δn^h copies of H.

We pass to a subgraph $G' \in G$ such that G'consists of the union of a maximum collection of edge-disjoint copies of H in G. As the removal of the edges of G' leaves an H-free subgraph of G, the graph G'has at least εn^2 edges.

Let

$$q = 2e\left(\frac{G'}{n^2}\right) \geq 2\varepsilon$$

At each stage of the proof, we have a partition $V = V_1 \cup \ldots \cup V_k$ of the vertex into parts such that almost all vertices are in parts of the same size.

Let

$$p_i = \frac{|V|}{n}$$

Recall that the mean entropy density is given by

$$\sum_{1 \leq i,, j \leq k} p_i p_j f(d(V_i, V_j))$$

where $f(x) = x\log x$ for $0 < x \leq 1$ and $f(0) = 0$.

Now, I use the definition of shattered sets and the Jensen defect inequality to show that the partition can be refined such that the mean entropy density increases by $\Omega(d)$ while the number of parts only goes up exponentially in $m(\varepsilon, z)\, s$, where $m(\varepsilon, z) = 2^{\left(\frac{z}{\varepsilon}\right)^{O\left(z^2\right)}}$.

So essentially, in each iteration, the number of parts is one exponential larger. This process must stop after $O\left(\log d^{-1}\right) = O\left(\log\varepsilon^{-1}\right)$ steps, as the mean entropy density is at least $d\log d$. The mean entropy density increases $\Omega(d)$ at each refinement, and is always nonpositive. We therefore obtain a bound on δ^{-1} in the theorem which is a tower of twos in height $O\left(\log\varepsilon^{-1}\right)$.

3.1 Example: Building Transparent Economic Models

The first example focuses on how building economic models can be made more transparent. The setting builds on Varian (2016). The author notes in Varian (2016:1): "In reality the process is much more haphazard than my description would suggest", implying that the use of random graphs is appropriate for this illustration of research transparency. The author also suggests that economists "...come up lots and lots of ideas and throw out all the ones that aren't good." (see Varian (2016:1)), which is partly consistent with my idea that dropping behavioral graphs relates to standard practice. The author also suggests that economists "not look at the literature too soon." Varian (2016: 3) which further implies that many modeling decisions are random.

The process begins with the generation of a research idea, often from media sources. Authors may document all of the sources and share these online using links to make such sources available to readers. This step could be skipped if a research idea comes from personal experience, although using a mobile app and voice recognition tools might serve this purpose. In the model, the first node v_1 is this research idea, and the problem of the researcher is to evaluate whether it is worth pursuing.

The first proposed step is to talk about the concept with a range of non-economists and receive a strong reaction. The economist shall next get feedback from mentors and audiences while refining the research. Suppose that the feedback requires changes that the economist accepts to be important, but hesitates to make, given the effort already exerted. The researcher starts by making the graph as simple as possible before generalizing it. The author uses Bayesian updating to make the graph simple: she only adds edges that do suit her prior from feedback, then collects more feedback and updates her priors. If new edges unnecessarily complicate the model (e.g. involving a triangle), these are dropped.

Generalizing the model to a more general version to analyze complex environments means aggreating the discussion to the level of graphs. If the researcher keeps graph copies, the theory is not transparent. Dropping graph copies leave the objective theory. All calculations and proofs, including rough work and failed attempts are recorded on an interactive pad and made available online to replicators, making this example possible.

Given the complexity involved in generating general versions of economic theoretical models, I use the dropping of homomorphisms as an analogy for the dropping of graphs discussed in the main results. This is a simple generalization of the prior discussion[7]. The argument is based on the following terms:

Homomorphic image: For a graph H, the homomorphic image is a graph F for which there is a surjective homomorphism from H to F. Each homomorphic image of H has$|H|$vertices.

Blow-up lemma (Komlós; Sárkőzy, and Szemerédi, 1998). The blow-up lemma states that a system of super-regular pairs contains all bounded degree spanning graphs as subgraphs that embed into a corresponding system of complete pairs.

Given a graph R that is characterized by the following

1. order r,
2. minimal vertex degree ρ
3. maximal vertex degree $\triangle$,'

Let there be an arbitrary positive integer, and replace the vertices of R with pairwise disjoint N-sets $V_1 \dots V_r$ (blowing up).

Construct two graphs on the same vertex set $V = \cup V_i$. The graph $R(N)$ is obtained by replacing all edges of R with copies of the complete bipartite graph $K_{N,N}$ and construct a sparser graph by replacing the edges of R with some (ε, ρ)-superregular pair. If a graph H with $\triangle(H) \leq \triangle$ is embeddable in $R(N)$, then it is already embeddable into G.

[7]One may think of this process as being very broadly analogous to homomorphic filtering, which is a computing technique used in image processing for image enhancement.

The *blow-up* $G(a_1 \ldots, a_k)$of G is the graph obtained from G by replacing each vertex by an independent set W_i of order a_i

Proposition (Removing homomophisms). For every graph H on h vertices and every $\varepsilon > 0$, there is $\delta > 0$ such that if G is a graph on n vertices with at most δn^h copies of H then εn^2 edges of G can be removed to obtain a graph $G'$for which there is no homomorphism from H to G'.

Proof. I provide a contradiction to the statement that G has at most δn^h copies of H.The δ in this homomorphism discussion is akin to the one in the graph removal discussion. The proof builds on the graph removal discussion and Szemerédi's regularity lemma. Let G be a graph on n vertices which has at most δn^h of H. Recall that the homomorphic image is a graph F for which there is a surjective homomorphism from H to F and that each homomorphic image of H has$|H|$vertices, so that the number of homomorphic images of H is finite.

Observe that to remove all homomorphisms from H to G it suffices to remove all copies of homomorphic images of H in G. Graph removal can remove few edges and remove all homomorphisms from H to G if there are few copies of G of each homomorphic images of H in G.

There must be a homomorphic image F of H for which there are many copies of F in H. Let this number of copies be cn^k such that $c > \delta^{h^{-h}}$, whereby k is the number of vertices of F.

Let f be a surjective homomorphism from H to F. For each vertex i of F, let a_i represent the number of vertices of H which map to vertex i in f.

The *blow-up* $F(a_1 \ldots, a_k)$of F is the graph obtained from F by replacing each vertex by an independent set I_i of order a_i. A pair of vertices in different parts I_i and I_j are adjacent if and only if i and j are adjacent in F.

Observe that H is a subgraph of the blow-up $F(a_1 \ldots, a_k)$.

Let $\mathcal{S} = (v_1 \ldots, v_k)$ refer to the set of sequences of k vertices of G which form a copy of F with v_i the copy of vertex i.

Observe that if $A_1 \ldots, A_k$ are vertex subsets of G with $|A_i| = a_i$, and all k-tuples in

$A_1 \times \cdots \times A_k$ belong to $\mathcal{S}$, then these vertex subsets form a copy of $F\left(a_1, \ldots a_k\right)$in G, and hence also make a copy of H in G.

Since G has cn^k copies of F, then from Erdos and Simonovits (1983), a convexity argument shows that if

$$c \gg n^{\frac{-1}{\left(a_1, \ldots a_k\right)}}$$

then $\mathcal{S}$ contains at least $\left(1 - o\left(1\right)\right) c^{a_1 \ldots a_k} n^h \geq \left(1 - o\left(1\right)\right) \delta \left(\frac{3^{\frac{1}{3}}}{h}\right)^h n^{^h} \geq h! \delta n^h$ labeled copies of H, where we use $a_1 \cdots a_k \leq 3^{\frac{h}{3}}$ since $a_1, \ldots a_k$ are positive integers which sum to h, and $c > \delta^{h^{-h}}$. This provides a contradiction to the statement that G has at most δn^h copies of H. $\square$

3.2 Example: Transparent Qualitative Research for Program Evaluations

This example expands on recent work on how qualitative research plays a significant role in program evaluation research in development economics (Glennerster and Takavarasha, 2013). Economists interview a proposed project with beneficiaries and attempt to incorporate their subjective experiences and insights into program evaluation design to understand the context. I assume that this is based on a theoretical model, as discussed previously. I discuss how such qualitative work can occur more transparently. To the best of my knowledge, there is little work on qualitative research transparency.

Qualitative research based on large amounts of ethnographic or descriptive information. The objective path is to pivot the paper to suit subject responses, and the behavioral outcome is to deviate from the objective trajectory.

Suppose that the qualitative data results and feedback requires program changes that the economist accepts to be important, but hesitates to make, given the effort already exerted in raising funds for a program evaluation. This can be thought of as an issue where the economist is minimizing $\sum_{e\in R} c(e)$ over $R \in R(u,v)$.

The qualitative data collection is modeled as a directed graph. If the researcher keeps graph copies, the ensuring program evaluation is not transparent. I assume that all interview transcripts are anonymized and recorded on an interactive pad and made available online to replicators.

Directed graphs: The interviewer interviews individual subjects or individual focus groups, so that the graphs are directed.

Directed graph removal lemma (Alon and Shapira (2004)). For each directed graph H on h vertices and $\varepsilon > 0$, there is $\delta = \delta(\varepsilon, H) > 0$ such that every directed graph $G = (V, E)$ on n vertices with at most δn^h copies of H can be made H-free by removing at most εn^2 edges.

Proof. This borrows from Theorem 1. I first find a subgraph $G'$of G which is the disjoint union of $\varepsilon' n^2$ copies of H, with $\varepsilon' \geq 2h^{-2}\varepsilon$.

There is a partition $V = V_1 \cup \ldots \cup V_h$ with at least $h^{-h}\varepsilon' n^2$ edge-disjoint copies of H with the copy of vertex i in V_i. Indeed, in a uniform random partition into h parts, each copy of H has the probability h^{-h} that its copy of vertex i lies in V_i for all $i \in [h]$.

I then let G''be the subgraph of G'which consists of the union of these at least $2h^{-h-2}\varepsilon n^2$ edge-disjoint copies of H.

The rest of the proof is the same as the proof of Theorem 1, with the exception that I start with the partition $V = V_1 \cup \ldots \cup V_h$ and refine it further at each step.

Ethnography. Recent work has documented how ethnographic research can be integrated into program evaluations (e.g. Gertler et al, 2016). Such information may consist of rich text to be grouped by topic themes and are significantly more complex than standard quantitative information represented by a spreadsheet. Moravcsik (2014) contains a discussion on the difficulties of using secondary sources such as ethnographic notes, primary responses, archives and interviews.

In this example, I propose that qualitative information gained from such ethnographic information can be defined by a hypergraph in the model context, as hypergraphs entail having an edge being able join any number of vertices. Since qualitative responses have less structure than quantitative ones, some qualitative responses may connect to several other study subject responses and it is important to transparently present these where relevant. Given several interconnecting themes and variables, it may not be clear how to evaluate a relevant program. I show how the approach can help such qualitative research be more transparent. The hypergraph removal lemma (see Tao 2006) for a variation) shows how hypergraph copies may be dropped. As such, even when qualitative information that may be highly endogenous (from the ability of an edge to join any number of vertices), it is still feasible to extract transparent information in a transparent way for integration into program evaluation design.

4 Concluding Discussion

The paper introduces and solves the behavioral question of research transparency in both theoretical and empirical economics research. The approach has certain characteristics that are worth emphasizing. Although the ethics of the research economist has not yet been explicitly modeled in the literature, it has mainly been imagined as a black-box response to unaccountability. This can be interpreted literally as an agent navigating a graph, or alternatively, as a sequence of decision-making agents that are at once rational and behavioral. However, it should be emphasized that in many settings only a sample of the agent's action is observable or recordable or even predictable ex-ante. When the universe of research decisions cannot be recorded, it seems appropriate for the maximum number of observable decisions to serve as a proxy.

Researchers may use preanalysis plans to preregister the graph of their analyses and report additional analyses as exploratory. Going further, whether or not a researcher uses a preanalysis plan at the start of their study or not, there is in principle the possibility of using timestamps to record every data analysis decision, and to use dynamic documents to record exploratory research decisions and report them. By editing the preanalysis plan, researchers can strengthen the transparency of their exploratory research. Although innovations in software and dynamic document notebooks make these technologically feasible, the incentives to foster adoption are very limited in scope. Indeed, the Berkeley Initiative for Transparency in the Social Sciences (<http://www.bitss.org/>) is one open network that exposes economists and other scholars on the use of cutting-edge software and other technologies that embed a research paper into data analyses and help make more research decisions common in economics transparent with data management and other tools. All of their tools and materials are available online for free download, and there is also a Massively Open Online Course (FutureLearn, 2017). However, a general study of fostering adoption and scale up of timestamped analyses that report the time of analysis (to be ranked in terms of their timeliness) is a topic of ongoing research.

In this paper, the decisions made during objective research is non-random, however, exploratory research is assumed to mimic the formation of a random graph. This allows for an identification of exploratory graphs and commitment to focusing on objective graphs (shows by the dropping of exploratory graphs by a commitment device) but this randomness assumption is not critical for the analysis. The role of the underlying process is summarized in the graph removal, and any process that gives rise to graph copies that must be deleted can be accommodated.

The agent is assumed to be willing to submit the universe of code and data queries in near-real time to a repository system and to rely on intuition in reporting final results in the research paper. Preanalysis plans can be accomodated in a straightforward way. This is because the researcher can prespecify hypotheses and methodologies prior to public data access if the timestamp of the data download can be reported from the data-hosting website. Indeed, replications of the paper can be allowed to depend on the data analysis choices of the original researcher, which would be available in much richer detail than in many present replication cases. In particular, we can continue to document all human-computer interactions associated with a research project because it will completely summarize the researcher actions that define a paper, and enable replicators to "back out" an objective trajectory in part or in full while filling in gaps. Other researchers may also perform similar secondary sensitivity analyses on the replications.

However, this is no longer true if the first replicator does not have access to the entire set of human-computer interactions of the original researcher with respect to the paper.

Then the replicator may provide their human-computer interactions to other replicators. Other methods such as video can be used to document the history of a project, and closer attention can be paid to software and version control. Indeed, a working paper version (Opoku-Agyemang, 2017) considers the version where the researcher has self-control problems in generating code and benefits from simpler point-and-click interfaces that automate the generation of code. That paper shows how even the text that accompanies data visualizations can be automated, to help scholars fight the temptation to exaggerate results. In this setting, the policy is identical to the one studied here.

The researcher's creativity is important here as it is during standard preanalysis plans. Formally, the dynamic documents enable the self-commitment principle to allow researchers to generate their individualized research perspective graphs and "factorized out" the underlying graph that is objective. It is worth emphasizing that even when the research is unable to fully commit to a preanalysis plan to be edited and timestamped along with the data analysis, the self-commitment principle is useful as an optimal benchmark for nonexperimental analyses that sub-optimal policies can support. In the absence of a commitment device such as a preanalysis plan, the environment raises the new possibility that human-computer interactions can serve that purpose. Other innovative approaches involve enabling publishers investing to promote scientific integrity (Lee and Moher, 2017). Generalizing the discussion to other economist settings of policy, consulting, teaching and service would make relevant future research.

Appendix A

PROOF OF THOREM 1: Let H be a fixed graph with h vertices. The proof proceeds by contradiction. Suppose that $G = (V, E)$ is a graph on n vertices for which εn^2 edges need to be removed to make it H-free, but that G contains fewer than δn^h copies of H.

We pass to a subgraph $G' \in G$ such that G'consists of the union of a maximum collection of edge-disjoint copies of H in G. As the removal of the edges of G' leaves an H-free subgraph of G, the graph G'has at least εn^2 edges. Let $\varepsilon_0 n^2$ represent the number of edge-disjoint copies of H in G' so that $e(G') = e(H)\varepsilon_0 n^2$.

Since there is at least one and at most δn^h copies of H, we have $n \geq \delta^{\frac{-1}{h}}$. Let $\mathcal{P}_0$ be an arbitrary partition $V = V_1 \cup \ldots \cup V_k$ of the vertex set of G'into parts of size $n_0 = \left[\frac{\varepsilon_0}{8}n\right]$, except possibly one remaining set of size less than $\frac{\varepsilon_0}{8}n$. The number p_0 of parts of $\mathcal{P}_0$is at most $8\varepsilon_0^{-1} + 1 \leq 5h^2\varepsilon^{-1}$. I next rely on the following lemma:

Lemma. Let $f : \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}$ be a convex function, $G = (V, E)$ be a graph and $d = d(V, V) = 2\frac{|E|}{|V|^2}$. Then, the following two conditions hold

1. For vertex subsets $A, B \subset V$ and partitions $\mathcal{A}$of A and $\mathcal{B}$ of B, $f(\mathcal{A}, \mathcal{B}) \geq f(A, B)$
2. if $\mathcal{P}$ is a partition of V, then $f(\mathcal{P}) \geq f(d)$.

The proof of the first condition is shown now. I obtain the following:
$f(\mathcal{A}, \mathcal{B}) = \sum_{A' \in \mathcal{A}, B' \in \mathcal{B}} f(A', B') = \sum_{A' \in \mathcal{A}, B' \in \mathcal{B}} \frac{|A'||B'|}{|V|^2} f(d(A', B')) = \frac{|A'||B'|}{|V|^2} \sum_{A' \in \mathcal{A}, B' \in \mathcal{B}} f(d(A', B')) \geq \frac{|A||B|}{|V|^2} f(d(A, B)) = f(A, B)$, while relying on

$$\sum_{A' \in \mathcal{A}, B' \in \mathcal{B}} \frac{|A'|\,|B'|}{|A|\,|B|} = 1$$

and Jensen's inequality. This completes the proof for the first condition.

To prove the second condition, observe that if $\mathcal{P}$ is a partition of V, then the first condition implies that $f(\mathcal{P}) = f(\mathcal{P}, \mathcal{P})$. Also, $f(\mathcal{P}, \mathcal{P}) \geq f(V, V)$ and $f(V, V) = f(d)$. This completes the proof for the second condition.

Integrating both conditions into Theorem 1, I obtain $f(\mathcal{P}_0) \geq f(d) = d\log d$, such that $d = \frac{2e(G')}{n^2} \geq 2\varepsilon$.

I rely on the following lemma to obtain a sequence of partition refinements:

Lemma (Fox, 2011).

- Let H be a graph on h vertices. Suppose $G = (V, E)$is a graph on n vertices whose edge can be partitioned into $\varepsilon_0 n^2$ of H.

- Let $n_0 \leq \frac{\varepsilon_0}{4} n$ be a positive integer.
- Let $\mathcal{P}$ be a partition of V into at most T parts with all parts of size at most n_0 and all but at most $\frac{\varepsilon_0}{8} n$ vertices in parts of size n_0.
- Let G have at most $2^{-\left(\frac{40}{\varepsilon_0}\right)^{h^2}} T^{-h} n^h$ copies of H.
- Let $f(x) = x\log x$ for $x > 0$ and $f(0) = 0$.

Then,

- There is a refinement of $\mathcal{P}'$ to $\mathcal{P}$ with at most s^T parts with $s = 2^{2^{\left(\frac{50}{\varepsilon_0}\right)^{h^2}}}$, so that ,$f(\mathcal{P}') \geq f(\mathcal{P}) + \frac{\varepsilon_0}{4h^2}$ and all but at most $\frac{\varepsilon_0}{8} n$ vertices are in part of equal size, and all other parts are of smaller size.

This last lemma is implemented repeatedly to obtain a sequence of partition refinements

$$\mathcal{P}_i = \mathcal{P}_0, \mathcal{P}_1 \ldots$$

such that p_i defines the number of parts of $\mathcal{P}_i$.

As soon as I obtain the partition$\mathcal{P}_i$, I can apply the lemma to obtain a refinement $\mathcal{P}_{i+1}$ of $\mathcal{P}_i$, as long as $\delta \leq 2^{-\left(\frac{40}{\varepsilon_0}\right)^{h^2}}$. After implementing i iterations, $f(\mathcal{P}_i) \geq f(\mathcal{P}_0) + i\frac{\varepsilon_0}{4h^2}$ and $p_i \leq s^{p_{i-1}}$ and where $s = s = 2^{2^{\left(\frac{50}{\varepsilon_0}\right)^{h^2}}}$.

At each iteration, the number of parts is approximately one exponential larger than in the previous iteration. Note that this process continues for at least $i_0 := \left[4h^4\log\varepsilon^{-1}\right]$ interations, because δ^{-1} is a tower of twos of height $5h^4\log\varepsilon^{-1}$. Since $h^2\varepsilon_0 > 2e(H)\varepsilon_0 = d$ and $d \geq 2\varepsilon$, $f(\mathcal{P}_{i0}) \geq f(\mathcal{P}_0) + i_0\frac{\varepsilon_0}{4h^2} \geq d\log d + \left(4h^4\log\varepsilon^{-1}\right)\frac{\varepsilon_0}{4h^2} = d\log d + h^2\varepsilon_0\log\varepsilon^{-1} > d\log\left(\frac{d}{\varepsilon}\right) > 0$.
This contradicts the statement that f applied to any partition is nonpositive.$\square$.

References

1. Aker, Jenny C., Ishita Ghosh, and Jenna Burrell. 2016. "The promise (and pitfalls) of ICT for agriculture initiatives." *Agricultural Economics* 47(S1): 35-48.

2. Anderson, Richard G. and Areerat Kichkha. 2017. "Replication, Meta-analysis, and Research Synthesis in Economics." *American Economic Review: Papers and Proceedings*, 107(5): 56-59.

3. Anderson, Michael L., and Jeremy Magruder. 2017. "Split-Sample Strategies for Avoiding False Discoveries." National Bureau of Economic Research Working Paper.

4. Bedev, Anton. 2013. "Discrete Games in Endogenous Networks: Theory and Policy." University of Pennsylvania Working Paper.

5. Berry, James, Lucas C. Coffman, Douglas Hanley, Rania Gihleb and Alistair J. Wilson. 2017. "Assessing the Rate of Replication in Economics." *American Economic Review: Papers and Proceedings*, 107(5): 27-31.

6. Card, Stuart K., Thomas P. Moran, and Allen Newell. 1980. "The keystroke-level model for user performance time with interactive systems". *Communications of the ACM*. 23 (7): 396–410.

7. Card, Stuart K., Thomas P. Moran, and Allen Newell. 1983. *The Psychology of Human-Computer Interaction*. New Jersey: Lawrence Erlbaum Associates.

8. Carlisle, James H. 1976. "Evaluating the impact of office automation on top management communication." *Proceedings of the June 7–10, 1976, National Computer Conference and Exposition*, 611–616.

9. Chandrasekhar, Arun and Matthew O. Jackson. 2014. "Tractable and Consistent Random Graph Models." Stanford University Working Paper.

10. Chandrasekhar, Arun and Matthew O. Jackson. 2016. "A Network Formation Model Based on Subgraphs." Stanford University Working Paper.

11. Christiansen, Garrett and Edward Miguel. 2017. "Transparency, Reproducibility, and the Credibility of Economics Research." *Journal of Economic Literature*, forthcoming.

12. Chang, Andrew C. and Phillip Li. 2017. "A Preanalysis Plan to Replicate Sixty Economics Research Papers That Worked Half of the Time." *American Economic Review: Papers and Proceedings*, 107(5): 60-64.

13. Coffman, Lucas C. and Muriel Niederle. 2015. "Pre-analysis Plans Have Limited Upside, Especially Where Replications Are Feasible." *Journal of Economic Perspectives*, 29(3), 81-98.

14. Coffman, Lucas C., Muriel Niederle and Alistair J. Wilson. 2017. "A Proposal to Organize and Promote Replications." *American Economic Review: Papers and Proceedings*, 107(5): 41-45.

15. Deaton, Angus. 2016. "Measuring and Understanding Behavior, Welfare, and Poverty." *American Economic Review*, 106(6): 1221–1243.

16. Durak, Nurcan, Ali Pinar, Tamara G. Kolda, and C. Seshadhri. 2012. "Degree relations of triangles in real-world networks and graph models." *CIKM'12.*

17. Duvendack, Maren, Richard Palmer-Jones and W. Robert Reed. 2017. "What Is Meant by "Replication" and Why Does It Encounter Resistance in Economics? *American Economic Review: Papers and Proceedings*, 107(5): 46-51.

18. Fachamps, Marcel, and Julien Labonne, 2017. "Replication Data for "Using Split Samples to Improve Inference on Causal Effects." *Political Analysis*, forthcoming.

19. Flexner, Abraham. (1939). "The Usefulness of Useless Knowledge." *Harper's Magazine*, 544-552.

20. Fox, Jacob. 2011. "A new proof of the graph removal lemma." *Annals of Mathematics,* 174(1): 561-579.

21. FutureLearn. 2017. *Transparent and Open Social Science Research Online Course.* https://www.futurelearn.com/courses/open-social-science-research

22. Gelman, Andrew and Eric Loken. .2013. "The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time" Unpublished Working Paper. http://www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf

23. Gertler, Paul J., Sebastian Martinez, Patrick Premand, Laura B. Rawlings, and Christel M. J. Vermeersch. 2016. *Impact Evaluation in Practice: Volume 2.* Washington DC: World Bank.

24. Goldreich, Oded, Shari Goldwasser, and Dana Ron. 1998. "Property testing and its connection to learning and approximation." *Journal of the ACM*, 45(4): 653-750..

25. Graham, Bryan. 2015. "Methods of Identification in Social Networks." *Annual Review of Economics*, 7: 465-485.

26. Hamermesh, Daniel S. 2017. "Replication in Labor Economics: Evidence from Data, and What It Suggests." *American Economic Review: Papers and Proceedings*, 107(5): 37-40.

27. Höffler, Jan H. 2017. "Replication and Economics Journal Policies." *American Economic Review: Papers and Proceedings*, 107(5): 52-55.

28. Kleinberg, Jon and Sigal Oren. 2014. "Time-inconsistent planning: a computational problem in behavioral economics." *In Proceedings of the Fifteenth ACM Conference on Economics and Computation*, 547-564).

29. Komlós, J.; G. N. Sárkőzy, and E. Szemerédi. 1998. "Proof of the Seymour Conjecture for Large Graphs." *Annals of Combinatorics*, 2, 43-60.

30. Laibson, David. 1997. "Golden eggs and hyperbolic discounting." *Quarterly Journal of Economics,* 112(2): 443-478.

31. Lee, Carole J., and David Moher. 2017. "Promote scientific integrity via journal peer review data." *Science,* 357(6348): 256-257.

32. Lindvall, T. 1992. *Lectures on the coupling method.* New York: Wiley.

33. Mele, Angelo. 2010. "A structural model of segregation in social networks." CEMMAP Working Paper.

34. Moravcsik, Andrew. 2014. "Transparency: The Revolution in Qualitative Research." *PS: Political Science and Politics,* 47(1), 48-53.

35. Merton, Robert K. 1973. *The Sociology of Science: Theoretical and Empirical Investigations.* Chicago: University of Chicago Press.

36. Opoku-Agyemang, Kweku. 2017. A Human-Computer Approach for Research Transparency. Cornell Tech Working Paper.

37. Opoku-Agyemang, Kweku, Bhaumik Shah and Tapan Parikh. 2017. Scaling Up Peer Education with Farmers in India. Cornell Tech Working Paper.

38. Opsahl, Tore and Pietro and Panzarasa 2009. Clustering in weighted networks. *Social Networks*, 31, 155–163.

39. Plant, A. L., L. E. Locascio, W. E. May, and P. D. Gallagher. 2014. "Improved reproducibility by assuring confidence in measurements in biomedical research." *Nature Methods,* 11, 895–898.

40. Rubinfeld, Ronitt, and Madhu Sudan. 1996. "Robust characterizations of polynomials with applications to program testing." *SIAM Journal on Computing* 25(2): 252-271.

41. Sené, Martyn, Ian Gilmore and Jan-Theodoor Janssen. 2017. Metrology is Key to Reproducing Results. *Nature,* 547, 397–399.

42. Simmons, Joseph P., Leif D. Nelson, and Uri Simonsohn. 2011. False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science* 22(11): 1359 –1366.

43. Sukhtankar, Sandip. 2017. "Replications in Development Economics." *American Economic Review: Papers and Proceedings*, 107(5): 32-36.

44. Szemerédi, Endre. 1975. "On sets of integers containing no k elements in arithmetic progression." *Acta Arithmetica.* 27, 199–245.

45. Szemerédi, Endre. 1976. "Regular partitions of graphs, in Problmes Combinatoires et Theorie des Graphes." (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), Colloq. Internat. CNRS 260, CNRS, Paris, 1978, 399–401.

46. Tao, Terence. 2006. "A variant of the hypergraph removal lemma." *Journal of Combinatorial Theory Series A*, 113, 1257-1280.

47. Varian, Hal. 2016. "How to Build an Economic Model in Your Spare Time." *American Economist,* 61(1): 81-90.