# The network science of philosophy

**Cody Moser**[1,2,✉], **Alyssa Ortega**[1,3], **and Tyler Marghetis**[1,✉]

[1]Cognitive and Information Sciences, University of California, Merced
[2]School of Collective Intelligence, Université Mohammed VI Polytechnique
[3]Psychology, University of Wisconsin Madison

**Philosophy is one of the oldest forms of institutional knowledge production, predating modern science by thousands of years. Analyses of science and other systems of collective inquiry have shown that patterns of discovery are shaped not only by individual insight but also by the social structures that guide how ideas are generated, shared, and evaluated. While the structure of scientific collaboration and influence can be inferred from co-authorship and citations, philosophical influence and interaction are often only implicit in published texts. It thus remains unclear how intellectual vitality relates to social structure within philosophy. Here, we build on the work of historians and sociologists to quantify the social structure of global philosophical communities consisting of thousands of individual philosophers, ranging from ancient India (c. 800 BCE) to modern Europe and America (1980 CE). We analyze the time-evolving network structure of philosophical interaction and disagreement within these communities. We find that epistemically vital communities become more integrated over time, with less fractionated debate, as a few centralizing thinkers bridge fragmented intellectual communities. The intellectual vitality and creativity of a community, moreover, is predicted by its social structure but not overall antagonism among individuals, suggesting that epistemic health depends more on how communities are organized than on how contentious they are. Our approach offers a framework for understanding the health and dynamism of epistemic communities. By extending tools from collective intelligence to the study of philosophy, we call for a comparative "science of philosophy" alongside the science of science and the philosophy of science.**

**Collective Intelligence | Philosophy | Networks | Network Dynamics | Intellectual History**

**Correspondence:** ***cmoserj@gmail.com and tyler.marghetis@gmail.com***
***Department of Cognitive and Information Sciences, University of California, Merced, CA 95343, USA***

## Introduction

Philosophy is one of the oldest forms of institutional knowledge production, predating the emergence of modern science by thousands of years. From ancient times to the modern era, in societies across the globe, philosophical debate has generated distinct schools or communities, each characterized by their own influential thinkers, intellectual focus, and sociopolitical origin (1). Philosophical debates are recursive, open-ended, and can sometimes persist across generations. Viewing philosophical communities as cognitive ecologies (2, 3) can offer insight into why some intellectual communities seem especially vital. From this perspective, the epistemic vitality of a philosophical system should depend not on individual genius, but on the interpersonal interactions that give rise to collective thinking (4–6). Recent work in the science of science (7, 8) has begun to quantify these collective processes in scientific systems, but philosophy, with its lack of clear empirical resolutions and deeper historical record, presents an opportunity to examine the social processes involved in purely intellectual debate. What are the social architectures that give rise to philosophical thought, and how do they sustain or constrain its development?

Comparative philosophers and sociologists of knowledge have developed several historical frameworks to explain the dynamics of philosophical traditions. The 19th century theologian and historian John Henry Newman (9) likened philosophical systems to living organisms, gaining maturity through sustained internal debate. The 20th century philosopher Richard Rorty (10) described philosophy as an open-ended system, sustained by "strong poets" who periodically reinvent its discourse by introducing new metaphors for enduring problems. At the turn of the current century, the sociologist Randall Collins (11) argued that philosophical traditions are structured by competition for collective attention among individual thinkers and their apprenticeship lineages. Echoing these perspectives, research in social epistemology and collective intelligence has emphasized several structural factors that support epistemic vitality. These include the importance of system-wide disagreement and diversity in generating new ideas (12, 13), the influence of radical disruptors on the trajectory of intellectual communities (14, 15), the innovative potential of peripheral actors within networks (16, 17), and the role of individuals who bridge otherwise disconnected groups in enabling synthesis and discovery (15, 18).

Each of these perspectives suggests that philosophical traditions reflect both the contributions of individual thinkers and the broader communities that structure their interactions. It remains unclear, however, which specific features of a community are most closely associated with epistemic vitality. One family of views holds that vitality arises from individual-level tension, disagreement, or diversity (13, 19). According to these accounts, the most productive communities are those marked by persistent disagreement (20) or by a high degree of individual autonomy, where thinkers are relatively isolated and free to pursue independent lines of inquiry (21, 22). Alternatively, epistemic vitality may depend on a balance between tight-knit collaboration and intellectual distance (11, 23). Productive communities, for instance, have been proposed to consist of a small number of highly integrated groups that generate intellectual energy through competition and debate (11). Synthesizing these perspectives, recent work emphasizes the dual importance of innovation within independent sub-communities and the role of 'bridg-

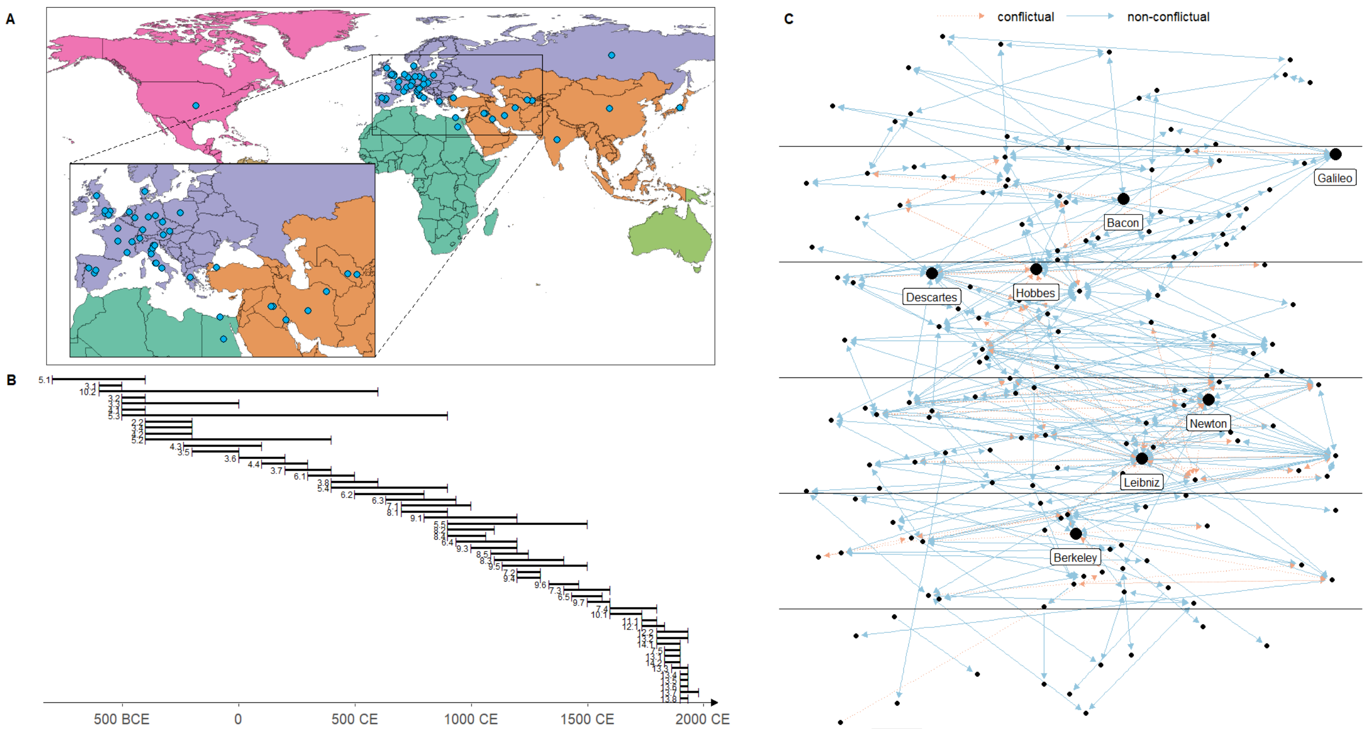


**Fig. 1.** Global philosophical communities throughout history. **(A)** Geographic locations of all philosophical networks included in the study. Inset shows the region of Europe and the Middle East with the highest density of philosophical communities, although there are communities from throughout Africa, America, Asia, and Europe. **(B)** The timeline of each philosophical community, with lines indicating their start and end dates. Numbers next to the lines correspond to figures in (11); the location and identity of each numbered network is available in the Supplemental Materials. **(C)** An example of one network, described in (11) as, "The European Network: The Cascade of Circles," which lasted from 1600 to 1735 CE. Nodes represent individual philosophers (key figures labeled); edges represent intellectual relationships (solid blue = non-conflictual; dotted orange = conflictual). Each network was divided into non-overlapping temporal periods of equal duration (indicated by horizontal black lines), following the temporal divisions introduced in Randall Collins's visual representations of the history of philosophy (11).

ing nodes' — individuals or institutions that connect otherwise disconnected subgroups — to facilitate collective innovation and prevent stagnation (23, 24). Viewed this way, productive communities should derive their vitality from structural convergence, in which different subgroups are brought together, thus reshaping the collective dynamics of disagreement.

Here, to investigate the social systems that generate epistemic vitality within philosophy, we analyzed a dataset of philosophers ($N = 3187$ philosophers) and their intellectual relationships ($N = 5415$ edges) across multiple historical communities ($N = 55$ networks) spanning over two millennia (Fig. 1). We digitized these social networks from a large-scale sociological analysis of the history of philosophy conducted by the sociologist Randall Collins (11). Within these networks — which often span multiple generations — individuals are connected through various relationships, such as student-teacher ties, alliances, or rivalries. Each community is thus characterized by a social network that emerges over time from a mix of differing relationships. These networks vary in complexity, with some containing over a hundred nodes (philosophers) and hundreds of edges (relationships) representing intellectual relationships between them.

Using this corpus of networks, we quantified the structure of philosophical social networks to examine how different patterns of connection may shape the dynamics of philosophical thought. This approach allowed us to investigate how a network's structure relates to the epistemic vitality of the community. We analyzed the temporal evolution of each network, investigating how structure emerges over time. Finally, we examined how individual actors and broader community structures have contributed to epistemic vitality throughout the history of philosophy, showing that the temporal trajectory of a community's network structure predicts its epistemic vitality.

## Results

**Quantifying the social networks of philosophical debate.** Based solely on the local structure of networks, foundational figures from the history of philosophy stood out as highly influential relative to their peers (Fig. 2A). For each philosopher, we operationalized their relative influence within a network as their normalized degree (i.e., the number of links to other philosophers, z-scored within each community). As a few illustrative examples, highest-degree nodes included Confucius, Socrates, Husserl, John Stuart Mill, and Leibniz; Kant was second only to Schiller.

The contributions of individual philosophers, however, will depend on the larger communities in which they are embedded. We thus quantified the social structure of philosophical debate across history by characterizing each community's social network using a set of structural measures. These included macroscale properties (e.g., modularity, number of k-cores, sparsity), microscale features (e.g., peak centrality, cyclicity, average clustering), and global distance-based fea-

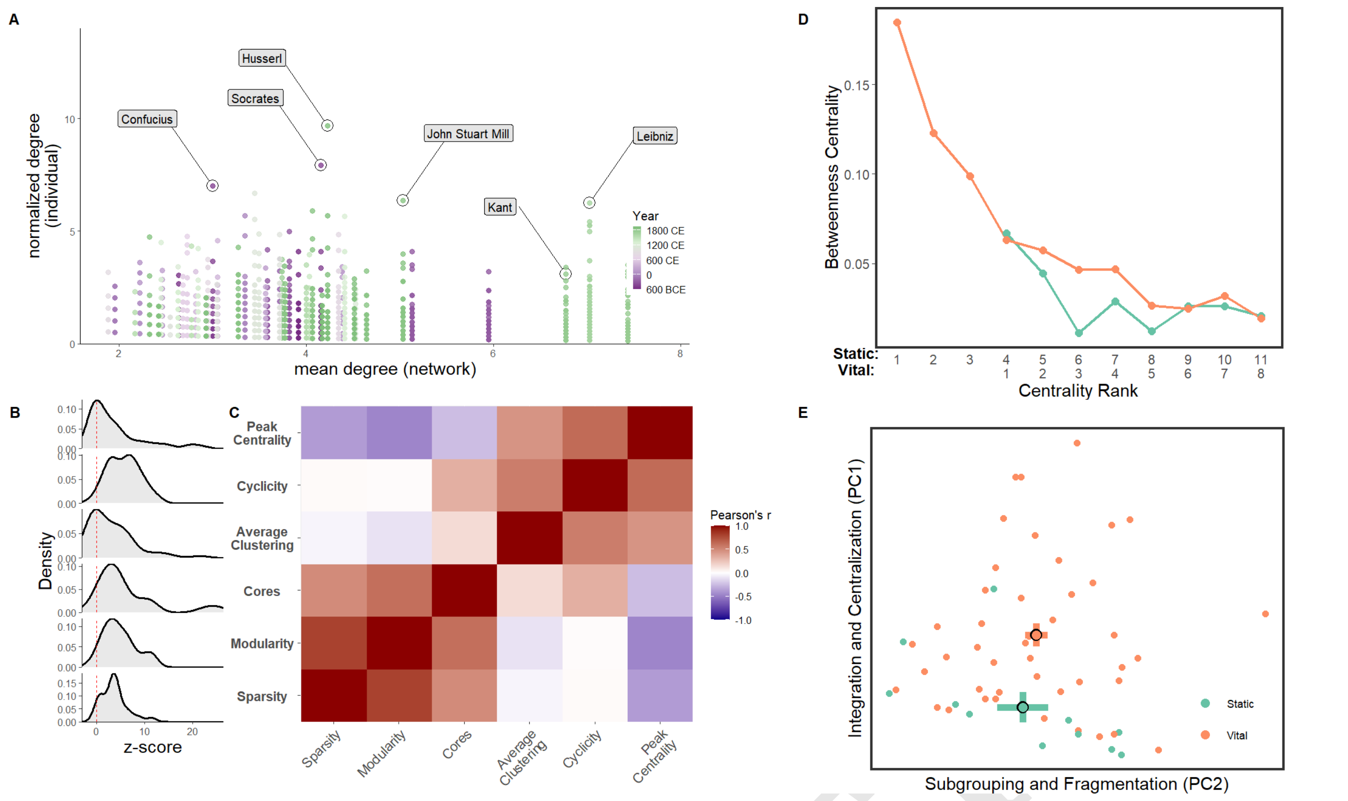


**Fig. 2.** Quantifying the structure of philosophical communities. **(A)** Influential figures were highly connected. Each point ($N = 3187$) represents one philosopher's normalized degree (y-axis) within their community, plotted against the mean degree of that community (x-axis). Columns of points represent philosophers within a particular community. Historically transformative philosophers (annotated) appear consistently as high-degree nodes relative to their peers. (Color indicates each community's mean year.) **(B)** Philosophical communities were highly structured. Density plots of key network measures. Measures for each network were z-scored relative to the network's null distribution derived from time-respecting null networks (see Methods for details). Red lines at 0 thus represent the expected value of the time-respecting null networks. (See Figure S2 for additional network measures). **(C)** Correlation heatmap of the same key network measures shown in (B). (See Figure S1 for additional network measures.) **(D)** Rank ordering of agent betweenness centrality in vital and static networks. For each network, the agent with the nth highest betweenness centrality was identified and points represent the group average at each rank across all networks. Static group points are slightly right-shifted, highlighting convergence between network types. Color indicates the network's epistemic vitality.) **(E)** Vital and static communities differed in network structure. Each network was embedded in a 2-dimensional structure space, created using Principle Component Analysis. PC1 primarily captures network integration and centralization, while PC2 relates more to community fragmentation. Color indicates the network's epistemic vitality. Crosses represent means ± standard errors.

tures (e.g., average path length, network diameter). These measures tended to group into three correlated clusters related to community structure, network centralization, and path length (Figure 2C; Figure S1; see Table S1 for descriptions of each measure and their correlations).

To contextualize these measures, we generated time-respecting null models by randomly rewiring each network's edges while preserving temporal structure (see Methods for details). Philosophical networks generally exhibited greater community structure and network centralization than expected by chance. Philosophical communities exhibited significantly greater modularity and sparsity than their null counterparts, with 96.2% of networks scoring above the median of their respective null distributions for both measures. Binomial tests confirmed that the proportion of networks exceeding their null counterparts was significantly greater than chance for both modularity ($p < 0.001$) and sparsity ($p < 0.001$), suggesting that philosophical communities tend to be both well-partitioned and loosely connected. Real networks also contained more communities (79% above nulls, $p < 0.001$), indicating a tendency toward intellectual factionalism; were more cyclical (96%, $p < 0.001$), reflecting discursive feedback between scholars; and more clustered (68%, $p = 0.006$), a sign of tight-knit local interactions. Additionally, peak centrality was higher in 70% of networks ($p = 0.003$), suggesting the presence of individuals who were especially influential within their respective systems.

**The network structure of epistemic vitality.** We next investigated whether certain network structures facilitate philosophical vitality. We classified philosophical communities into periods of epistemic vitality or of stasis, based on the account of global philosophy in Collins' *Sociology of Philosophies* (11), which synthesized historical research on philosophical communities throughout history. Epistemically vital periods were characterized by the emergence and synthesis of new ideas, while static periods either maintained existing ideas or involved unresolved disputes that did not lead to synthesis (Figure S3). We validated this classification of epistemic vitality using a pretrained machine learning model (see Methods); in the Supplemental Materials (Supplementary Note S4), we show that our results are robust to whether we use human or machine classification of epistemic vitality.

Overall, vital and static communities did not differ significantly in size ($M_{static} = 62$ vs. $M_{vital} = 57, t_{53} = 0.4, p = .69$) or number of edges ($M_{static} = 77$ vs. $M_{vital} = 104$,

$t_{53} = -0.9, p = .35)$ . The question, then, is whether these nodes and edges were *organized* differently.

Several theoretical accounts suggest that a community's epistemic vitality is shaped by its social structure, including factors such as interconnectedness (23), levels of disagreement (20), centralizing individuals (18), and productive fragmentation (13). We evaluated these proposals using measures of the social network's structure: connectivity and average degree (representing interconnectedness of each community), peak centrality (indicating the existence of centralizing figures), and the scaled number of sub-communities (capturing fragmentation versus integration). As these they were correlated (Fig. 2), we also reduced these variables to two dimensions using Principal Component Analysis and examined whether networks differing in vitality occupied distinct regions in this reduced feature space.

To test the role of interconnectedness, we examined network connectivity and degree. Vital communities differed reliably from static ones (Fig. 2D, E). Specifically, individual philosophers in vital communities were more integrated across the whole network (network connectivity: $M_{static} = 0.18, M_{vital} = 0.49 t_{53} = -4.58, p =< 0.001$) and were more locally connected (average degree: $M_{static} = 1.33, M_{vital} = 1.88 t_{53} = -3.87, p =< 0.001$).

To test the role of centralizing figures, we examined the peak centrality of philosophers within each community. As predicted, vital communities contained individuals who were more centralizing (peak centrality: $M_{static} = 0.07, M_{vital} = 0.18 t_{53} = -3.88, p =< 0.001$). Inspection of the rank-centrality distribution confirmed that vital communities were characterized by a few extreme outliers who were highly central; in the absence of those highly central individuals, the distribution of node centrality was qualitatively the same for vital and static communities (Fig. 2D).

Epistemically vital communities were thus more globally integrated, more locally connected, and contained extreme individuals who were centralizing. These three features all loaded highly on the first principle component, which reliably distinguished epistemically vital communities from static ones (linear regression predicting PC1: $b = 1.624 \pm 0.543$ SE, $p = 0.004$; Fig. 2E).

By contrast, vital and static communities did not differ in the tendency of philosophers to disagree among themselves or fragment into factions. To test the role of antagonism or disagreement, we examined the tendency to disagree within each community (ratio of conflictual to non-conflictual edges). Vital and static communities did not differ in antagonism ($M_{static} = 0.45, M_{vital} = 0.26 t_{53} = 1.41, p = 0.19$). This suggests that it is not the sheer amount of disagreement that distinguishes vital communities, but rather how that disagreement is structured. To test the role of fragmentation, we measured the tendency of communities to cluster into sub-communities (number of sub-communities, scaled by network size). Vital and static networks did not differ along this dimension (: $M_{static} = 0.19, M_{vital} = 0.16 t_{53} = 1.23, p = 0.24$). This measure loaded highly on the second principal component, which did not reliably distinguish vital communities from static ones (linear regression predicting PC2: $b = 0.148 \pm 0.285$ SE, $p = 0.61$). Epistemic vitality may depend less on how much communities disagree or fragment into factions, and more on how disagreements and sub-communities are organized to support productive tensions at the collective level.

Thus, while philosophers in vital and static communities were equally likely to fragment into distinct sub-communities (PC2 in Fig. 2D), in vital communities the overall structure was more integrated and centralized (PC1 in Fig. 2D). This is consistent with theories of epistemic vitality that foreground the importance of structural coherence and centralizing figures, who can bridge communities and facilitate the circulation of ideas across otherwise disconnected subgroups (18, 23).

**The temporal dynamics of epistemic vitality.** We next investigated the temporal emergence of these networks. To analyze the temporal trajectory of network structure, networks were divided into non-overlapping temporal periods (Fig. 1c), following the temporal divisions introduced in Randall Collins's visual representations of the history of philosophy (11). Within each network, these periods were of equal duration, typically corresponding to approximately one generation, although sometimes of longer duration for communities that persisted for hundreds of years. To examine how centrality and community structure shape the dynamics of agreement and disagreement in systems with differing levels of epistemic vitality, we analyzed how these distinguishing measures evolved over time, including the scaled number of communities, peak centrality, average degree, and the ratio of disagreement to agreement. To investigate the temporal evolution of these measures within each network, for each measure we fit a mixed effects model, with fixed effects of epistemic vitality, time within each community, and their interaction (Table 1; see Methods for details).

Our analysis above of holistic community structure revealed that epistemic vitality was associated with greater network integration. We thus investigated how integration emerged over the life of a philosophical community (Table 1; plotted in Supplementary Note 5.). We find that, at their origin, vital and static networks did not differ in integration, as measured by connectivity (effect of epistemic vitality on connectivity at communities' origination ($b = -0.01 \pm 0.08$ SEM, $t = -0.83, p = 0.94$). Over time, however, the connectivity of non-vital communities remained stable (change over time in non-productive communities: $b = -0.04 \pm 0.13$ SEM, $t = -0.29, p = 0.77$), while epistemically vital communities became more connected (productive communities: $b = 0.30 \pm 0.07$ SEM, $t = 4.34, p < 0.001$). At their culmination, therefore, epistemically vital communities were more connected than static communities (effect of epistemic vitality on connectivity at the communities' culmination: $b = 0.33 \pm 0.13$ SEM, $t = 2.60, p = 0.013$).

There was a similar pattern for community density as measured by average degree (Table 1 and Fig. 3A, bottom left), with no significant difference between productive and non-productive communities at the beginning ($b = 0.29 \pm$

**Table 1.** Mixed effects model results predicting network structural measures over time.

| Predictor | Connectivity | Avg Degree | Peak Centrality | Sub-Communities | Disagreement Ratio |
|---|---|---|---|---|---|
| Intercept | 0.23 (0.07) [3.22] | 0.61 (0.14) [4.44] | 0.04 (0.03) [1.22] | 0.66 (0.06) [11.27] | 0.63 (0.26) [2.42] |
| Vitality (Start) | -0.01 (0.08) [-0.08] | 0.29 (0.16) [1.81] | 0.06 (0.04) [1.64] | -0.06 (0.07) [-0.87] | 0.18 (0.30) [0.62] |
| Vitality (End) | **0.33 (0.13) [2.60]*** | **0.82 (0.22) [3.75]***** | **0.13 (0.04) [3.31]**** | **-0.23 (0.06) [-4.11]***** | -0.07 (0.16) [-0.42] |
| Time | -0.04 (0.13) [0.29] | 0.14 (0.21) [0.65] | 0.01 (0.05) [0.25] | -0.07 (0.08) [-0.87] | -0.12 (0.33) [-0.36] |
| Prod × Time | **0.34 (0.14) [2.35]*** | **0.54 (0.25) [2.18]*** | 0.07 (0.06) [1.16] | **-0.18 (0.09) [-2.02]*** | -0.25 (0.37) [-0.67] |

Note: Standard errors and t-statistics are in parentheses and square brackets, respectively. Statistical significance is indicated by asterisks (*, < .05; **, <.01, ***, <.001).

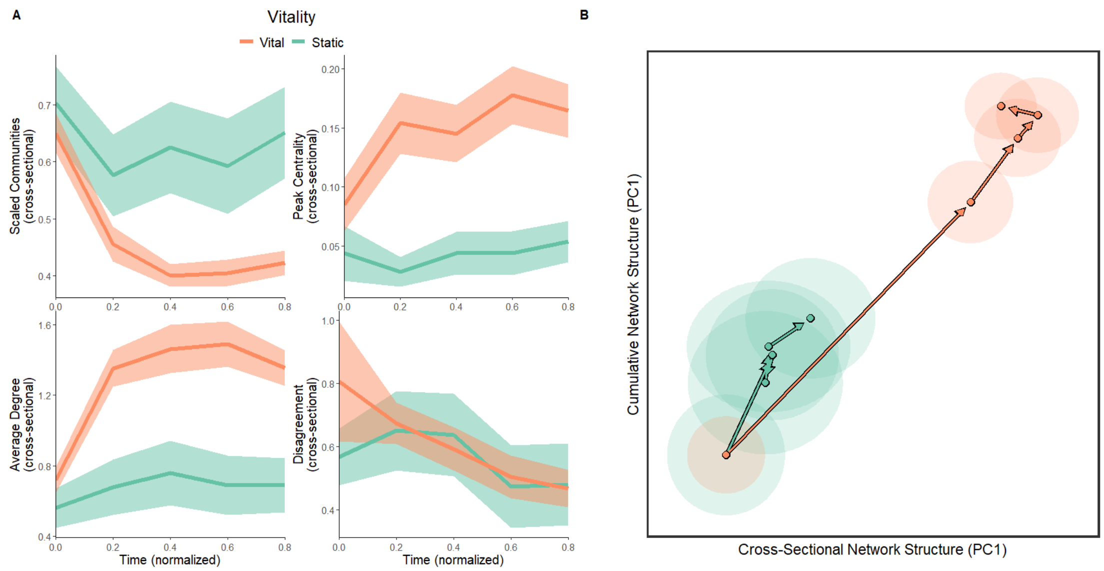


**Fig. 3.** The temporal emergence of epistemic vitality. **(A)** Temporal evolution of network structure within cross-sectional time periods. Vital and static communities did not differ at their origin. Over time, however, vital communities became less fractionated (scaled communities), more interconnected (degree), and contained more highly central individuals (peak centrality). By contrast, vital and static communities had similar amounts of disagreement throughout their existence. (Lines = means; shaded ribbons = standard errors.) See Figure S5 for additional network measures. **(B)** Temporal differentiation of the network structure of vital (orange) and static (green) communities. The x-axis shows the network structure (integration and centralization) within each non-overlapping, cross-sectional time period. The y-axis shows the network structure accumulated from a community's origin up until that time period. This allows us to visualize the temporal evolution of network structure, both cross-sectional (x-axis) and cumulative (y-axis). Static communities show little change over time in the integration and centralization of philosophers within a given time period (x-axis) and only a slight increase in cumulative structure. Vital communities, by contrast, grew more integrated and centralized both with each subsequent time periods (x-axis) and overall (y-axis). (Points = means. Shaded circles = standard errors.)

0.16 SEM, $t = 1.81, p = 0.080$), but significantly increased in average degree in productive communities ($b = 0.68 \pm 0.12$ SEM, $t = 5.5, p < 0.001$), so that productive communities had higher degree at their culmination. In other words, while both static and vital communities *begin* with similar measures of local and global connectivity, vital communities became more integrated over time, with individual nodes being more connected to nodes in disparate parts of their communities and overall more connected as individuals.

Vital communities were also distinguished by by the temporal emergence of centralizing individuals who connect different parts of the network (Table 1 and Fig. 3A, top right). In the analyses above of the whole networks, we found that epistemic vitality was associated with the presence of highly centralizing figures, as measured by peak centrality. Initially, however, both vital and static communities exhibited similar peak centrality (centrality at origin: $b = -0.06 \pm 0.04$ SEM, $t = 1.64, p = 0.11$), but at their culmination the productive communities had higher peak centrality ($b = -0.13 \pm 0.04$ SEM, $t = 3.31, p = 0.002$). This emergent difference reflected a significant increase over time in the peak centrality of vital communities ($b = 0.08 \pm 0.03$ SEM, $t = 2.7, p = 0.008$) but not in static communities ($b = 0.01 \pm 0.05$ SEM, $t = 0.25, p = 0.81$), though this difference was not itself significant ($b = 0.07 \pm 0.06$ SEM, $t = 1.16, p = 0.254$). Thus, while vital and static communities were equally connected by centralizing figures at their origin, only in vital communities do highly centralizing figures emerge who bridge divisions and consolidate influence, occupying key positions within the network.

Productive discourses may emerge from the merging of many disparate communities into a few highly-concentrated ones. Our analyses of whole networks found no association between epistemic vitality and the tendency for individuals to organize into sub-communities. The epistemic benefits of tight-knit sub-communities, however, may be specific to sub-communities that co-exist within the same time period. We thus analyzed the change over of time of sub-communities calculated cross-sectionally (i.e., among individuals within the same time period; Table 1 and Fig. 3A, top left). When they first originated, epistemically vital communities had the same amount of sub-community structure as other communities (effect of epistemic vitality on number of sub-communities, scaled by community size of temporal period at the communities' origination: $b = -0.06 \pm 0.07$ SEM, $t = -0.87, p = 0.39$). Over time, however, the community structure of non-vital communities remained stable (change over time in non-productive communities: $b = -0.07 \pm 0.08$ SEM, $t = -0.87, p = 0.39$), while epistemically vital communities gradually coalesced into a smaller number of communities (productive communities: $b = -0.24 \pm 0.04$ SEM, $t = -5.63, p < 0.001$). At their culmination, therefore, epistemically vital communities had significantly fewer communities than other communities (effect of epistemic vitality on scaled sub-communities at the communities' culmination: $b = -0.23 \pm 0.06$ SEM, $t = -4.11, p < 0.001$). These results suggest that while static and vital communities begin with comparable levels of fragmentation, only vital communities undergo a process of structural consolidation over time. This dynamic, not reflected in the results from the static whole network measures, arises as communities at each successive point in time became more integrated with one another compared to previous periods, showing how epistemically vital communities became increasingly integrated across generations.

The productive and non-productive communities did not differ along all dimensions. Some theories of epistemic vitality, for instance, argue that epistemically vital communities are characterized by high levels of disagreement (20). To test such theories, we looked at the relative amount of disagreement and agreement within each network, measured as the ratio of conflictual edges to non-conflictual edges (Table 1 and Fig. 3A, bottom right). We found no evidence that epistemic vitality is associated with individual-level antagonism on its own. At their temporal origination, vital and static communities did not differ in the amount of disagreement ($b = 0.18 \pm 0.3$ SEM, $t = 0.62, p = 0.54$). Both types of communities typically decreased in disagreement over time, but the amount of decrease did not differ significantly ($b = -0.25 \pm 0.37$ SEM, $t = -0.67, p = 0.50$). As a result, at their culmination, vital communities and static communities did not differ in the amount of disagreement ($b = -0.07 \pm 0.16$ SEM, $t = -0.42, p = 0.67$). The amount of disagreement on its own, therefore, is insufficient for epistemic vitality. Vital communities are not merely cantankerous; they organize disagreement into tight-knit communities that are connected by centralizing figures.

To synthesize these changes in network structure over time, we projected these network measures onto the first principal component derived from the Principal Component Analysis of whole-network measures described in the previous section. We did this both for network measures calculated cross-sectionally (i.e., for each non-overlapping time period) and cumulatively (i.e., for the entire network, from its origin to the current time period). This yielded a two-dimensional embedding of the network's time-evolving structure (Fig. 3B). At their origin, vital (orange) and static (green) communities were indistinguishable on the basis of their cumulative (y-axis) or cross-sectional (x-axis). Over time, static communities changed only minimally in their cumulative and cross-sectional structure, with their core structural attributes remaining largely stable (green trajectory in Fig. 3B). In contrast, epistemically vital communities showed consistent structural change in both their cumulative and cross-sectional structure. As a result, at their culmination, vital communities differed significantly from static ones in both their cumulative structure (measured at the final time point: $M_{static} = -3.82$ vs. $M_{vital} = -1.36$, $t_{35} = -4.78, p < 0.001$) and their cross-sectional structure ($M_{static} = -4.37$ vs. $M_{vital} = -2.08$, $t_{34} = -3.77, p < 0.001$). This suggests that philosophical vitality is associated with increasing structural integration and centralization over time, while static systems maintain a more fragmented and diffuse configuration.

## Discussion

What kinds of community structures facilitate epistemic vitality? By quantifying the social network structure of nearly 3,000 years of philosophical debate, we found that epistemic vitality is more likely to emerge in communities characterized by greater interconnectedness, increasing integration, and the presence of centralizing figures. Philosophical communities were more likely to sustain epistemic vitality if they consisted of integrated groups connected by centralizing figures. Diffuse communities that lacked clear intellectual lineages or centralizing figures remained epistemically static. While epistemically vital and static communities began with similar social structures, the features associated with vitality, such as the presence of strong centralizing figures and connectivity, tended to emerge over time, pointing to the importance of bridging nodes and integrative figures. At their origin, both vital and static communities consisted of fragmented subgroups. But while static communities maintained this fragmentation, vital communities became integrated by centralizing individuals, suggesting that a key historical driver of epistemic vitality is the bridging of previously disconnected perspectives into a cohesive and "energetic" intellectual discourse (Fig. 2A) (11).

Studies of collective intelligence have proposed at least three accounts of epistemic vitality: (1) that vital communities are sustained by persistent, individual-level disagreement (12, 13); (2) that they emerge from loosely connected sub-communities pursuing parallel lines of inquiry (21, 22); and (3) that they depend on central actors who bridge these communities and integrate their insights (11, 23). Our findings speak to all three accounts. Our analyses of historical philosophical communities suggested that — while disagreement may play a role — the structure of this disagreement, rather than its mere presence, is more critical. We also found that while most philosophical communities begin with multiple sub-communities, vital communities show evidence of consolidation of this community structure over time time. Furthermore, this consolidation appears to be driven, in part, by the emergence of centralizing individuals who play a key role in linking and synthesizing across sub-communities. Notably, the community structures associated with epistemically vital communities differ from theories that emphasize sparsity and persistent discord (13). Instead, our findings highlight the importance of restructuring over time, where diverse interactions became increasingly integrated — a dynamic that may guide future research in collective intelligence.

Since communities can differ in their epistemic goals, different communities may benefit from different network structures. Philosophical traditions that emphasize "progress" or practical application, for instance, may be more likely to produce communities marked by synthesis and centralization. Others may prioritize the preservation of inherited systems of thought. For instance, if a community is attempting to preserve, unchanged, an older tradition of thought, then vitality as defined here may be antithetical to this goal.

Our approach was informed by research on the role that social structures play in shaping the vitality of other knowledge systems and was thus idea-agnostic, focusing not on the content of philosophical thought but on the structural dynamics that shape how philosophical systems evolve and interact over time (5, 25–28). We have not examined the specific content of different philosophical systems or the specific intellectual role that centralizing thinkers may have played in these systems, whether as radical innovators, synthesizers, or disruptors. This allowed us to first isolate the potential role of network topology itself, a factor that has been implicated in the dynamics of other epistemic systems (7, 16, 23). Future work may help bridge this gap by incorporating the sociopolitical contexts and ideological contents of these traditions, as emphasized in the original sociological analysis of Randall Collins (11) and work in comparative philosophy (1). For instance, to clarify how epistemic commitments manifest in both discourse and community structure, one could integrate community-level analyses, such as those presented here, with natural language processing of the products of those communities (e.g., publications, letters, etc.) (29, 30).

One fertile direction for future research lies in a comparative project examining structural differences between philosophy, science, and other epistemic enterprises. Studies of epistemic communities of the sort pursued here and within the science of science could be extended to a wide range of discursive knowledge systems, including those found in mathematics, theology, and other so-called "thought collectives" (3). Such work could help clarify whether science and philosophy represent fundamentally distinct enterprises and how each contributes to the broader social systems through which humans understand and engage with the world. Such a project holds promise to integrate the "science of science" (8) with the "philosophy of science" (4, 31, 32) to inform our scientific understanding of the dynamics of knowledge writ large — a "science of philosophy" (26–28).

## Methods

**Digitizing the social networks of philosophy.** Networks were digitized from Randall Collins's comparative history of global philosophy, *The Sociology of Philosophies* (11). On the basis of historical documents, Collins reconstructed networks of philosophical agreement, disagreement, and master-pupil relationships for philosophical communities across history ($N = 3187$ nodes representing philosophers and $N = 5415$ edges across $N = 55$ networks). Communities in the dataset span over 2700 years, from 800 BCE to 1980 CE; cover a wide geographic range, stretching across Europe, North Africa, the Middle East, India, China, and Japan, and more; and encompassing philosophical movements from ancient Buddhism to modern French existentialism. In total, the dataset includes 55 networks:

- 24 from Asia, including 10 with philosophers in ancient China, 7 in India, 8 in Japan, and 4 encompassing Persian, Middle Eastern, and Central Asian thinkers,
- 30 from Europe (9 from Ancient Greece, 5 from Rome, and the remainder spread across Western and Central Europe) and 5 from North Africa and the Middle East,
- 1 from North America (the American Pragmatist movement).

The dataset thus reflects a diversity of eras, regions, and intellectual traditions, across multiple languages and cultural contexts, from Classical Chinese, Sanskrit, and Arabic, to Latin, German, French, and English.

We digitized the network diagrams in Collins' text (11). For each network, two independent coders translated Collins' visualization into a directed, weighted edge list. In each edge list, philosophical disputes were assigned a directed negative weight ($-1$), master-pupil relationships received a directed positive weight ($+1$), and acquaintances were coded as bi-directional positive edges ($+1$ in both directions). In cases of coding discrepancies, a third coder adjudicated. We also divided each network into non-overlapping temporal periods, following the temporal divisions introduced in Randall Collins's visualizations (11). Within each network, these periods were of equal duration, typically corresponding to approximately one generation, although sometimes of longer duration for communities that persisted for hundreds of years. To validate the historical accuracy of these edge lists, we used a pretrained large language model (OpenAI's GPT-4o) to independently assess relationships between philosophers. The model agreed substantially with the edge lists (Cohen's $\kappa = 0.69$, $p < .0001$; see SM Methods and Fig. S6 for details).

Collins also evaluated the epistemic vitality of each community using a comparative methodology that spanned a large temporal and geographical range. This allows for a bird's eye evaluation of different philosophical community's relative productivity and vitality. A coder reviewed Collins' discussion of each community and made a made a holistic determination of the community's epistemic vitality. To validate these judgments, we used a pretrained machine learning model, OpenAI's GPT-4o, to independently assess the epistemic vitality of each network using Monte Carlo cross-validation (repeated random subsampling). We prompted the model with a definition of epistemic vitality, along with a a labeled random subset of communities (75% each of vital and static communities). Communities were described only by name (e.g., "Network of Greek Philosophers from Socrates to Chrysippus") and date range. We then asked the model to decide whether or not each of the remaining communities was epistemically vital. This process was repeated 101 times, and for each community we calculated how often it was classified as epistemically vital by the model. This ML-derived measure of epistemic vitality was significantly correlated with the human classification of epistemic vitality classification derived from Collins (2000) (11) ($r = .30, p = .02$). In the Supplemental Materials, we reproduce our main results using this machine classification of epistemic vitality.

**Quantifying network topology.** For each network we calculated 13 measures of structural and relational properties: average degree, average clustering, peak centrality, connectivity, flow hierarchy, average path length, network diameter, sparsity, modularity, communities, cores, coreness, and the ratio of disagreement to agreement. (See Table S1 for a description of all network measures.) We calculated these measures for each entire network. We also calculated the temporal evolution of these measures within each community in two ways: (1) cross-sectionally, for each non-overlapping time period, and (2) cumulatively, incorporating all network interactions from the beginning up to the given time period.

To capture variation along theoretically-important dimensions of interconnectedness (23), centralizing individuals (18), and productive fragmentation (13), we performed Principal Component Analysis (PCA) on four key variables (connectivity, scaled communities, peak centrality, and average degree) calculated on the whole networks. The first principal component (PC1) explained $74\%$ of the variance. We also used this principle component to capture the temporal evolution of network structure by projecting the cross-sectional measures and the cumulative measures on to PC1.

**Time-respecting null models.** For each of these growing networks, we constructed time-respecting null models ($N = 1000$) (33) in which edges are randomly shuffled while respecting the network's temporal structure. To do so, edges are shuffled using a Newman edge-rewiring algorithm (34) in two ways. Edges between nodes within the same time period are shuffled to connect random nodes within the same time period. Edges between a given time period and a preceding time period (i.e., between philosophers and their predecessors) are shuffled so they connect the same two periods. This approach preserves the time-respecting nature of the networks and the balance of edges both within and between time periods, while randomizing the overall network structure and degree distribution.

**Statistical analyses.** To assess the structural distinctiveness of the real networks, we calculated all measures for each time-respecting null model and compared the real networks to these null distributions (i.e., z-scored using the null distributions).

To analyze the networks' temporal evolution, we used linear mixed effect models to predict each measure of network structure. Models included fixed effects for time (normalized within each network to range from 0 to 1), epistemic vitality (static = 0, vital = 1), and their interaction, and random intercepts and slopes for time for each network. To estimate the effect of epistemic vitality at the networks' culmination (i.e., when time = 1), we rebaselined time so it runs from -1 (origin) to 0 (culmination). To estimate the change over time for epistemically vital networks, we rebaselined epistemic vitality (i.e., vital = 0, static = 1).

**ACKNOWLEDGEMENTS**

This work received support from the Institute for Humane Studies under grant no. IHS017802 awarded to CM.

## References

1. Hajime Nakamura. *A Comparative History of Ideas*. Kegan Paul International, London, 1992. ISBN 071030384X.
2. Edwin Hutchins. Cognitive ecology. *Topics in cognitive science*, 2(4):705–715, 2010.
3. Jonathan Harwood. Ludwik fleck and the sociology of knowledge. *Social studies of science*, 16(1):173–187, 1986.
4. Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, 1962.
5. David L. Hull. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. University of Chicago Press, Chicago, 1988. ISBN 9780226360507.
6. Paul Thagard. Societies of minds: Science as distributed computing. *Studies in History and Philosophy of Science*, 24(49-67):251–257, 1993.
7. Mark EJ Newman. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2):404–409, 2001.
8. Santo Fortunato, Carl T Bergstrom, Katy Börner, James A Evans, Dirk Helbing, Staša Milojević, Alexander M Petersen, Filippo Radicchi, Roberta Sinatra, Brian Uzzi, et al. Science of science. *Science*, 359(6379):eaao0185, 2018.
9. John Henry Newman. *An essay on the development of Christian doctrine*. James Toovey, 1846.
10. Richard Rorty. *Philosophy and the Mirror of Nature*. Princeton University Press, Princeton, NJ, 1979.
11. Randall Collins. *The sociology of philosophies: A global theory of intellectual change*. Harvard University Press, 2000.
12. Kevin JS Zollman. The epistemic benefit of transient diversity. *Erkenntnis*, 72(1):17–35, 2010.
13. Paul E Smaldino, Cody Moser, Alejandro Pérez Velilla, and Mikkel Werling. Maintaining transient diversity is a general principle for improving collective problem solving. *Perspectives on Psychological Science*, 19(2):454–464, 2024.
14. Jacob G Foster, Andrey Rzhetsky, and James A Evans. Tradition and innovation in scientists' research strategies. *American sociological review*, 80(5):875–908, 2015.
15. Lingfei Wu, Dashun Wang, and James A Evans. Large teams develop and small teams disrupt science and technology. *Nature*, 566(7744):378–382, 2019.
16. Jesse Milzman and Cody Moser. Decentralized core-periphery structure in social networks accelerates cultural innovation in agent-based modeling. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 2213–2221, 2023.
17. Peter Csermely, András London, Ling-Yun Wu, and Brian Uzzi. Structure and dynamics of core/periphery networks. *Journal of Complex Networks*, 1(2):93–123, 2013.
18. Janet C Long, Frances C Cunningham, and Jeffrey Braithwaite. Bridges, brokers and boundary spanners in collaborative networks: a systematic review. *BMC health services research*, 13:1–13, 2013.
19. Jingyi Wu and Cailin O'Connor. How should we promote transient diversity in science? *Synthese*, 201(2):37, 2023.
20. Feng Shi, Misha Teplitskiy, Eamon Duede, and James A Evans. The wisdom of polarized crowds. *Nature human behaviour*, 3(4):329–336, 2019.
21. Yiling Lin, James A Evans, and Lingfei Wu. New directions in science emerge from disconnection and discord. *Journal of Informetrics*, 16(1):101234, 2022.
22. Christina Fang, Jeho Lee, and Melissa A Schilling. Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization science*, 21(3):625–642, 2010.
23. Cody Moser and Paul E Smaldino. Innovation-facilitating networks create inequality. *Proceedings of the Royal Society B*, 290(2011):20232281, 2023.
24. James A Evans and Donghyun Kang. Lack of diffusion of popular scientific ideas marks the presence of epistemic'bubbles'. *Nature Human Behaviour*, 2025.
25. Damon Centola. The network science of collective intelligence. *Trends in Cognitive Sciences*, 26(11):923–941, 2022.
26. Pablo Contreras Kallens, Daniel J Hicks, and Carolyn Dicey Jennings. Networks in philosophy: Social networks and employment in academic philosophy. *Metaphilosophy*, 53(5):653–684, 2022.
27. Pablo Contreras Kallens and Jeff Yoshimi. Bibliometric analysis of the phenomenology literature. In *Horizons of Phenomenology: Essays on the State of the Field and Its Applications*, pages 17–47. Springer International Publishing Cham, 2023.
28. Martin Davies and Angelito Calma. Australasian journal of philosophy 1947–2016: a retrospective using citation and social network analyses. *Global Intellectual History*, 4(2):181–203, 2019.

29. Dan Edelstein, Paula Findlen, Giovanna Ceserani, Caroline Winterer, and Nicole Coleman. Historical research in a digital age: reflections from the mapping the republic of letters project. *The American Historical Review*, 122(2):400–424, 2017.
30. Ruth Ahnert and Sebastian E Ahnert. *Tudor networks of power*. Oxford University Press, 2023.
31. Patrick Grim, Frank Seidl, Calum McNamara, Isabell N Astor, and Caroline Diaso. The punctuated equilibrium of scientific change: a bayesian network model. *Synthese*, 200(4): 297, 2022.
32. Karin Knorr Cetina. Culture in global knowledge societies: Knowledge cultures and epistemic cultures. *Interdisciplinary science reviews*, 32(4):361–375, 2007.
33. Zhuo-Ming Ren, Manuel Sebastian Mariani, Yi-Cheng Zhang, and Matúš Medo. Randomizing growing networks with a time-respecting null model. *Physical Review E*, 97(5):052311, 2018.
34. Mark EJ Newman, Steven H Strogatz, and Duncan J Watts. Random graphs with arbitrary degree distributions and their applications. *Physical review E*, 64(2):026118, 2001.

## Supplementary Note 1: Pairwise correlations of all network metrics

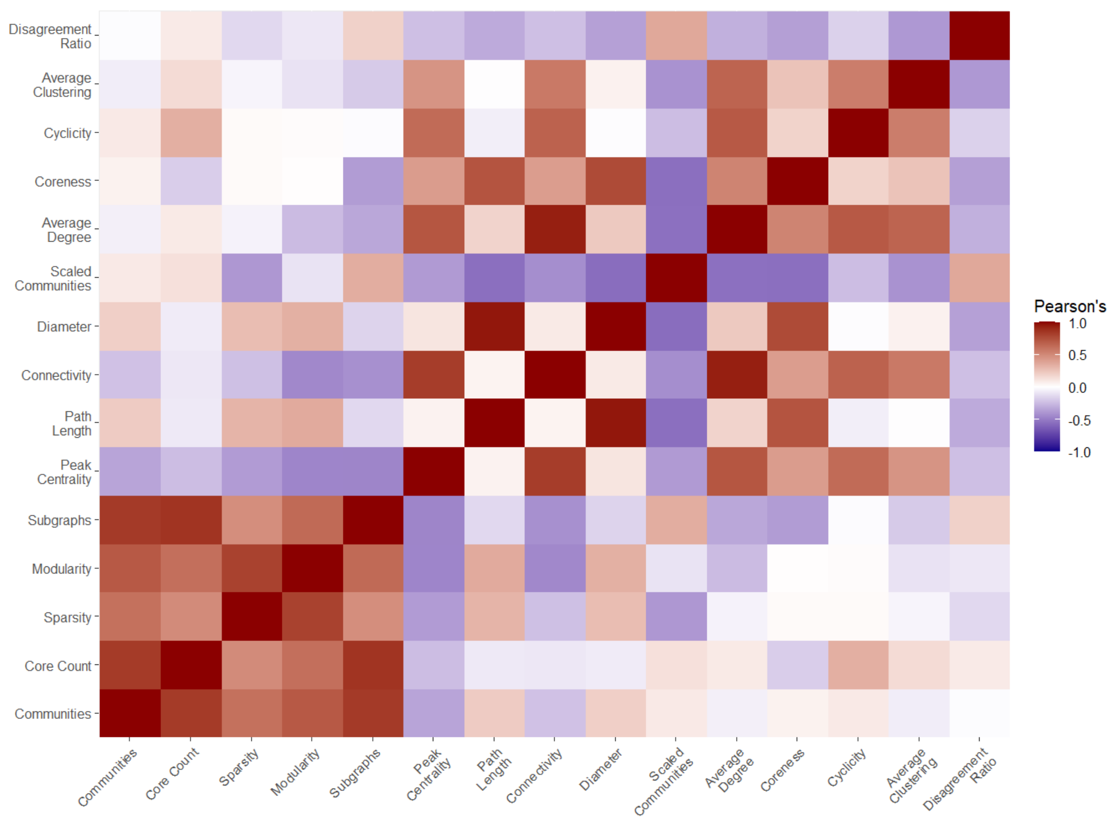


**Fig. S1.** Correlation heatmap of all network metrics used in the study, showing relationships among structural and community-level features. Red indicates positive correlations; blue indicates negative correlations (Pearson's $r$).

## Supplementary Note 2: Observed and null distributions of network-based structural, community, and path length metrics

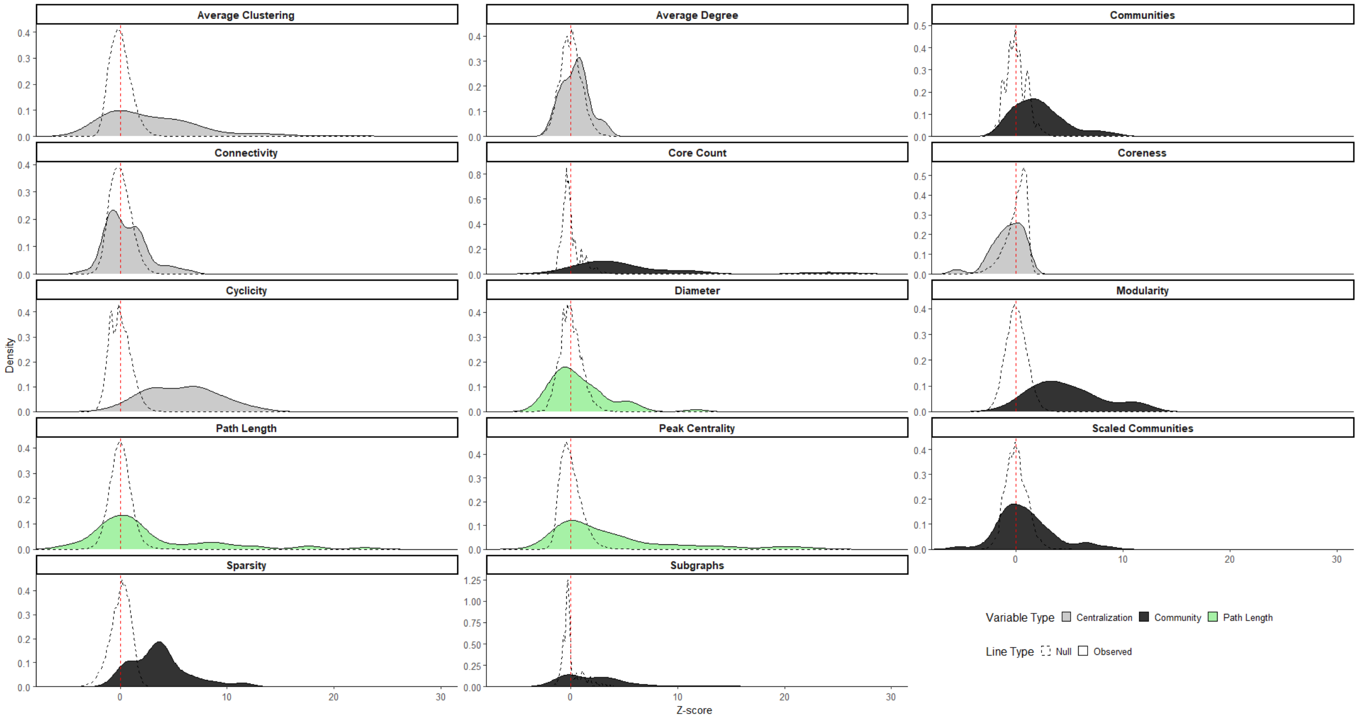


**Fig. S2.** Density plots of z-scored network metrics across graphs. Metrics are categorized by variable type from correlation metrics between all variables (Fig S1): Community (black), Centralization (gray), and Path Length (green), while dashed curves indicate null model distributions. Dashed vertical lines mark the mean of null distributions.

## Supplementary Note 3: Timeline of vital and static communities in the history of philosoph (600BCE-1940CE

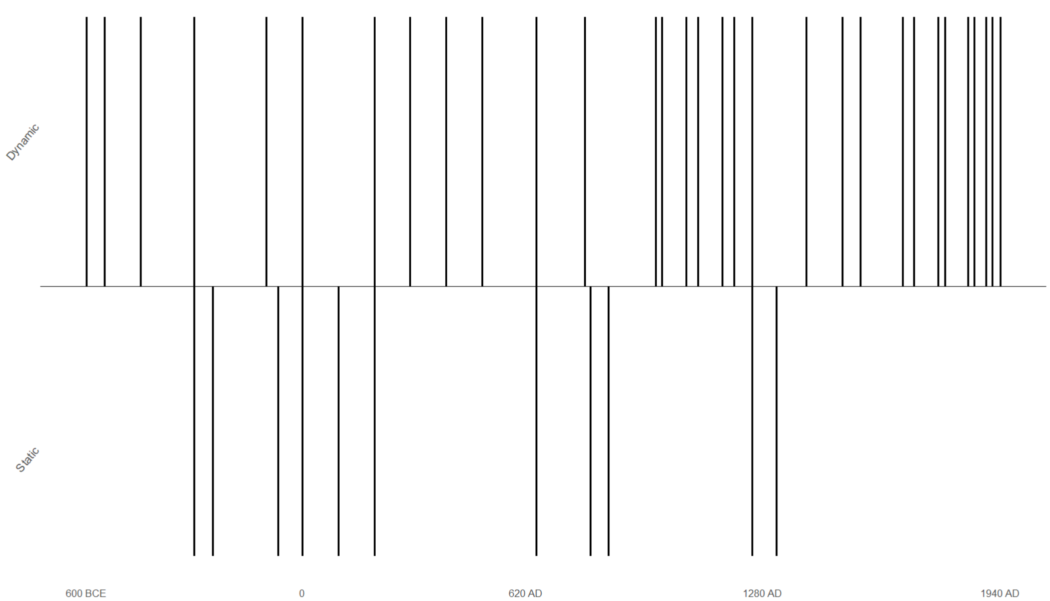


**Fig. S3.** Timeline of static and dynamic philosophical spanning from 600 BCE to 1940. Vertical lines represent the middle point of each tradition classified as either static or dynamic.

## Supplementary Note 4: Validating Epistemic Vitality

In the Main Text, we use a measure of epistemic vitality based on human annotation of Randall Collins' global history of philosophy (11) (hereafter, "human-judged-vitality). To validate this measure of epistemic vitality,, we used a large language model (OpenAI's GPT-4o) to independently classify the epistemic vitality of every philosophical community (hereafter, "ML-vitality"). We employed Monte Carlo cross-validation (repeated random subsampling). In each iteration, the model was given a definition of epistemic vitality along with a labeled random subset of communities ($75\%$ each of vital and static communities). We then asked the model to classify the remaining communities as either epistemically vital or static. This process was repeated 101 times, and for each community, we calculated the proportion of runs in which it was classified as vital. These proportions were thus a continuous measure of each community's vitality assessment ("ML-vitality").

We first asked whether ML-vitality was consistent with the human-judged-vitality measures used in the Main Text. A linear regression predicting ML-vitality from human-judged-vitality showed a significant positive relationship ($b = 0.271 \pm 0.120$ SE, $p = 0.027$). A Welch two-sample t-test confirmed this association, with significantly higher ML-vitality values for human-judged-vital communities ($M_{\text{static}} = 0.39$ vs. $M_{\text{vital}} = 0.66$, $t_{53} = -2.2$, $p = 0.043$). We thus examined whether our main findings replicate when we use ML-vitality instead of human-judged-vitality.

At the network level, we found that ML-vitality predicted PC1 of the Principle Component Analysis of network structure, which captures integration and centralization (linear regression predicting PC1: $b = 0.894 \pm 0.270$ SE, $p = 0.002$; Fig. 2), thus replicating the finding in the Main Text. ML-vitality did not predict PC2, which captures fragmentation and antagonism (PC2: $b = 0.051 \pm 0.139$ SE, $p = 0.368$), once again replicating the result in the Main Text. Likewise, we find more directly that ML-vitality does not predict the levels of antagonism in communities (linear regression predicting antagonism: $b = -0.065 \pm 0.057$ SE, $p = 0.262$). Thus, epistemically vital and static communities differ significantly in their structural integration and centralization, but not in fragmentation or antagonism, whether communities' vitality was judged by human or ML model.

We next examined the temporal evolution of these networks. To convert the continuous ML-vitality measure into discrete categories, we used tertile-defined categories. The temporal evolution of these ML-defined vital and static communities was qualitatively similar to temporal evolution of human-judged categories (Fig. S4). At their culmination, ML-defined vital communities differed significantly from static ones in their cumulative structure ($M_{\text{static}} = -2.71$ vs. $M_{\text{vital}} = -0.69$, $t_{35} = -3.07$, $p = 0.004$). However, departing from the Main Text, they did not differ significantly in their cross-sectional structure at the very end ($M_{\text{static}} = -3.02$ vs. $M_{\text{vital}} = -1.60$, $t_{34} = -1.85$, $p = 0.071$).

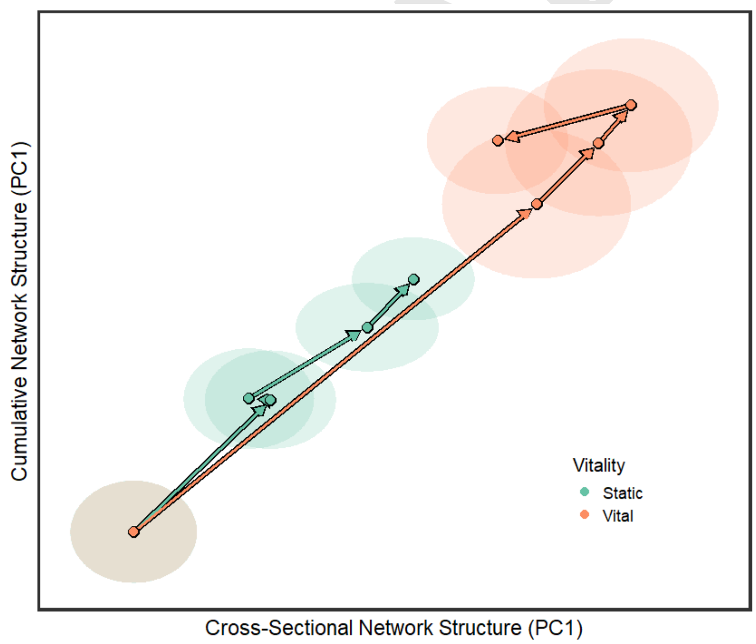


**Fig. S4.** Temporal differentiation of the network structure of vital (orange) and static (green) communities, using the first and third tertiles of GPT-4o's continuous vitality ratings. The x-axis shows the network structure (integration and centralization) within each time period. The y-axis shows the network structure accumulated from a community's origin up until that time period. This allows us to visualize the temporal evolution of network structure, both cross-sectional (x-axis) and cumulative (y-axis). Both static and vital communities increase in cumulative structure over time (y-axis), reflecting growing integration and centralization. However, vital communities show more consistent and pronounced changes across both cross-sectional structure (x-axis) and cumulative integration (y-axis), whereas static communities exhibit more modest or variable shifts in network structure across time periods. (Points = means. Shaded circles = standard errors.

## Supplementary Note 5: Temporal evolution of network metrics by epistemic vitality

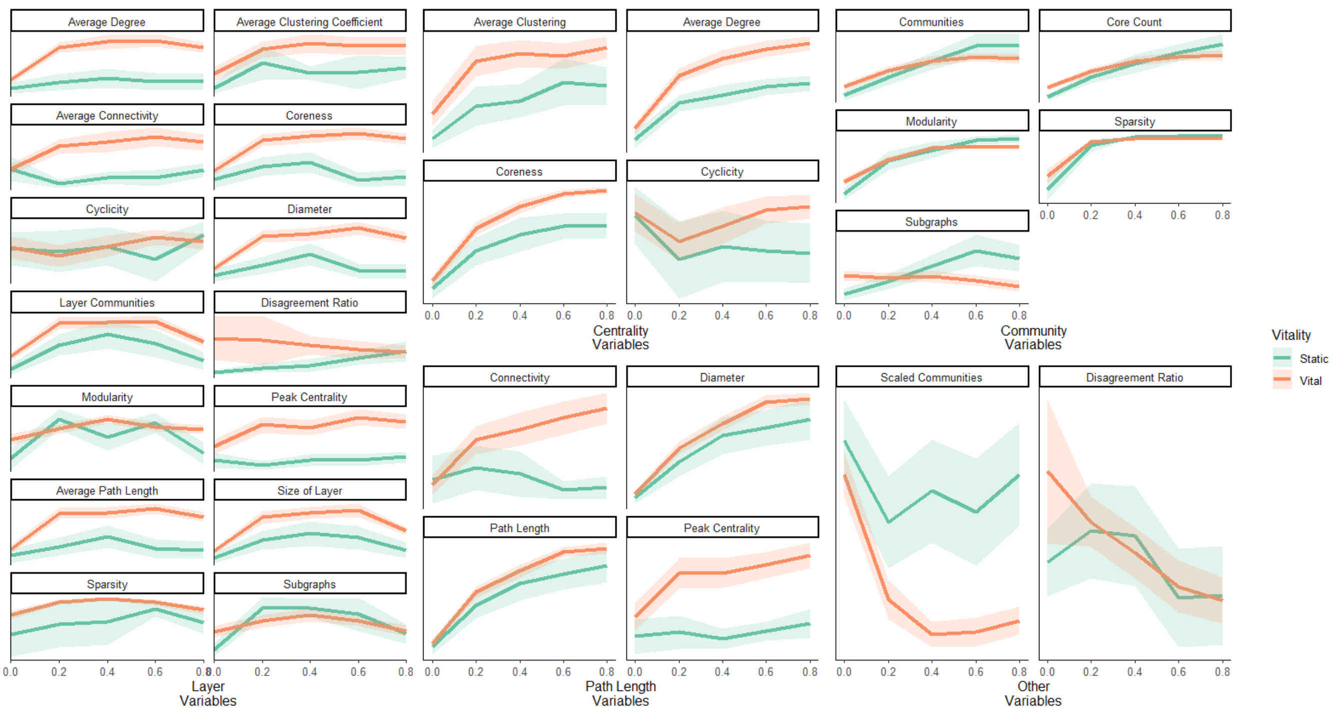


**Fig. S5.** Line plots of all network metrics used in the study. Shaded areas represent standard errors of the mean. Lines are grouped by network vitality condition (Static vs. Vital).

## Supplementary Note 6: Historical edge validation using LLM prediction

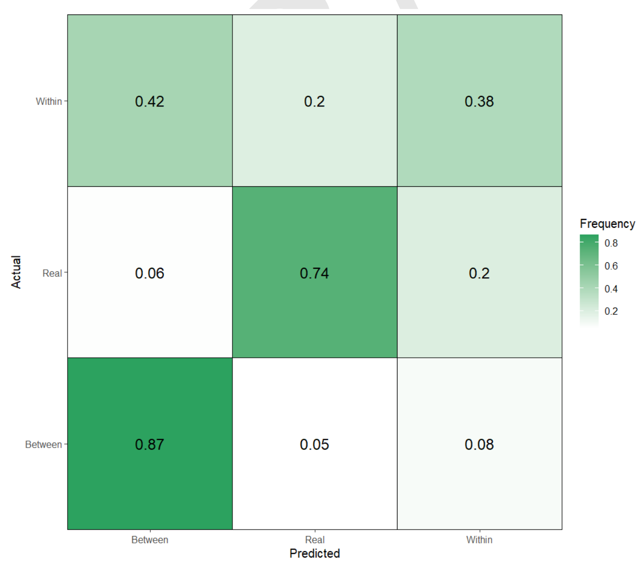


**Fig. S6.** To validate the edge list, we combined the list of all edges between major philosophers ($N = 884$), or those with names in the dataset, with a surrogate list of philosopher duos who were not connected. Approximately half of these surrogate duos consisted of philosophers from the same network ($N = 220$), with the rest from different networks ($N = 275$). We then queried a pretrained large language model (OpenAI's GPT-4o) about every pair, both real and surrogate, asking whether the two philosophers had interacted or influenced each other, belonged to the same community but had not interacted or influenced each other, or belonged to entirely different communities. There was substantial agreement between the model and the edge list we derived from Collins's text (11); (weighted Cohen's $\kappa = 0.69, p < .001$).

## Supplementary Note 7: Features Table

| Feature | Label (Whole/Temporal) | Description | Variable Type | Feature Cluster |
|---|---|---|---|---|
| Average Clustering Coefficient | average_clustering, layer_clustering | Measure of proportion of how many of a node's neighbors are connected to one another to form a complete clique. | Node average | Centralization |
| **Average Degree** | average_deg, average_deg_layer | Measure of the number of neighbors an individual node has. | Node average | Centralization |
| Coreness | cp_ratio, cp_layer | Graph-based measure of the number of nodes with shortest path lengths equivalent to the graph's diameter to those which are not. | Graph-based | Centralization |
| Cyclicity | cyclicity, cyclicity_layer | The inverse of a graph's flow hierarchy (or the fraction of edges not participating in a cycle), where a cycle is a path that starts and ends at the same node, in a graph. | Graph-based | Centralization |
| **Sub-Communities** | communities, layer_communities, comm_scale | Number of distinct communities detected in the graph using the modularity-based Louvain optimization procedure. | Graph-based | Community |
| Cores | cores | Graph-based measure of the number of cores in a k-core decomposition of the graph, wherein a k-core is a group of nodes that each possess at least *k* connections. | Graph-based | Community |
| Modularity | modularity, mod_layer | Measures the strength of division of a network into distinct communities based on the density of edges within communities compared to between them. Communities were defined using the Louvain optimization algorithm. | Graph-based | Community |
| Size of Layer | size_layer | The number of nodes in a temporal layer of a graph. | Layer-based | Community |
| Sparsity | sparsity, sparsity_layer | The number of edges in a network divided by the number of possible edges in the network. | Graph-based | Community |
| Subgraphs | subgraphs, subgraphs_layer | Number of connected components in the graph, where each subgraph is a set of nodes that are connected. | Graph-based | Community |
| **Average Connectivity** | connectivity, connectivity_layer | Measure of the minimum number of nodes between two pairs which must be removed to disconnect the pair. | Node average | Path-based |
| Average Path Length | path_length, pl_layer | Measure of the shortest number of edges between a given pair of nodes. | Node average | Path-based |
| Diameter | diameter, diameter_layer | The maximum eccentricity (the longest shortest path length) in a graph. | Graph-based | Path-based |
| **Peak Centrality** | peak_centrality, peak_centrality_layer | The highest betweenness centrality of a node in a graph, where betweenness centrality measures how often a node lies on the shortest paths between other nodes. | Graph-based | Path-based |
| Disagreement Ratio | disagreement_ratio, layer_disagreement | The ratio of antagonistic to non-antagonistic edges in a graph. | Graph-based | Edge-based |

**Table S1.** Feature descriptions and types, sorted by Feature and cluster from correlation matrices. Bolded variables are those used in the Principal Components Analysis.