

An Experimental Study of the Coloring Problem on Human Subject Networks

Michael Kearns,* Siddharth Suri, Nick Montfort

Theoretical work suggests that structural properties of naturally occurring networks are important in shaping behavior and dynamics. However, the relationships between structure and behavior are difficult to establish through empirical studies, because the networks in such studies are typically fixed. We studied networks of human subjects attempting to solve the graph or network coloring problem, which models settings in which it is desirable to distinguish one's behavior from that of one's network neighbors. Networks generated by preferential attachment made solving the coloring problem more difficult than did networks based on cyclical structures, and "small worlds" networks were easier still. We also showed that providing more information can have opposite effects on performance, depending on network structure.

It is often thought that structural properties of naturally occurring networks are influential in shaping individual and collective behavior and dynamics. Examples include the popular notion that "hubs" or "connectors" are inordinately important in the routing of information in social and organizational networks (1, 2). A long history of research has established the frequent empirical appearance of certain structural properties in networks from many domains, including sociology (1, 3–5), biology (6, 7), and technology (8). These properties include small diameter (the "six degrees of separation" phenomenon), local clustering of connectivity (9), and heavy-tailed distributions of connectivity (10). Theoretical models have sought to explain how some of these may interact with network dynamics (11).

The relationships between structure and behavior are difficult to establish in empirical field studies of existing networks. In such studies, the network structure is fixed and given, thus preventing the investigation of alternatives. A different approach is to conduct controlled laboratory studies in which network structure is deliberately varied.

We have been performing human subject experiments in distributed problem-solving from local information on a variety of simple and complex networks. Subjects each simultaneously control a single vertex in a network of 38 vertices and attempt to solve the challenging graph coloring problem (12) on the network. In this problem, the collective goal is for every player to select a color for their vertex that differs from the colors of all of their network neighbors. The number of colors made available is the minimum necessary to color the entire network without conflicts (edges connecting two vertices with the same color), known as the chromatic number of the network.

The graph coloring problem is a natural abstraction of many human and organizational problems in which it is desirable or necessary to distinguish one's behavior from that of neighboring parties. As a specific scenario, consider the problem faced by faculty members scheduling departmental events—recurring classes, one-time seminars, exams, and so on—in a limited number of available rooms. We can view the events to be scheduled as the vertices in a network, with an edge connecting any pair of events that temporally overlap, even partially. Clearly, two such events must be assigned to different rooms or "colors," thus yielding a natural graph coloring problem. Furthermore, even when there is a centralized first-come, first-serve sign-up sheet for rooms, this mechanism is simply the starting point for the negotiation of a solution, and the problem is still solved in a largely distributed fashion by the participants: Faculty members routinely query the current holder of a room whether they might be able to switch to a different room, whether their event will really require their entire time slot, and the like. Other coloring-like problems arise in a

variety of social activities (such as selecting a cell phone ringtone that differs from those of family members, friends, and colleagues); technological coordination [selecting a channel unused by nearby parties in a wireless communication network (13, 14)]; and individual differentiation within an organization (developing an expertise not duplicated by others nearby). Graph coloring also generalizes many traditional problems in logistics and operations research (12).

The coloring problem was chosen for both its simplicity of description and its contrast to other distributed network optimization problems. Unlike the well-studied navigation or shortest-paths problem, optimal coloring is notoriously intractable from the viewpoint of even centralized computation (12, 15). In fact, even weak approximations (in which many more colors than the chromatic number are permitted) are known to be equally difficult (16, 17).

We report here on the findings from two extensive experimental sessions held in January 2006 with 55 University of Pennsylvania undergraduate students (18). Subjects were given a series of coloring experiments in which the network had one of six topologies, each chosen according to recently proposed models of network formation (Fig. 1 and Table 1). Three of these six begin with a simple cycle and then add a varying number of randomly chosen chords while preserving a chromatic number of two. These "small worlds" networks (9, 19) are intended to model the mixture of local connectivity (as induced by geography) with long-distance connectivity (as induced by travel or chance meetings) often found in social and other networks. The fourth cycle-based network adopted a more engineered or hierarchical structure, with two distinguished individuals having inordinately high connectivity. The fifth and sixth networks were generated according to the well-studied preferential attachment

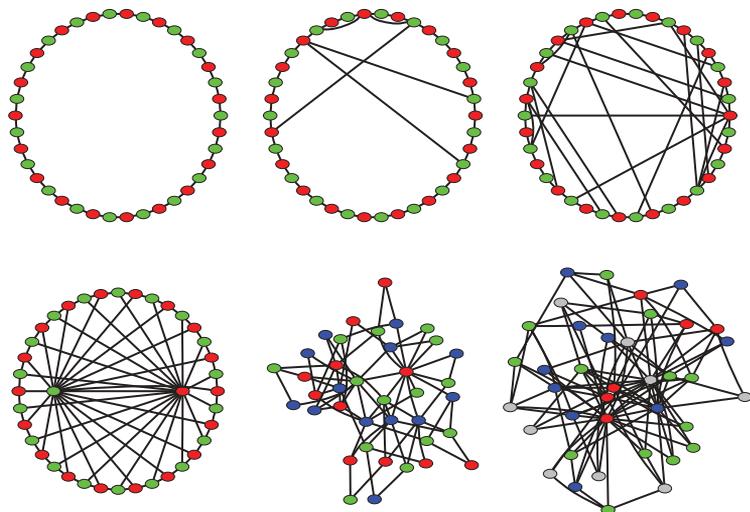


Fig. 1. Network topologies with sample colorings found by subjects. From left to right and top to bottom: simple cycle, 5-chord cycle, 20-chord cycle, leader cycle, and preferential attachment with two and three links initially added to each new vertex.

University of Pennsylvania, Department of Computer and Information Science, 3330 Walnut Street, Philadelphia, PA 19104, USA.

*To whom correspondence should be addressed. E-mail: mkearns@cis.upenn.edu

model (10), in which vertices already highly connected are more likely to receive further connections as the network is formed incrementally. Such networks are known to generate a number of structural properties frequently documented empirically, including the presence of highly connected “hubs.”

Six or seven trials (20) of each of the six networks were performed under varying informational conditions. Subjects sat at workstations running a browser-based client of a distributed computer system built for the experiments. The interface provided each subject with a local view (including their own color and those of their neighbors) of the current state (Fig. 2). In a minority of trials, subjects were given a global view. Subjects were familiarized with the coloring problem but given no guidance on how to play; subjects could update their colors at any time from a fixed menu that provided the minimum number of colors required to solve the problem for the

given network. The experimental protocol forbade all communication outside the confines of the system, and physical partitions prevented subjects from seeing beyond the information view provided on their own workstation. In accordance with standard practices in behavioral game theory and economics (21), participants received \$5 for each experiment in which they were “successful” (22).

Subjects could indeed solve the coloring problem across a wide range of network structures. Of the 38 experiments we conducted, 31 (82%) resulted in an optimal coloring of the network in less than the allotted 5 min (300 s), with the mean completion time of the solved networks only 82 s (SD, 75 s) and the median just 44 s, indicating considerable skew toward low solution times. All six of the network structures were solved at least twice when subjects could see only their neighbors’ color choices (low-information view).

Collective performance was strongly affected by network structure. The networks generated by preferential attachment proved considerably more difficult than any of the cycle-based networks: Six of the seven experiments that ended without an optimal coloring after 5 min were on networks of the former type, and the mean experiment duration [which includes 300-s values for unsolved networks (23)] for preferential attachment graphs was higher than for all others (Table 1). The duration times for the cycle-based networks and those for the preferential attachment networks passed a two-tailed variance *t* test for different means at *P* = 0.03. Within the cycle-based family, there was a monotonic relationship between solution time and network average distance (the average shortest distance, measured in number of links traveled, across all pairs of vertices), with smaller average distance leading to shorter solution times. (For the cycle-based networks, the correlation between aver-

Table 1. For each of the six experimental networks, the first six columns provide statistics summarizing structural properties, including the chromatic number (smallest number of colors required for solution), and statistics on the distribution of the degree (number of links) of each vertex. Network average distance is the average shortest-path distance, measured

in number of links traveled, over all pairs of vertices. Also displayed are the average experiment duration for each network, along with the fraction of trials on which it was solved within 300 s and the number of steps (measured in color changes) for a natural distributed computer heuristic. Pref. att., preferential attachment.

	Graph statistics						Avg. experiment duration (s) and fraction solved		Distributed heuristic (No. of color changes)
	Colors required (No.)	Min. links (No.)	Max. links (No.)	Avg. links (No.)	SD	Avg. distance (No. of links)			
Simple cycle	2	2	2	2	0	9.76	144.17	5/6	378
5-chord cycle	2	2	4	2.26	0.60	5.63	121.14	7/7	687
20-chord cycle	2	2	7	3.05	1.01	3.34	65.67	6/6	8265
Leader cycle	2	3	19	3.84	3.62	2.31	40.86	7/7	8797
Pref. att., $v = 2$	3	2	13	3.84	2.44	2.63	219.67	2/6	1744
Pref. att., $v = 3$	4	3	22	5.68	4.22	2.08	154.83	4/6	4703

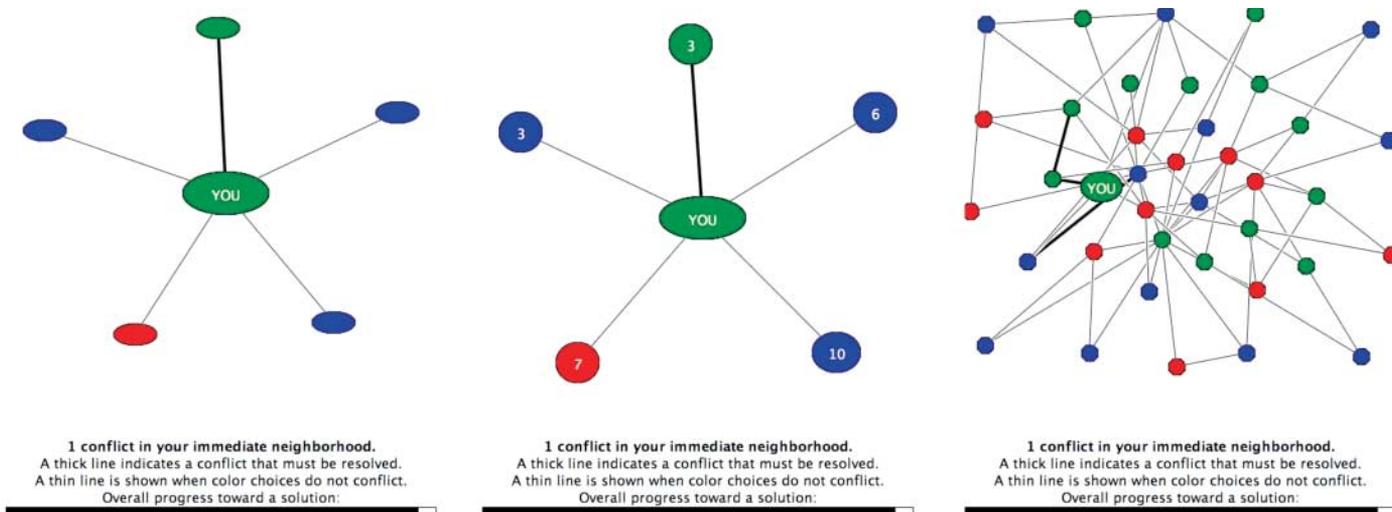


Fig. 2. In the low-information view (left), subjects could see only the color they chose for their own vertices and the colors of their immediate neighbors in the graph. The medium-information view (center) is similar, but each neighbor is labeled by its number of links. In the high-information view (right), each subject could see the current color choices of the entire

network. The text at the bottom of each screenshot reads as follows: “1 conflict in your immediate neighborhood. A thick line indicates a conflict that must be resolved. A thin line is shown when color choices do not conflict.” The bar at the bottom of each screen gave subjects an indication of global progress toward a solution.

age distance and experiment duration was 0.45, with $P = 0.02$.) Thus, the highly symmetric leader cycle (average distance 2.3) led to the fastest optimal colorings, the simple cycle (average distance 9.8) to the slowest. The addition of random chords to the simple cycle systematically reduced solution time. Although the addition of chords makes the problem more difficult from the isolated viewpoint of any individual subject (because they must now coordinate with a potentially larger number of neighbors but still are permitted only two colors), it apparently makes the collective problem easier by reducing the number of links coloring conflicts must travel to be resolved. This establishes a second and rather different problem [along with navigation (11, 24)], for

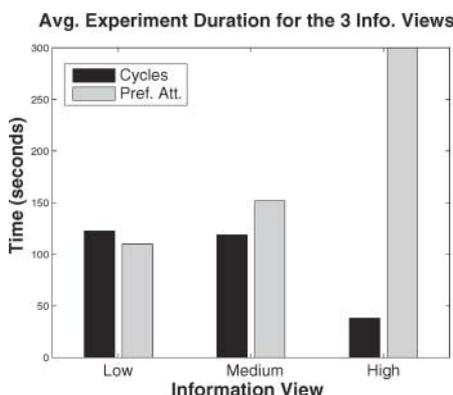


Fig. 3. More information resulted in a decrease in experiment duration for the four cycle-based graphs and an increase in duration for the two preferential attachment graphs. If we introduce an ordinal variable assuming a value of 1 for the low-information condition, 2 for the medium-information condition, and 3 for the high-information condition, the correlation between this variable and experiment duration is -0.40 for the cycle-based networks ($P = 0.04$), and 0.65 ($P = 0.02$) for the preferential attachment networks.

Fig. 4. Population convergence to one of the two possible proper colorings for the four cycle-based graphs. The y axis measures distance from the two colorings (in terms of number of disagreements), and the x axis measures time. Points below the horizontal line at $y = 19$ are closer to one of the two solutions, whereas points above this line are closer to the other. A y axis value of 0 or 38 indicates a completed experiment in which the corresponding coloring was found by the population. All experiments begin equidistant ($y = 19$) from both solutions. In the experiments under the low-information view (left), the population

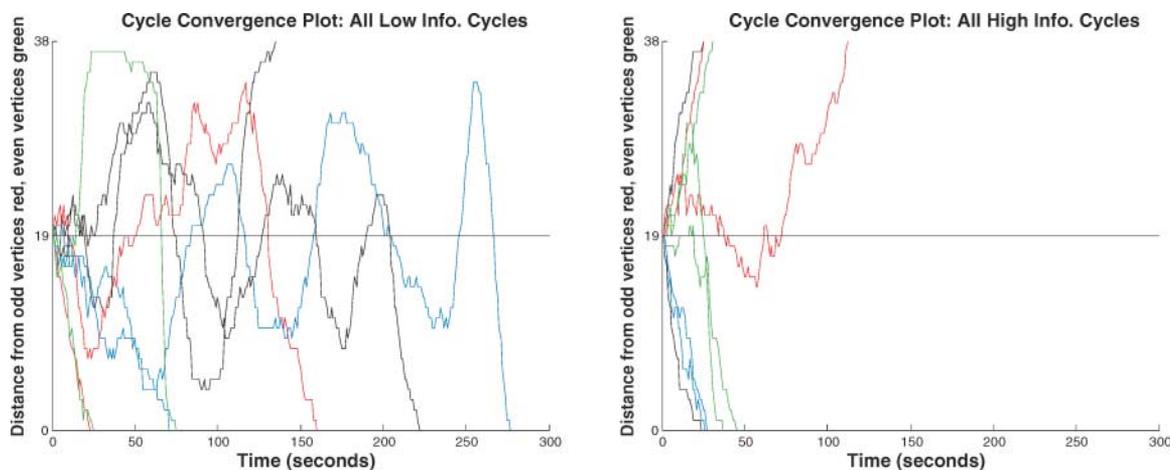
which reduced average distance seems to have a beneficial influence on collective behavior.

It is interesting to compare the influence of network structure on collective human behavior with its influence on natural distributed heuristics. Here, we consider the simplest such heuristic, in which a vertex v is randomly selected among those with a color conflict (that is, a neighboring vertex of the same color). If there are one or more unused colors in the neighborhood of v , one is selected at random for v , thus eliminating the conflict; if not (conflict is inevitable), a new color for v is simply chosen randomly among all possible colors. Because the direct comparison of experimental duration for human populations with computation time is not meaningful, in the rightmost column of Table 1 we report the mean number of color changes (averaged over 10,000 trials) required for this heuristic to successfully color each of the six networks. (We note that for the human subjects, the number of color changes made by the population and experiment duration are proportional and nearly perfectly correlated.) The differences with the collective human behavior are striking, with the order of difficulty within the cycle family exactly reversed (lower average distance increases difficulty for the heuristic), and the preferential attachment networks being relatively easy for the heuristic.

Our final result concerns the effects of varying the locality of information provided to subjects. Our system provided subjects one of three information views during play. Two of these were highly localized, allowing participants to see only their own and neighboring colors, with one of them additionally providing static connectivity information about neighbors (Fig. 2). A minority of experiments used a third information view that allowed subjects to see the global coloring state at all times. In each experiment, the same type of information view was given to all subjects. The most striking

finding again highlights apparently strong differences between our two main categories of networks: Although increasing the amount of information provided sharply reduced solution times for cycle-based networks, it sharply increased it for preferential attachment (Fig. 3). For the cycle-based networks, there are only two possible proper colorings (each a cyclic shift of the other), and subjects seem to have a strong understanding of the collective effort required to rapidly coordinate and converge to one of them when provided with a global view (Fig. 4). In contrast, no such common understanding appears to exist for preferential attachment, and the global information view seems to greatly hamper subjects. All four trials of preferential attachment graphs with global views ended without solution after 5 min. Possible explanations include “information overload” due to the apparent complexity of the networks or their visual layout, combined with the rapid dynamics of the global color selection process. Alternately, it could simply be that time spent by subjects examining the activity in more distant regions of the network distracts them from attending to their own local subtask in the global coordination problem, thus slowing collective solution. With further study, such findings may have implications for areas such as information sharing across large organizations and the design of user interfaces for complex systems for multiparty coordination.

The discussion so far has emphasized collective behavior and performance, but it is also of interest to understand the individual strategies for play used by subjects. Toward this goal, we can apply both the detailed experimental data (which logs every color change by every player, along with their times of occurrence, for each experiment) and the self-reports of the subjects themselves, who were given an exit survey asking them what strategies they employed. These surveys reveal frequent and independent adoption



often oscillates between approaches to the two solutions, whereas in the high-information view experiments (right), there is rapid convergence to one of the two solutions with almost no oscillation.

of certain natural heuristics. These include choosing colors that will result in the fewest local conflicts (mentioned on 11 surveys), as well as attempting to avoid conflicts with neighbors with high connectivity (mentioned on 39 surveys, and obviously applicable to only those two information views that revealed such neighboring information), presumably on the logic that highly connected vertices present the most constrained and difficult problems for subjects. Surveys and the experimental data also revealed a number of instances of signaling behavior by subjects, but here there was less consistency. Some subjects clearly alternated between two colors that were unused in their neighborhood in an attempt to inform neighbors of this fact. Others would alternate between colors in an attempt to call attention to conflicts. Although such signaling behaviors are apparent in the data, it is unclear whether they ever had their intended effects. Many subjects also reported introducing conflicts into their local neighborhood even when they had an available color, in an attempt to perturb the global state from a perceived stasis or local minimum. Even excluding perturbations introduced 2 s or less after the absence of any local conflicts (to account for reaction time), the 38 experiments together had 181 such incidents. This behavior might be viewed as a human analog of the deliberate injection of randomization or “thermal noise” into common optimization algorithms such as simulated annealing. Further work is needed to integrate these observations with statistical methods applied to the experimental data, in order to develop plausible stochastic models of individual behavior in the coloring problem. Ideally such models, when run in multiple independent copies, could predict which networks would be easy or difficult for human populations.

Although the results presented here are suggestive, they are limited in a variety of important ways. The human subject networks were small, a perhaps necessary consequence of the carefully controlled, simultaneous play experiments. It is tempting to contemplate Web-based studies (25) on a much larger scale, which will require addressing incentives, attrition, communication, and many other issues. The network topologies examined here were but a sampling of the rich space of possibilities and recent network formation models. Rather than imposing a chosen network structure on subjects, it would also be interesting to consider scenarios in which the subjects themselves participated in the network formation process, while still allowing some variability of structure. Future work should consider an even wider range of natural collective problems and activities. Candidates include problems of agreement or consensus rather than differentiation, and problems involving the formation of local teams or subgroups specifying certain properties (such as being fully connected or having at least one member of each of a fixed number of types or roles).

References and Notes

- J. Travers, S. Milgram, *Sociometry* **32**, 425 (1969).
- M. Gladwell, *The Tipping Point* (Little, Brown, Boston, 2000).
- S. Milgram, *Psychol. Today* **61**, 60 (1967).
- M. S. Granovetter, *AJS* **78**, 1360 (1973).
- B. A. Huberman, L. A. Adamic, in *Complex Networks*, E. Ben-Naim, H. Frauenfelder, Z. Toroczkai, Eds. (Springer, Berlin, 2004), pp. 371–398.
- O. Sporns, G. Tononi, G. M. Edelman, *Cereb. Cortex* **10**, 127 (2000).
- A. Wagner, D. Fell, *Proc. R. Soc. London B. Biol. Sci.* **268**, 1803 (2001).
- A. Broder *et al.*, in *Proceedings of the Ninth International World Wide Web Conference* (Elsevier, Amsterdam, 2000), pp. 309–320.
- D. J. Watts, *Small Worlds: The Dynamics of Networks Between Order and Randomness* (Princeton Univ. Press, Princeton, NJ, 1999).
- A.-L. Barabási, R. Albert, *Science* **286**, 509 (1999).
- J. Kleinberg, *Nature* **406**, 845 (2000).
- M. R. Garey, D. S. Johnson, *Computers and Intractability* (W. H. Freeman, New York, 1979).
- J. A. Zoellner, C. L. Beall, *IEEE Trans. Electromagn. Compatibil.* **19**, 313 (1977).
- J. Janssen, D. Krizanc, L. Narayanan, S. Shende, *J. Algorithms* **36**, 119 (2000).
- R. M. Karp, in *Complexity of Computer Computations*, Proc. Sympos. IBM Thomas J. Watson Res. Center, Yorktown Heights, NY (Plenum, New York, 1972), pp. 85–103.
- S. Khot, in *Proceedings of 42nd IEEE Symposium on Foundations of Computer Science*, (2001).
- U. Feige, J. Kilian, *J. Comput. Syst. Sci.* **57**, 187 (1998).
- An earlier session of similarly controlled human subject experiments was conducted in September 2005 as a field test for the system. Although we do not report on these here because they used different network structures, they largely foreshadowed our main findings and influenced the design of the January 2006 experiments. The same was true of a graph-coloring exercise held in a University of Pennsylvania class in February 2005 in which communication between subjects was restricted to obey the network structure (but was otherwise unconstrained, and included open discussion and negotiation). Details are in the supporting material (26).
- D. J. Watts, S. H. Strogatz, *Nature* **393**, 440 (1998).
- For each network structure, trials were conducted under three information conditions. For each network, we also examined two different incentive or payment schemes for subjects. The variation across these two design variables yielded $3 \times 2 = 6$ trials of each network. Two networks were given an additional trial to examine potential learning effects, yielding a total of seven trials. Details in (26).
- C. F. Camerer, *Behavioral Game Theory* (Princeton Univ. Press, Princeton, NJ, 2003).
- One session of experiments used collective or global incentives, in which all 38 participants were each paid \$5 for every experiment that resulted in a proper coloring within the 5 min allotted. The other session used individual or local incentives, in which a participant received \$5 if an experiment ended (due either to proper coloring or the termination of 5 min) with their own color being different from all their network neighbors. There were no noteworthy differences between the two incentive schemes in any of the measures discussed.
- There is a downward bias introduced in the average experiment durations because they were capped at 300 s. However, the distribution of unsolved experiments was such that allowing these experiments to continue to solution would only have strengthened the results reported here; see (26).
- J. Kleinberg, in *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2000), pp. 163–170.
- P. Dodds, R. Muhamad, D. J. Watts, *Science* **301**, 827 (2003).
- Materials and methods are available as supporting material on Science Online.
- The authors thank D. Watts, C. Camerer, and T. Senator for their encouragement and for many helpful suggestions regarding this research. We are also grateful to H. Ni for his contributions to the early stages of the project. The work was supported by a grant from the Defense Advanced Research Projects Agency.

Supporting Online Material

www.sciencemag.org/cgi/content/full/313/5788/824/DC1
Materials and Methods
SOM Text

9 March 2006; accepted 15 June 2006
10.1126/science.1127207

Insignificant Change in Antarctic Snowfall Since the International Geophysical Year

Andrew J. Monaghan,^{1*} David H. Bromwich,¹ Ryan L. Fogt,¹ Sheng-Hung Wang,¹ Paul A. Mayewski,³ Daniel A. Dixon,³ Alexey Ekaykin,⁴ Massimo Frezzotti,⁵ Ian Goodwin,⁶ Elisabeth Isaksson,⁷ Susan D. Kaspari,³ Vin I. Morgan,⁸ Hans Oerter,⁹ Tas D. Van Ommen,⁸ Cornelius J. Van der Veen,² Jiahong Wen¹⁰

Antarctic snowfall exhibits substantial variability over a range of time scales, with consequent impacts on global sea level and the mass balance of the ice sheets. To assess how snowfall has affected the thickness of the ice sheets in Antarctica and to provide an extended perspective, we derived a 50-year time series of snowfall accumulation over the continent by combining model simulations and observations primarily from ice cores. There has been no statistically significant change in snowfall since the 1950s, indicating that Antarctic precipitation is not mitigating global sea level rise as expected, despite recent winter warming of the overlying atmosphere.

Global sea level (GSL) has been increasing by 1.7 mm year⁻¹ over the past century (1) and 2.8 mm year⁻¹ over the past decade (2). One of the greatest

uncertainties in predictions of GSL rise is the contribution of the Antarctic ice sheets (3). The Antarctic ice budget is balanced by the buildup of snowfall in the interior and wastage due to