Cue validity, availability, and strategies: Finding experts in networked teams using the Team Map tool

R. Dirk Beer, Harvey S. Smallman, and Frank C. Lacson
Pacific Science and Engineering Group
San Diego, California

Team members in rapidly assembled and often spatially distributed command and control teams have little time to get to know each other. Yet they need access to each other's expertise to work effectively. Current tools such as organizational charts are inadequate because they do not provide sufficient cues to indicate who has what expertise. The Team Map tool is a "souped-up" organizational chart in the form of a table that can reveal information such as previous jobs, training or other attributes in an easy-to-use way. As the first step in its design, we investigate how much information can usefully be shown in the Team Map user interface, given several possible team organization types. Using a theoretical framework of cues and cue foraging strategies, we find that a Team Map with just one additional "summary" cue leads to robust expert-finding performance across organization types, regardless of validity of already available cues. Computer simulations of the foraging task showed participants used less of the information available in the user interface the more cues are provided, and that participants had relatively simple cue-use strategies.

INTRODUCTION

Problem

Command and control (C2) teams spend considerable time fielding information requests and finding information as they plan for, monitor and report the status of missions and operations. A challenge to rapidly assembled C2 teams, such as those standing-up quick-reaction operations centers, is that members of those teams are unfamiliar with each other and the expertise that each of them brings to the team. Yet these teams are expected to be operational almost immediately leaving little time for team members to become familiar with each other. In some organizations, this unfamiliarity is compounded by deliberate staff cross-training (for team resilience and flexibility), resulting in teams where knowledge doesn't reside in easily identified experts, but rather is shuffled across the team. Current tools for navigating teams, such as organizational charts, only indicate role assignments and therefore don't address the problem of shuffled expertise.

Theoretical framework

This lack of knowledge of each others' roles and expertise reduces the effectiveness of team information exchange and decision making (see, for example, Carroll et al., 2003; Alge, et al., 2003). One way to conceive of this problem is that the team has hidden its expertise (Stasser & Titus, 2003). Given that, how can the team's expertise profile be revealed and harnessed for C2 information foraging (Pirolli, 1999)? Here, we apply the cue utility framework of the late Egon Brunswik to the problem (Brunswik, 1956; Goldstein, 2006). In Brunswik's framework, the perceived expertise of team members depends on the validity and availability of cues (Figure 1, right side), and on the cue-use strategies employed by team members (Figure 1, left side). Within this framework, we focused our design and experimentation on making more, and more valid cues available in a way that supports known information-foraging and fast and frugal decision-making strategies (e.g. Pirolli, 2006; Gigerenzer & Goldstein, 1996).

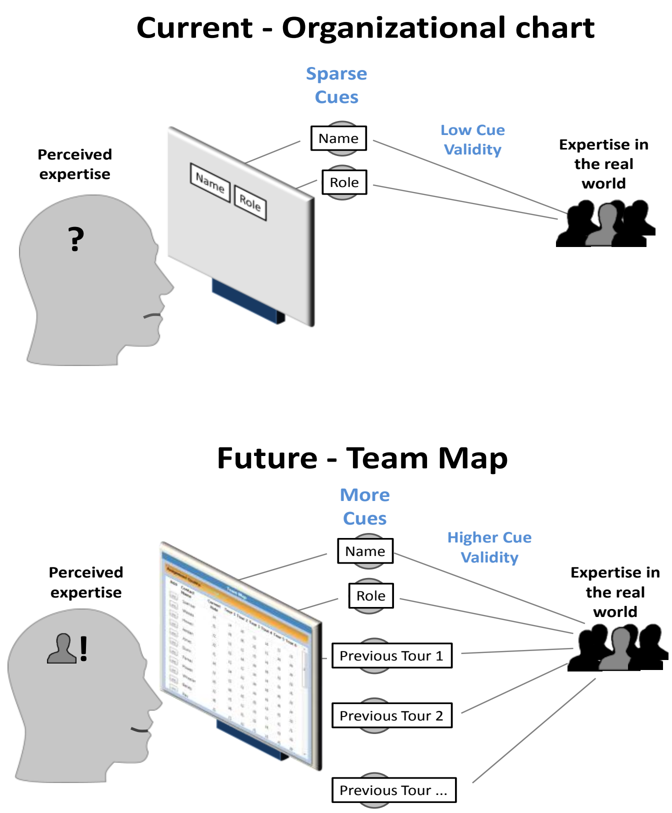

Figure 1. Currently, few cues to expertise are available when searching for hidden experts in a team (top). The Team Map tool would make more cues available in a way that supports known cue-use strategies, in effect revealing the hidden experts (bottom). We investigated the effect of cue availability and cue validity on expert-finding performance and strategies.

Team Map tool

We designed a "Team Map" software tool concept that increases the number and validity of cues such as information on the previous experience and qualifications of team members. The Team Map is designed to support foraging for expertise in the team and therefore improve the exchange of knowledge and use of a team's expertise. The design of this tool is a multi-faceted problem, raising at least two major questions: 1) What determines expertise, and what expertise cues best support mapping of questions to experts? and 2) what is the best way to expose expertise cues in a user interface to supports foraging without bogging users down with low yield information? In this first stage of Team Map design and experimentation we focused on the second question -- how can we best expose expertise cues in a Team Map user interface to support foraging for expertise? We left aside -- for the moment -- the important question of how expertise is best defined and characterized. Our investigation considered the utility of the Team Map for several types of teams by varying the degree to which team members were assigned based on expertise, so that our results can be generalized to a variety of organizations.

EXPERIMENT

Method

Participants. Thirty-six college students or graduates were recruited from www.craigslist.org and were paid $30 for their participation. A second participant sample consisted of 36 active-duty military officers at the Naval Postgraduate School in Monterrey, CA. This group received class credit and was not paid for their participation.

Design. The experiment was a single-participant experience with no co-participants or confederates. The two independent variables were varied in a 3 x 3 within-participants design. The resulting nine conditions were blocked by *Expertise-assignment* to provide a consistent organizational context while the *Team Map type* was varied within each block. Order of the blocks and order of the conditions within the blocks was fully permuted across participants. Each of the nine conditions consisted of three practice trials and nine recorded trials, for a total of 648 recorded trials per condition in the experiment.

Independent variables. We varied the degree of *Expertise assignment* in the team (100% expertise-assigned vs. 75% expertise-assigned vs. 50% expertise-assigned). This is the extent to which team members are assigned based on their expertise, and determines the *validity* of the "current role" cue (for details, see "*Current role assignment*", below). We also varied the cue *availability* in user interface using three *Team Map types* (Figure 2). The Org Chart condition provides the single expertise cue found in existing tools: "current role". The Org Chart + Summary condition contains the current role plus an additional cue, "dominant role", defined as the most frequently served previous tour. Our dominant role cue is an example of an easily calculated, easily used, rough and ready cue. It does not require or use any complex model of expertise, and does not perfectly predict expertise. Nevertheless, we hypothesized that this combination of adequate validity and compact representation of experience would effectively support participant's fast and frugal (Gigerenzer & Goldstein, 1996) cue-use strategies. The Org Chart + Complete conditions contained current role, and a history of the team members' previous six roles. This condition provided the complete detailed information needed to assess expertise, but at the risk of giving participants more information that they could process.

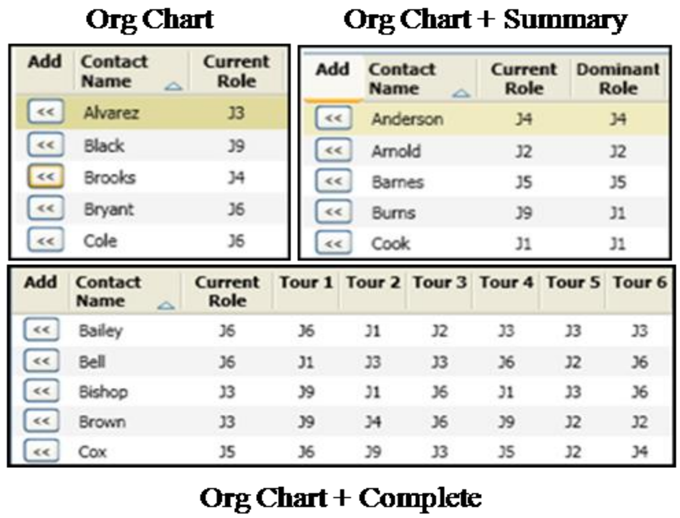
Org Chart

Add	Contact Name	Current Role
<<	Alvarez	J3
<<	Black	J9
<<	Brooks	J4
<<	Bryant	J6
<<	Cole	J6

Org Chart + Summary

Add	Contact Name	Current Role	Dominant Role
<<	Anderson	J4	J4
<<	Arnold	J2	J2
<<	Barnes	J5	J5
<<	Burns	J9	J1
<<	Cook	J1	J1

Add	Contact Name	Current Role	Tour 1	Tour 2	Tour 3	Tour 4	Tour 5	Tour 6
<<	Bailey	J6	J6	J1	J2	J3	J3	J3
<<	Bell	J6	J1	J3	J3	J6	J2	J6
<<	Bishop	J3	J9	J1	J6	J1	J3	J6
<<	Brown	J3	J9	J4	J6	J9	J2	J2
<<	Cox	J5	J6	J9	J3	J5	J2	J4

Org Chart + Complete

Figure 2. Team Map representations and available expertise cues.

Task environment. Participants played the role of a planner in a fictitious Navy command center, whose task was to field email-like information requests by using the Team Map to find the three members of the team most likely to be experts on the emailed questions. To simulate the expert-finding task, participants were provided with an interface consisting of a message box that displays incoming questions, a Team Map which was a sortable table containing information on the current team, and a response box in which they placed their team member selections and received feedback. This response/feedback box was an experimental stand-in for the real-life process of contacting potential experts and assessing the information gained from them.

Staff codes. For the purposes of our present study, each of the incoming questions had a clear one-to-one mapping to one of seven possible "staff codes", which also defined the current role and previous roles of the fictional team members. Staff codes (also called "J-codes" in joint-service U.S. commands) are military designations for the subdivisions of a command staff. For example, J3 is operations, J5 is planning, and J2 is intelligence. For the participants, this meant that a question containing the keyword "intelligence" was clearly a "J2" question, and could best be answered by fictional team members with previous J2 tours. No other staff code experience would be relevant in this case.

Fictional team. The team in which participants were foraging for experts consisted of 28 *virtual* team members, each of whom had a randomly determined history of six previous J-code tours. These six previous roles or tours determined each team member's expertise on any question, calculated by the method described in "*Expertise score*", below. In addition, each team member was assigned a current role as described in "*Current role assignment*", below. The current role did *not* contribute to their expertise, since the scenario was that the command center had just been assembled. Instead, this expertise predictor depended on the expertise-assignment condition.

Expertise score. To provide a quantitative measure of participant performance, we used a model based on the forgetting function to generate an expertise score as a function of relevant experience. Tours matching the question J-code were weighted with an exponential time decay to model both forgetting and loss of relevance of older experience. Weighted experience was then added. So, for example, a team member with two recent relevant tours could have a higher expertise score on a question than another team member with three relevant tours long ago. This rough model produced sensible distribution of expertise scores and allowed us to record a quantitative measure of participant performance: score on any trial was the sum of the expertise scores of the three experts chosen from the Team Map.

Current role assignment. The current role assignment determined to what extent team members are assigned based on their expertise, and therefore how much of the team's expertise would be visible with currently used tools (Org Chart condition). We first used the Kuhn-Munkres algorithm to assign team members to available roles in a way that maximized total *visible* team expertise, meaning that team members' current roles indicated their area of maximum expertise. A team where all team members were assigned in this way was a 100% expertise-assigned team where the current role cue could be trusted to help find the most qualified experts. We also created more realistic 75% and 50% expertise-assigned teams by randomly shuffling 7 of 28 or 14 of 28 of the team members, respectively, into non-optimal roles. This in effect *hides* their expertise, because their current roles do not indicate their expertise. In these teams, the current role cue is less valid. For example, someone in a J2 (intelligence) role would have only a 75%, or a 50%, chance of being an expert in intelligence. According to informal interviews with active duty military officers, 50% expertise-assigned teams are typical, and we consider this the real-world baseline.

Procedure. For the experimental task, a trial was initiated by an incoming question. This question contained a keyword associated with only one staff code. After reading the question, participants decided which staff code expertise was needed. To ensure that participants correctly identified the staff code associated with each question, they were given a reminder sheet of staff "J" codes, their descriptions, and their associated keywords. After determining the appropriate staff code, participants then interacted with the Team Map to select three team members they judged to have the most staff code expertise. After finalizing their team member selection, participants received feedback on their selection in the form of a text and graphical representation of each expert's expertise score. This procedure lasted about 75 minutes.

Dependent variables. The expertise found on each question was measured by summing the expertise scores of the three chosen experts. In addition, we measured time per question, defined as the length of time needed to read the question, to forage for and select three experts, and to submit the selection.

Results

Additional cues. Providing additional cues to expertise increased the amount of expertise found regardless of the validity of the existing "current role" cue (team degree of expertise assignment). The effect was greatest when the existing cue validity was lowest -- in the realistic 50% expertise-assigned condition. This increase in expertise found came at a time cost of as little as 2 seconds to much as 15 seconds per trial (see Figure 3).

Team Map cue types. We investigated two Team Maps with additional cues: an Org Chart + Summary, which contained the rough and ready "dominant role" cue, and an Org Chart + Complete, which contained a complete tour history. In typical teams (Figure 3, solid line) participants found much expertise using these two Team Map types, with the largest increase due to the addition of the single "dominant role" cue. Another difference between the Team Map types is the time to use them – it took slightly longer than baseline to find three experts with the Org Chart + Summary (12 s versus 7 s), and but three times longer than baseline (21 s versus 7 s) using the Org Chart + Complete.

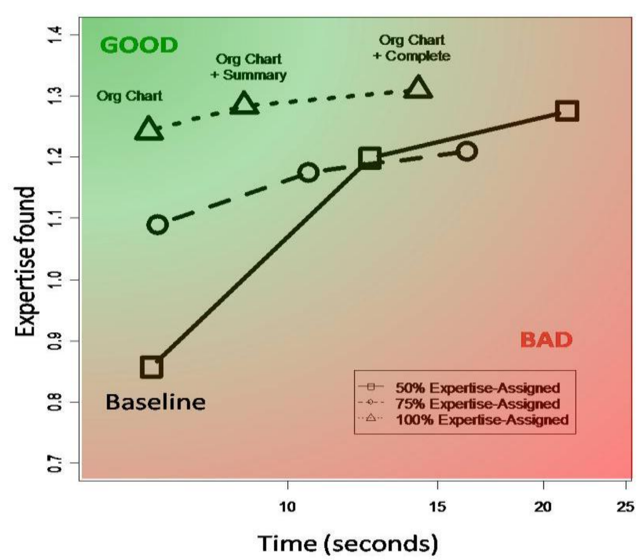

Figure 3. Adding cues increases the expertise found per trial, at some cost in the time required. The benefit of adding cues depends on the validity of the "current role" cue, which is determined by the degree to which the team is assigned based on expertise. The "Org Chart + Summary" Team Map maintains good performance across all team types, providing increase in expertise found with modest increase in task time. Data points represent the mean of trials across participants; 95% confidence intervals are approximately the marker size.

Computer-simulated participant. The performance data above reflects the effect of both the cue environment (cue availability and validity) and of the participants' cue use efficiency and strategies -- both sides of the Brunswik framework in Figure 1. To disentangle the two we created a computer simulated participant. This simulated participant runs through the experiment trial-by-trial, applying a cue-use strategy specified by the authors. One eight-minute run of the simulation applies the specified (potentially condition-specific) algorithm to all trials of the experiment and produces a dataset of expertise scores of the same size as our participant data.

Efficiency of information use. How efficiently do participants use additional cues? We compared participant performance to a simulated participant that used an optimal condition-appropriate set of strategies. This simulated participant is an "ideal observer" (Barlow & Lal, 1980), or ideal information processor, that uses all information available in the user interface and combines that information optimally. The set of optimal cue-use strategies was found using trial and error to determine which of a limited set of plausible strategies result in the best possible simulated performance.

Comparison of real to simulated ideal participant performance shows that in the Org Chart (baseline) condition, participants use all available information (Figure 4). When the additional "dominant role" cue is added, participants use slightly less (95%) of the available information. When given the complete set of tour history cues (Org Chart + Complete condition) participants used only about 78% of the information available in those cues.

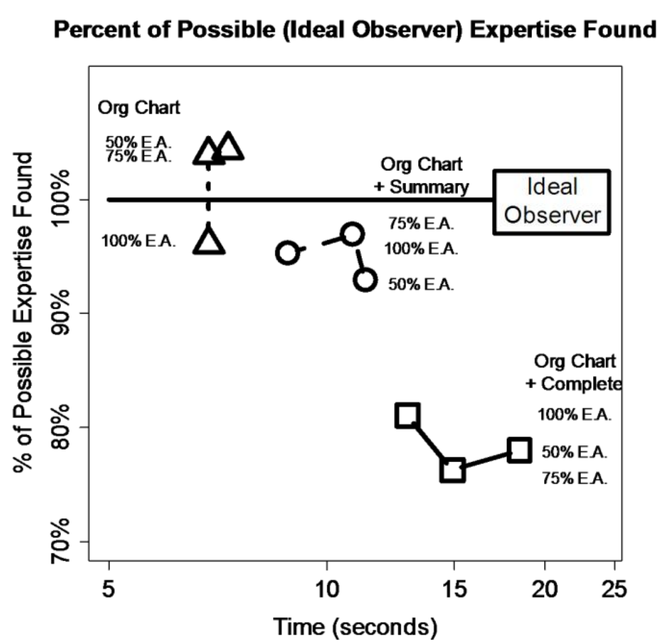

Figure 4. Adding cues leads to a reduction in the efficiency with which those cues are used. Participants search for expertise as well as a simulated ideal observer in the Org Chart and Org Chart + Summary condition (100% and 95%, respectively). However, when presented with a complete set of cues, participants only find 78% as much expertise as the ideal observer. Data points represent the mean of trials across participants; 95% confidence intervals are approximately the marker size.

Participant strategies. We also compared real participant data to data generated by plausible *imperfect* cue-use strategies. By simulating progressively more complex strategies, we were able to determine the least complex strategies necessary to explain participant data, and thereby get a handle on how participants use information in the user interface.

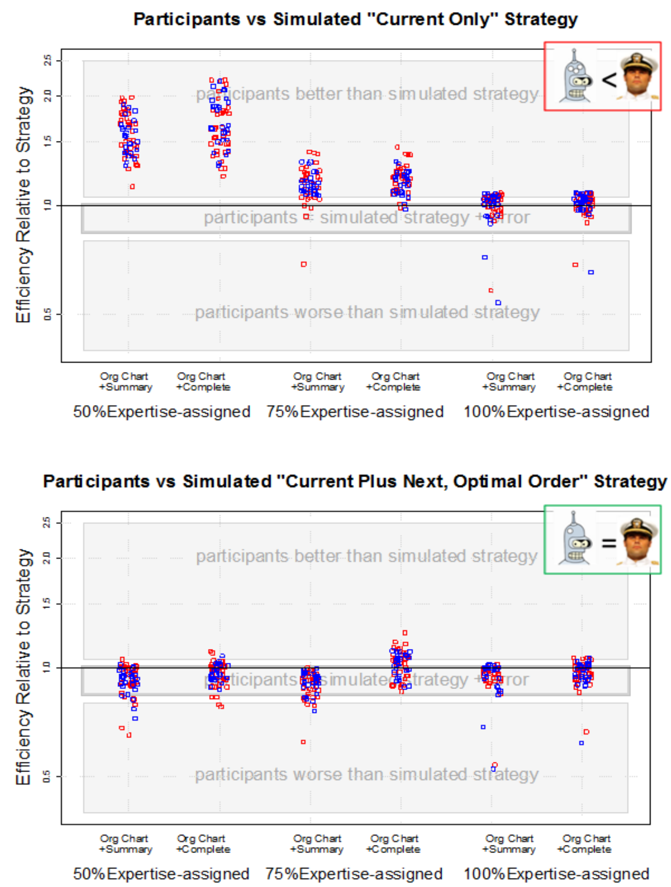

Figure 5. Simulation of progressively more complex cue-use strategies shows that a relatively simple strategy adequately models performance of participants across all experimental conditions. Participants find more expertise than a simulated single-cue strategy (efficiencies > 1, top), but about the same amount of expertise as a two-cue, cue-validity sensitive strategy (efficiencies ~= 1, bottom). Data points represent the mean within each condition for each participant. Blue markers represent Naval Postgraduate School participants; red symbols represent college-educated civilian participants.

In Figure 5, we show a comparison of real participant data to simulated data for two plausible participant strategies. A limited strategy of only using the "current role" and ignoring additional cues is insufficient to explain participant performance (Figure 5, top). On the other hand, a strategy that considers the "current role" cue and the next available cue (either "dominant role" or "tour 1" depending on Team Map type) better explains participant performance (Figure 5, bottom). This strategy uses the most valid of the two cues for the initial search, and the less valid cue to narrow the choices to three.

We conclude that participants likely used this simple strategy for all but one condition because the strategy predicted their data and few, if any, plausible alternative strategies exist in this well defined and bounded task and cue environment. Participants appear to use simple strategies even in conditions when more complex strategies are potentially beneficial.

DISCUSSION

Rapidly stood-up command and control teams are common for military, humanitarian assistance and disaster relief, and emergency response. In most cases, a large proportion of team members are assigned to roles based on factors other than their expertise, or they have expertise not related to their assigned role. Team member expertise is hidden. The Team Map tool is designed to provide additional cues to team members' expertise, making that expertise available to the team and improving team performance.

In this experiment and analysis, we concentrated on the user interface: what is the effect of increasing the availability of cues, given that the validity of existing cues may vary from team to team? We found that providing additional cues to expertise increased the amount of expertise found regardless of the validity of the existing cue (current role), but that additional cues were more helpful when existing cue validity was low. Since the degree of expertise-based assignment is what determines the validity of the current role cue, we also conclude that additional cues are helpful no matter how the team is assigned.

We also investigated participant strategies of cue use in the user interface. How efficiently are additional cues used, and what strategies are applied to them? An ideal observer computer simulation showed that adding more than a few (possibly more than just one) extra cues meant those additional cues are used inefficiently. A simulation of progressively more complex strategies indicated that participants use relatively simple strategies, even when complex cues are available. The parameter-free computer strategy simulation we created was a powerful tool in this well-defined task and cue environment. Conclusions about strategy that are based on the simulation could be validated in future studies using an independent measure such as eye tracking.

The present study concentrated on the user interface, and on participants' abilities and limitations in using the cues presented. Another component of the design of a Team Map tool will be to understand how questions can be mapped onto knowledge in a specific domain, and which team member attributes are the best cues to help find an expert with that knowledge. We are currently working with military lawyers to develop an ontology of knowledge in the military operational law domain, and to develop a Team Map with cues that provide indexes into that domain, while building on the user interface principles found here.

ACKNOWLEDGEMENTS

This work was sponsored by the Office of Naval Research, program manager Dr. Gerald Malecki, through SPAWAR Systems Center Pacific, managing scientist, Dr. Michael B. Cowen.

REFERENCES

Alge, B., Wiethoff, C., & Klein, H. (2003). When does the medium matter? Knowledge-building experiences and opportunities in decision-making teams. *Organizational Behavior and Human Decision Processes*, 91(1), 26-37.

Barlow, H., & Lal, S. (1980). The Absolute Efficiency of Perceptual Decisions. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* (1934-1990), 290(1038), 71-82.

Brunswik, E. (1956). *Perception and the representative design of psychological experiments.* Berkeley: University of California Press.

Carroll, J., Neale, D., Isenhour, P., Rosson, M., & McCrickard, D. (2003). Notification and awareness: synchronizing task-oriented collaborative activity. *International Journal of Human-Computer Studies*, 58(5), 605-632.

Gigerenzer, G., & Goldstein, D. (1996). Reasoning the Fast and Frugal Way: Models of Bounded Rationality. *Psychological Review*, 103, 650-669.

Goldstein, W. (2006). Introduction to Brunswikian Theory and Method. In A. Kirlik (Ed.). *Adaptive Perspectives on Human-Technology Interaction: Methods and Models for Cognitive Engineering and Human-Computer Interaction*. New York: Oxford University Press.

Pirolli, P. (2006). The Use of Proximal Information Scent to Forage for Distal Content on the World Wide Web. In *Adaptive Perspectives on Human-Technology Interaction: Methods and Models for Cognitive Engineering and Human-Computer Interaction*: Oxford University Press, USA.

Pirolli, P., & Card, S. (1999). Information Foraging. *Psychological Review*, 106, 643-675.

Stasser, G., & Titus, W. (2003). Hidden Profiles: A Brief History. *Psychological Inquiry*, 14(3&4), 304-313.