TITLE PAGE

Title: A mechanism of abiogenesis based on complex reaction networks organized by seed-dependent autocatalytic systems

Author Line

Zhen Peng (彭镇)[1], Jeff Linderoth[1,2], David A. Baum[1,3*]

Author Affiliations

[1]Wisconsin Institute for Discovery, University of Wisconsin-Madison, Madison WI 53706, USA

[2]Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison WI 53706, USA

[3]Department of Botany, University of Wisconsin-Madison, Madison WI 53706, USA

Corresponding Author

David A. Baum

Address: Department of Botany, Birge Hall 216/220, 430 Lincoln Drive, Madison, WI 53706, USA

Phone number: (+1) 608-265-5385

Email address: dbaum@wisc.edu

Keywords

complex systems, origin of life, reaction networks, autocatalysis, chemical ecosystem, evolution

ABSTRACT

The core of the origin-of-life problem is to explain how a complex dissipative system could emerge spontaneously from a simple environment, perpetuate itself, and complexify over time. This would only be possible, we argue, if prebiotic chemical reaction networks had autocatalytic features organized in a way that permitted the accretion of complexity even in the absence of genetic control. To evaluate this claim, we developed tools to analyze the autocatalytic organization of food-driven reaction networks and applied these tools to both abiotic and biotic networks. Both networks contained seed-dependent autocatalytic systems (SDASs), which are subnetworks that can use a flux of food chemicals to self-propagate if, and only if, they are first seeded by some non-food chemicals. Moreover, SDASs were organized such that the activation of a lower-tier SDAS could render new higher-tier SDASs accessible. The organization of SDASs is, thus, similar to trophic levels (producer, primary consumer, etc.) in a biological ecosystem. Furthermore, similar to ecological succession, we found that higher-tier SDASs may produce chemicals that enhance the ability of the entire chemical ecosystem to utilize food more efficiently. The SDAS concept explains how driven abiotic environments, namely ones receiving an ongoing flux of food chemicals, can incrementally complexify even without genetic polymers. This framework predicts that it ought to be possible to detect the spontaneous emergence of life-like features, such as self-propagation and adaptability, in driven chemical systems in the laboratory. Additionally, SDAS theory may be useful for exploring general properties of other complex systems.

MAIN TEXT

1. INTRODUCTION

Life is the paradigmatic example of a complex system [1,2]. It consists of diverse chemical components, some of which are large, energy-expensive molecules. Furthermore, these components are not simply lumped together but are organized in such a way that the entire system can resist environmental perturbations, grow or divide to give rise to more life, and evolve over time. Such organization is not cheap, of course, which is why life must consume energy to maintain its internal order. It is almost magical that complex organisms can use "untargeted" energy sources such as light to convert simple materials (e.g., water, carbon dioxide, and minerals) into more living matter of the same kind. The mystery of the origin of life is how a system with sufficient complexity to conduct such improbable conversions could emerge spontaneously, when there was no prior design to follow or template to copy.

The solution to this conundrum lies, we believe, not in looking for particular molecules or reactions, but in looking at the emergent properties of chemical reaction networks. Autocatalysis is a process of which a product is also a catalyst for that process. If real chemical reaction networks contain many autocatalytic motifs, systems with a pre-existing ability to sustain themselves may arise readily, and then become progressively more complicated through the accretion of more autocatalytic modules. In our recent work [3], we used analyses of "toy" reaction networks to show that interacting autocatalytic cycles can exhibit features that closely resemble those seen when biological species interact ecologically. This result implies that driven chemical reaction systems, or "chemical ecosystems," can show complex dynamics, including succession. Here, we sought to extend this analysis to real chemical reaction networks to see if they have the features needed for stepwise complexification. Specifically, we focused on reaction networks that "bracket" the origin of life: a radiolytic and geochemical reaction database [4] serving to represent chemistry without life, and a curated subset of biochemical reactions [5] serving to constrain the metabolic network of the Last Universal Common Ancestor (LUCA). Our reasoning is that features shared by both networks are very likely to apply at the time when systems deserving the label "life" first emerged.

In this study, we focus on the concept of a seed-dependent autocatalytic system (SDAS), which is a network motif that can be activated by a rare chemical event but can, once activated, sustain

itself. When a network with multiple potential SDASs is driven by a suitable flux of food, it seems possible that it could gradually acquire new components, some of which may affect the overall efficiency of food utilization and/or open-up opportunities for yet further complexification. We suggest that, when actualized in spatially and temporally structured environments, a multi-SDAS network may be able to show adaptive evolution without producing genetic polymers.

We start by formalizing network expansion, the basic procedure needed to map out a network's architecture, and then describe an algorithm for identifying SDASs within a reaction network represented by a stoichiometric matrix. Using the algorithm, we show that SDASs are found in both the abiotic and biotic chemical reaction networks and are organized in tiers, which allows for the stepwise accretion of complexity. Furthermore, the higher-tier SDAS that we detected in the larger (biotic) network includes particular chemicals that would be expected to increase the efficiency of the system as a whole and open-up new ecological opportunities. This analysis supports the contention that reaction networks can evolve adaptively, in the sense of transitioning periodically to new quasistable states that use resources more efficiently. We end by discussing implications of our results, several remaining challenges, how our theory can be used to guide laboratory experiments, and the significance of our theory to complexity science in general.

2. RESULTS

2.1 Databases of abiotic and biotic reactions

The abiotic reaction database we analyzed was extracted from seven decades of published data [4]. This database includes free radical reactions, mineral geochemical reactions, amino acid production, chloride radical and polar reactions, nitrile radical and polar reactions, RNA nucleotide assembly, nuclear decay, and physicochemical reactions. To this database, we added some additional well-known abiotic organic reactions, most notably the classical formose reaction [6], but without formaldehyde dimerization because it is very slow and its reaction mechanism is yet to be determined [7,8]. Sources of radiation such as X-rays, ultraviolet, and visible light are treated as reactants and products, which is why we will use "entity" to refer to both chemicals and electromagnetic energy sources in this specific database. In part because of

these energy sources, most of the reactions in the abiotic database are irreversible. Although only covering a small portion of known abiotic chemistry, these reactions may still be used to test the applicability of our theoretical framework to real chemical reaction networks.

The biotic reaction database we studied is a subset of the KEGG biochemical reaction database [9–11] that was obtained by removing reactions that only occur in eukaryotes and reactions dependent on O_2 [5]. Xavier et al. (2020) [5] claimed that the resulting reaction network could be a proxy for the primordial metabolism of LUCA. We added five spontaneous, reversible reactions that are missing from KEGG, for example $H_2O \leftrightarrow H^+ + OH^-$ and $H_2CO_3 \leftrightarrow H^+ + HCO_3^-$, and one reaction (R06974) that is in KEGG but was filtered out by Xavier et al. (2020) [5]. As discussed in Section 4.1.2, this reaction is not oxygen-dependent and entails a reaction mechanism that is very similar to the reaction KEGG R06975, which is already included in the database. All reactions in the reaction network are assumed to be reversible, so each of them was split into two unidirectional reactions (Section 4.1.2).

We acknowledge that, because most of the reactions in the KEGG database are catalyzed by enzymes, it is likely that many of them could not occur at sufficiently high rates in a prebiotic world to have been relevant before biological catalysts evolved. Nonetheless, since all these reactions are chemically feasible, we reasoned that the relationships between reactants and products in such a curated biochemical reaction network is meaningful and that features shared by it and the abiotic network should be relevant to the origin of life. Thus, our goal is not to claim that the chemical reactions in this database are those that applied during the origin of life but to explore how a network with this topology would be expected to change over time when driven by food flux.

We preprocessed the databases such that every reaction appears only once and had clear stoichiometry. After curation, the abiotic reaction database consists of 277 entities and 717 unidirectional reactions for a reaction density of 2.59 reactions per entity (Table S1). The resulting biotic reaction database consists of 4216 chemicals and 8402 unidirectional reactions for a reaction density of 1.99 reactions per chemical (Table S2). Fig. S1 shows the histograms of the numbers of reactions that an entity or chemical species is involved in for the two databases. Both distributions fit a power law, being highly right-skewed, meaning that most entities or chemicals are involved in a small number of reactions.

2.2 Network expansion and tier-0 systems

Reaction databases are just collections of reactions allowed to occur. To generate an organized subnetwork for analysis, we define a network expansion operation, $\Xi(\boldsymbol{S_O}, \boldsymbol{R})$, which calculates all reactions and chemical species that can be accessed given a starting set of chemical species $\boldsymbol{S_O}$ and a set of all allowed reactions $\boldsymbol{R}$. The network expansion (see Section 4.2) starts with $\boldsymbol{S_O}$ and scans $\boldsymbol{R}$ for any reactions that are not yet in the subnetwork, but whose reactants are all present in the current subnetwork. These reactions are added to the subnetwork and the expansion iterates until no more chemical reactions can be added to the subnetwork. The reaction subnetwork resulting from the expansion is described by the tuple $(\boldsymbol{S_E}, \boldsymbol{R_E})$, where $\boldsymbol{S_E}$ is the set of chemical species and $\boldsymbol{R_E}$ is the set of reactions. The network expansion operation can be visualized using a stoichiometric matrix, where each row represents a chemical species and each column represents a unidirectional reaction, with stoichiometric coefficients as the entries (Fig. S2(a)-(d)). In general, reactants have negative stoichiometric coefficients while products have positive ones. However, some reactions may involve species that are present in both reactants and products (e.g., catalysts), which could make the stoichiometric coefficients of such species non-negative even though such species are required for the reactions to occur. To avoid downstream problems when catalysts are present (as is not the case here), columns should be annotated to indicate if any chemicals besides those with negative signs need to be present for the reaction to occur.

Since we are interested in systems that are maintained out of equilibrium by a constant flux of food or energy from the environment, the first step is to specify the ultimate food set. Note that chemicals in the ultimate food set might also be produced by reactions within the subnetwork. A full expansion starting from the ultimate food set generates a reaction subnetwork including all chemicals that would be expected to be generated at a non-zero rate in an environment receiving flux of the ultimate food species. We will call this subnetwork the tier-0 system, reflecting the fact that no triggering events are needed except for the provision of the ultimate food (Fig. S2(a)-(d)).

Although there is still no wide consensus on the environmental conditions of the prebiotic sites where life originated [12–20], to illustrate the approach, we selected simple entities as ultimate food. For all analyses with the abiotic reaction database, we chose {H_2, CH_4, NO, FeS_2, visible light} as the ultimate food set. A full expansion starting from this set generated a tier-0 system

with no additional entities and zero reactions. For all studies of the biotic reaction database, we chose {H_2O, CO_2, NH_3, H_2S, H_2SO_4, H_2SO_3, HSO_3^-, H_3PO_4, $H_4P_2O_7$} as the ultimate food set. The full expansion starting from this set generated a tier-0 system with 30 chemical species and 44 unidirectional reactions (i.e., 22 reversible reactions). This network contains substructures that can be drawn as autocatalytic cycles [3,21], an example of which is shown in Fig. S3. However, because all the members in the tier-0 system are either provided as the ultimate food or spontaneously synthesized from the ultimate food, these autocatalytic cycles lack significance beyond suggesting that the system might show non-linear dynamics upon initiation or perturbation.

2.3 Seeds and tier-1 seed-dependent autocatalytic systems (SDASs)

The materials and energy for building a tree already exist in the environment, but without a seed buried in the soil, materials like water, carbon dioxide, minerals, and energy from sunlight will not spontaneously form a new tree. It is the seed that provides the information for organizing the flow of simple materials and energy into a structure able to grow and self-propagate. Here we will show that seeding, or at least a prototype of seeding, exists in both abiotic and biotic reaction networks, and can induce systems more complex than tier-0 systems.

A chemical seed is a chemical or a set of chemicals that are extremely rare or not produced at all in a focal food-driven microenvironment but have a non-zero chance of being introduced. For example, some reactions not included in the network might be possible in the current environment but so slow that there is a long and uncertain time before seeding occurs. Alternatively, seeds could disperse into the site from another environment that allows for different chemical reactions due to different physicochemical conditions (e.g., temperature, pressure, pH), a different food set (e.g., a surficial *versus* a submarine site), or different resident catalysts (e.g., different minerals). What network features would be needed for a seed to push the system into a new, quasistable state?

Imagine introducing a candidate seed, P, to an activated tier-0 system (Fig. S2(e)). P may react with some of the chemical species in the tier-0 system to generate new chemical species. These, in turn can result, through network expansion, in activation of a bigger reaction subnetwork (Fig. S2(f)-(h)). The chemical species and reactions that are added on top of the tier-0 system may be called a tier-1 system (Fig. S2(h)), representing the fact that one seeding event is required to

induce such a system. While P allows all chemicals in the tier-1 system to be formed, we will only consider P to be a true "seed" if the system it initiates is one that has the potential to sustain itself in an driven environment, namely one experiencing constant dilution and influx of food. Thus, the only way for a tier-1 system to persist in a driven environment is to have a network topology that allows species in the tier-0 system to be converted into a stoichiometric excess of P. This means that the tier-1 system induced by P will only be viable if it contains at least one linear combination of reactions that are collectively autocatalytic, meaning that there must be a stoichiometrically explicit net reaction equation with only tier-0 chemicals among the reactants, and with all tier-1 chemicals in the products. Such a viable tier-1 system is here defined as a seed-dependent autocatalytic system (SDAS). SDASs are similar to pRAFs in the RAF theory [22], because neither can be constructed simply from the food set. However, they are not identical because SDASs require specific stoichiometric relationships among the involved reactions while RAF theory does not consider stoichiometry as its prerequisite.

To identify SDASs, we developed a linear programming algorithm inspired by Blokhuis et al. (2020) [21]. The key criterion is that for the submatrix consisting of all columns (the *q*th to *n*th columns in Fig. 1) and rows (the *p*th to *m*th rows in Fig. 1) induced by the seed, there exists a vector of non-negative elements $\boldsymbol{x} = (x_q, x_{q+1}, \ldots, x_n)$ such that for every row in the submatrix, the dot product of this row and $\boldsymbol{x}$ is positive (Fig. 1), or

$$\sum_{j=q}^{n} x_j s_{ij} > 0 \ (x_j \geq 0) \quad \forall i \in [p, m], \tag{1}$$

where s_{ij} is the entry at the *i*th row and *j*th column of the stoichiometric matrix. Linear programming can be used to determine whether the condition described by (1) can be satisfied (Section 4.4).

For the abiotic reaction database, 12 of the 272 non-tier-0 entities can serve as singleton seeds capable of inducing a tier-1 SDAS (Table S3). Further analysis showed that all 12 seeds induce the same tier-1 system, which includes 91 entities and 220 unidirectional reactions that were absent from the tier-0 system (Table S4). We define the set of entities that can each induce the same SDAS as a clique. Thus the 12 entities, which seed the same 91-entity autocatalytic subnetwork, comprises a clique.

All chemicals in a SDAS have the potential to be produced in a sustained manner in an open environment because the SDAS must contain at least one autocatalytic core, a combination of reactions that convert a subset of the external food to all non-food chemicals in the core, which we call "members" (Peng et al. 2020). However, it should be noted that an autocatalytic core does not necessarily involve all reactions or chemicals in the SDAS. For example, a member chemical of an autocatalytic core may be constantly converted to a non-member chemical by a side reaction that is also in the SDAS. To explore the autocatalytic structure within a SDAS, we developed an integer programming procedure to identify autocatalytic cores (Section 4.5). Among the many cores that may coexist within SDASs, our algorithm can detect those satisfying a desired criterion such as containing the fewest reactions.

The smallest autocatalytic core within the tier-1 SDAS entails 10 reactions. There are 10 alternative (largely overlapping) 10-reaction autocatalytic cores and at least 30 alternative 11-reaction autocatalytic cores (Table S5(b)). As an example, one 10-reaction autocatalytic core uses CH_4, NO, and visible light as food and produces H_2CNH and infrared light as waste (Fig. 2, Table S5(a)).

For the biotic reaction database, 304 of the 4186 non-tier-0 chemical species are viable singleton seeds that can induce a tier-1 SDAS (Table S6). These 304 species belong to three cliques: Clique-1a contains 267 species that each triggers SDAS-1a (301 species; 736 reactions; Table S7); Clique-1b contains 34 species that each triggers SDAS-1b (357 species; 916 reaction; Table S8); Clique-1c contains 3 species that each triggers SDAS-1c (1414 species; 4114 reaction; Table S9). The three SDASs are nested, with SDAS-1c including the entirety of SDAS-1b, which includes the entirety of SDAS-1a. All three SDASs share the same set of minimal autocatalytic cores: there are 24 alternative 22-reaction autocatalytic cores (Table S10(b)). One of them feeds on H_2O, CO_2, and $H_4P_2O_7$ and generates H_3PO_4 and H_2O_2 as waste (Fig. 3, Table S10(a)).

Clique-1a includes simple chemicals with as few as 2 carbon atoms, such as acetylene and glycoladehyde. In contrast, every chemical in Clique-1b is a pyrimidine nucleoside or a derivative thereof and contains at least 9 carbon atoms, and Clique-1c consists of three species (NAD^+, NADH, and deamino-NAD^+) that each contains 21 carbon atoms. Since seed chemicals need to be produced somewhere to enter the system, all things being equal, we would expect

smaller organic chemicals to stochastically emerge with higher probabilities in the prebiotic environment. Therefore, we will assume for further analyses that SDAS-1a would be initiated rather than SDAS-1b or SDAS-1c.

The biotic and abiotic networks can also be used to illustrate additional features of the seeding process. First, there is nothing in the concept of seeding that requires a "seed" to be a single chemical species: simultaneous seeding of multiple chemicals might be needed to initiate a SDAS, making the seeding event even less probable. Given a reaction network and an ultimate food set, interdependent seeding of a set of non-food chemicals, $\boldsymbol{U}$, can be detected if (a) $\boldsymbol{U}$ can induce a viable SDAS that sustainably synthesizes $\boldsymbol{U}$ from the external food and (b) for any non-empty $\boldsymbol{V} \subset \boldsymbol{U}$, $\boldsymbol{V}$ cannot induce a viable SDAS supporting $\boldsymbol{V}$. We can illustrate the concept of seed interdependence with a few examples. For the abiotic reaction database, neither the formyl radical (HCO) nor O_2 can seed a viable SDAS (Table S3). Nevertheless, the set {HCO, O_2} can seed the same tier-1 SDAS as the 12-member singleton clique (Table S11). Likewise, for the biotic reaction database, neither formaldehyde (H_2CO) nor acetate (CH_3COOH) is a viable seed (Table S6), yet together they can seed a viable SDAS identical to SDAS-1a (Table S12). The interdependence between members of a seed set is conceptually linked to the fact that for the life as we know it, multiple chemical species need to be "seeded" together to allow the conversion of abiotic food to more cells. However, in the absence of a mechanism to package seed chemicals together, as happens in a cell, SDASs requiring multiple interdependent seeds would be less likely *a priori* to be triggered than those that can be activated by a singleton seed of comparable complexity.

Second, it is possible for a chemical to induce a viable SDAS that does not produce the seeding chemical itself. In an open environment, such chemicals, which we will call "pseudo-seeds," would be expected to disappear over time. For example, in the abiotic reaction database, C_2H_3 is a pseudo-seed because it can induce a viable singleton seed, OH, but the tier-1 SDAS it triggers cannot produce C_2H_3. Similarly, for the biotic reaction network, ATP is a pseudo-seed because it can induce a viable seed, UMP, which triggers SDAS-1b, but SDAS-1b cannot produce ATP. The existence of pseudo-seeds is worth noting because it means that, in the same way that scaffolds are not seen after an arch is built, key transitions during the origin of life might have

been triggered by chemicals that are no longer produced by biochemical systems, potentially confounding historical inference.

These results show that for both the abiotic and biotic reaction networks, SDASs exist. This is significant because by receiving a rare seed, a chemical ecosystem can be pushed into a new quasistable state that is distinct from, and more complex than, the state prior to seeding. In a sense, the ecosystem "remembers" a seeding event, implying that SDASs may be the earliest and simplest general mechanism of heritability, a prerequisite for evolution.

2.4 Higher-tier systems

An important feature of biological systems is that some events must happen in a specific temporal order. For example, in primary succession, lichens dominate the environment before grasses and trees, and herbivores can survive only when plant populations are large enough. Such ordering arises when later stages require conditions that can only be provided by the earlier stages. Furthermore, when systems arising at different stages coexist, they frequently form hierarchies, as illustrated by the trophic hierarchy of biological ecosystems.

Both chemical reaction networks show such hierarchical structure. Once a tier-1 SDAS is established, all members of the tier-1 system, together with the tier-0 members, are now available to "feed" additional higher-tier SDASs. The same procedures used to detect a tier-1 SDAS can be used to look for SDASs that are viable in the case that a tier-1 SDAS has been activated. If such additional tiers exist then there is a natural ordering: tier-1 SDASs feed on tier-0 systems, and tier-2 SDASs feeds on the tier-0 and tier-1 systems, etc.

For the abiotic reaction database, once the tier-1 SDAS has been induced by any of its potential seeds, there exists a 13-member clique that induces a 14-entity, 35-reaction tier-2 system containing an autocatalytic core (Table S13, Table S14). This tier-2 SDAS includes the formose reaction, which feeds on H_2CO generated by the tier-1 SDAS.

As discussed above, we will assume that of the three biotic tier-1 SDASs, SDAS-1a is more likely to be activated first. Because SDAS-1a is nested within SDAS-1b and SDAS-1c, it stands to reason that chemicals in Clique-1b and Clique-1c could seed their respective SDASs even after SDAS-1a had been activated. A more interesting question is whether activation of SDAS-1a could make additional seeds viable. In particular, we were curious about whether chemicals

smaller than those needed to seed SDAS-1b and SDAS-1c could be viable seeds after SDAS-1a had been activated.

According to this rationale, we found a 6-member clique (Clique-2a) that seeds a 56-species, 180-reaction viable tier-2 SDAS (SDAS-2a), which includes all the remaining reactions in SDAS-1b (Table S15, Table S16, Fig. 4(b),(d)). The smallest chemical in Clique-2a is the 4-carbon-atom molecule cytosine ($C_4H_5N_3O$), as contrasted with the smallest chemical in Clique-1b, the 9-carbon-atom deoxycytidine ($C_9H_{13}N_3O_4$).

Other than the three chemicals in Clique-1c, which each contains 21 carbons, we did not find any singleton seeds that can activate the remainder of SDAS-1c from SDAS-1a and the tier-0 system. However, we observed that adenine, a 5-carbon-atom molecule, together with picolinic acid, a 6-carbon-atom molecule, are interdependent seeds of a viable tier-2 SDAS (SDAS-2b), which expands SDAS-1a to SDAS-1c, and has 1113 chemicals and 3378 reactions (Fig. 4(c),(e); Table S17). Within SDAS-2b, an autocatalytic core feeds on tier-0 and tier-1 chemical species to produce ATP (Fig. 5, Fig. S4). The existence of SDAS-2b shows that even if the direct seeding of the 21-carbon NADH were unlikely, sequential seeding by three simpler molecules, such as glycolaldehyde ($CHO\text{-}CH_2OH$), adenine ($C_5H_5N_5$) and picolinic acid ($C_6H_5NO_2$), could nonetheless achieve an autocatalytic system synthesizing NADH and other complex molecules. These results show that instead of requiring a large molecule to seed a complex SDAS, seeding may be activated via intermediate SDASs that use smaller molecules as seeds.

2.5 The potential for emergent behavior following activation of a new SDAS

In a biological ecosystem, it is well-known that the introduction of a new species into a community can alter the flux of energy and matter into and through the entire ecosystem (e.g., [23]). We exploited the well-characterized chemistry of the biotic network to evaluate whether the seeding of a higher-tier SDAS might similarly affect the efficiency or productivity of a chemical ecosystem as a whole.

The increased diversity of chemicals, which is maintained once higher-tier SDASs are activated, increases the chances of finding a chemical that catalyzes reactions at lower tiers. Catalysts enable a series of intermediate chemical reactions that, in net, lower the energy barrier for producing the same products as the uncatalyzed reaction, while regenerating the catalysts. To

identify potential catalysts we can, thus, look for chemicals that provide a new path for completing a pre-existing reaction. An example is provided by the SDAS-2b chemical FAD, which opens up a new way to assimilate phosphate (Fig. 6) and would act as a catalyst provided that the composite reaction rate of the FAD-dependent path is higher than the direct reaction, which seems plausible. In an environment where phosphate was limiting, therefore, the production of FAD by SDAS-2b might increase the total mass of SDAS member chemicals at all tiers.

SDAS-2b also shows the possibility that new chemicals might allow previously unexploited food to be used or previously exploited food to be used in new, more efficient ways. For example, before SDAS-2b is activated, the ultimate food H_2S serves as a source of thiol groups but does not function as a hydrogen donor. However, due to the reaction, $NAD^+ + H_2S \leftrightarrow NADH + H^+ + S$, the presence of NAD^+ in SDAS-2b allows H_2S to also serve as a terminal hydrogen donor, allowing for thermodynamically more favorable reduction reactions. This way of exploiting H_2S was not possible before the "discovery" of a molecule like NAD^+, making this another example of a feedback to overall ecosystem function. Furthermore, the presence of NADH also allows a new carbon-fixing autocatalytic cycle (Fig. 7) that is shorter, and thus potentially more efficient, than the carbon fixation processes of SDAS-1a (Fig. 3). This closely resembles the ecological principle that a community's efficiency of resource utilization tends to increase over the course of ecological succession [24,25]. Furthermore, the new SDAS-2b carbon-fixing cycle would likely make the entire chemical ecosystem more robust to environmental change. For example, SDAS-2b would still be viable if conditions changed so as to make reaction R00602 impossible (e.g., due to loss of an environmental catalyst), despite the fact that R00602 is essential for Clique-1a to persist. This suggests that some conditions necessary for lower-tier SDASs to survive in an early stage could become unimportant once higher-tier SDASs are activated, providing an additional scaffolding mechanism besides pseudo-seeds (Section 2.3).

In addition to providing alternative or more efficient paths to access resources, chemicals produced by higher tiers may confer new emergent properties to the entire system. To illustrate this possibility, SDAS-2b includes relatively long-chain amphiphiles such as hexadecanoic acid (Table S17). Were amphiphiles to attain a sufficient concentration, they might be able to spontaneously form a membrane or vesicle. Provided that the membrane was permeable to food

and waste, which is debatable, this innovation might reduce diffusive loss of autocatalytic cores and/or enable dispersal to distant sites. Thus, it seems plausible that chemical ecosystems might be expected to complexify over time not only in the trivial sense of coming to support the persistence of a greater number of chemical species, but also in the sense of acquiring and optimizing new, emergent ecosystem-level properties.

3. DISCUSSION

3.1 Seed-dependent autocatalytic systems provide a general framework for understanding complexification and evolution before genes

In this paper, we extend the framework of chemical ecosystem ecology [3] to look for autocatalytic motifs within real chemical reaction networks. Using a stoichiometric criterion, we detected multiple potential seed-dependent autocatalytic systems (SDASs) in both abiotic and biotic reaction networks, providing evidence that both networks share many basic features related to abiogenesis.

Most SDASs we found were associated with a clique of singleton seeds. Additionally, SDASs could be triggered by interdependent seed sets or by singleton pseudo-seeds. The size of the singleton seed clique does not obviously correlate with the size of the SDAS or the complexity of its members but might predict how resistant the SDAS is to perturbations. This follows because, once activated, extinguishing a SDAS would require the simultaneous loss of all potential seeds, which would seem less likely when a clique is large.

Our finding of SDASs in real reaction networks is an important advance because it provides a mechanism for memorizing past environments in the dynamical state of a reaction network. This provides a basic mechanism of heritability in the absence of genetic polymers because a seeding event can drive a transition to a new steady state that will tend to persist through time. Furthermore, since a steady-state chemical ecosystem will produce the seeds of its activated SDASs, the dynamical state of an ecosystem could also be transported to new environments by diffusion.

The chemical networks we analyzed both manifested a trophic hierarchy that allows for the stepwise accretion of complexity. Furthermore, as illustrated by the biotic SDAS-2b, acquisition of higher-tier motifs may enhance overall ecosystem functions via catalysis, the exploitation of new resources, or compartmentalization. This means that, despite requiring the sustained production of more chemical species, activation of higher-tier SDASs might sometimes increase the steady-state concentrations of lower-tier chemicals, which would make the system's dynamical state more robust to stochastic perturbations and, due to there being a higher concentration of seeds, more easily transported to new environments.

Combined with pre-existing models that allow for analog or compositional inheritance [26–29], our SDAS theory supports an ecosystem-first model for the origin of life entailing the following steps: (a) Planetary processes, such as solar irradiation, local redox disequilibria driven by tectonics, or global redox disequilibria driven by the loss of hydrogen to space [30], generate a steady flux of food molecules or entities into local environments. (b) Chemical seeds are intermittently introduced by rare reactions or stochastic events (e.g., delivery from space), triggering low-tier SDASs that support the persistent production of new, more complex chemicals in the environment. (c) Each additional SDAS expands the pool of food available to potential SDASs at yet higher tiers. (d) Some newly seeded SDASs cause the productivity of the network as a whole to increase, for example when complex molecules catalyze more efficient pathways for extracting energy from the food set. (e) Such chemical succession continues until the ecosystem hits upon catalytic and template-guided replicated polymers, which allows for even more complexity to accrete through conventional gene-based evolution.

The idea that catalytic control by polymers could have been acquired gradually during abiogenesis gains support from the structure of SDAS-2b that we identified within the biotic network. This system includes molecules such as NADH, which is a nucleotide dimer, and coenzyme A, which can be seen as a heterogenous polymer of one amine, one beta-amino acid, one alpha-hydroxy acid, and one nucleotide. The relationship between SDAS-2b and SDAS-1a is, abstractly, similar to the relationship between genetic/catalytic polymers (e.g., DNAs, RNAs, proteins) and metabolism in cellular life (Fig. 8). Thus, the earliest appearance of short catalytic nucleic acids and peptides might have begun fortuitously with a feedback to ecosystem-level fitness. Imagining that the life had originated from the biotic network we studied (which is surely

not true), it is noteworthy that multiple types of amino acids and nucleotides are synthesized by SDAS-1a and SDAS-2b. In a prebiotic world, they might have formed short peptides or nucleic acids via condensation reactions that are suppressed in modern enzyme-driven metabolism. As a result, these hypothetical polymers might have formed tier-3 SDASs and above. It should be noted, however, that catalytic feedback from short peptides and/or nucleic acids on ecosystem function would not solve the combinatorial problem of producing sufficient polymers of a particular sequence beyond a trivial length. Adding mechanisms to chemical ecology that can overcome the problematic combinatorics of polymerization remains an important focus of future work.

3.2 Limitations of the analysis

The two reaction systems that we analyzed are small fragments of the much larger network of reactions that would have been available to early life. Analyses of carbonaceous chondrites [31] and prebiotic synthesis systems [32–34], as well as *in silico* chemical network inferences [35], suggest that the prebiotically relevant chemical reaction network includes hundreds of thousands or even millions of organic compounds. The fact that we studied such a small subset of chemical space probably limited our ability to document the true scope of stepwise, seed-driven complexification. For example, the complete prebiotic network surely allowed for more than two SDAS tiers (even without considering polymers and catalytic feedbacks). Likewise, it is very likely that different seeds could trigger exploration of non-overlapping regions of chemical space, which would allow systems to diverge over time in an historically contingent manner.

Our analyses of the biotic network have additional limitations, because the network mainly consists of enzymatic reactions. One inference from this might be that only a tiny subset of the reactions in the network we studied would occur at an appreciable rate in the absence of enzymes, implying that the true network might be less interconnected than we supposed. However, the converse argument could also be made. In the absence of catalysts optimized for particular reactions, prebiotic reaction networks might contain many more reactions that have similar rate constants resulting in greater connectedness than we assumed. Given, this uncertainty, we really do not know how representative the biotic reaction database is. Nonetheless, since the results from the biotic and abiotic networks were qualitatively similar

despite differing in size and reaction density, we suspect that our conclusions regarding the overall structure of the biotic network are robust to the presence/absence of enzymes.

The criterion of autocatalysis we used is based on network topology and stoichiometry, while thermodynamics and kinetics were ignored. Because we are imagining a chemical system driven by the constant provision of food, which can potentially drive reactions that are individually disfavored under standard conditions, the lack of thermodynamic and kinetic factors does not invalidate the autocatalytic motifs detected. However, for our theoretical framework to simulate how a reaction network may evolve over time and to further direct experimental studies by estimating appropriate reaction conditions, it would be desirable to complement this work with simulations that include realistic thermodynamic and kinetic parameters (Section 3.3).

One consequence of ignoring thermodynamics and kinetics is that some of the detected autocatalytic cores might lack viability in practice. For example, carbon fixation within the biotic SDAS-1a depends on a module consisting of two thermodynamically unfavorable reactions, namely water reducing formaldehyde to methanol ($H_2CO + 2\ H_2O \rightarrow CH_3OH + H_2O_2$) and a reverse Cannizzaro reaction where methanol and formic acid react to generate two molecules of formaldehyde ($CH_3OH + HCOOH \rightarrow 2\ H_2CO + H_2O$) (Fig. 3). It seems unlikely, therefore, that this cycle could sustain itself under natural conditions. On the other hand, we should remember that the networks we studied do not include all possible chemical reactions. For example, the tier-1 autocatalytic core could be viable if there were other prebiotic pathways available to reduce formic acid to formaldehyde, which is likely [36].

Orgel (2008) [37] argued that, because autocatalytic cycles are only viable when the rates of production of all members of the cycle are higher than their rate of loss through dilution and side-reactions, specialized catalysts are essential for any viable metabolism. Thus, he proposed that prebiotic chemical reaction networks could not sustain themselves or evolve without enzymes, and hence some kind of genetic encoding. However, recent work has shown that many core metabolic reactions that depend on enzymes in modern life can proceed at reasonable rates through the actions of metal or mineral catalysts [38–46]. Combined with our finding that many autocatalytic cycles exist within networks that are tiny compared to the immensity of all prebiotically-relevant chemistry, it seems inevitable that many autocatalytic motifs would have been viable in the prebiotic world. Thus, we believe that it would be premature to rule out the

possibility of autocatalysis-based adaptive evolutionary dynamics in the absence of enzymes and the genes that encode them.

3.3 Future theoretical and experimental research

The two most obvious extensions of the analyses we conducted would be to add realistic kinetics so that full dynamical simulations could be explored, and to expand the network to better represent the full space of chemical possibilities. Unfortunately, it is almost impossible to build these extensions on empirical data. Only a limited number of reactions have been investigated experimentally, and those represent a non-random sample, for example being focused on reactions relevant to biochemistry, combustion, or organic synthesis. Thus, the best hope for extrapolating our theory to chemistry as a whole would be to apply empirically-grounded rule-based algorithms to generate reaction networks *in silico*. Recent advances in computational systems chemistry are very promising [35,47], but much more work is needed, especially to add realistic rate constants to these reactions.

The current SDAS framework allows for seeding but does not directly explain the source of seeds. We know that a seed molecule can only arise by a chemical reaction happening somewhere, implying that the reactions that generate seeds ought to be included in the reaction network. One approach would be to include kinetics and use models, such as the Gillespie algorithm [48], that can translate low reaction rates into discrete stochastic events. Another approach might be to utilize two or more distinct networks of allowed reactions, each interpreted as the set of feasible reactions in a particular environmental context, and then allow rare transfer of materials from one environment into another. Finally, the ultimately preferable approach would be to combine explicit kinetics and explicit spatial structure so that improbable local reactions and rare dispersal events can be allowed to simultaneously affect chemical ecosystem dynamics.

In addition to further theoretical research, we hope that our work will stimulate laboratory experiments to test our main conclusions. If our model is correct, SDASs might be a common feature of "messy" chemical reaction systems. Therefore, an obvious choice is to drive a chemical system with a flux of simple food chemicals using a continuously-stirred tank reactor (CSTR), or its poor-man's equivalent, a serial batch-transfer-with-dilution experiment [49–51], and intermittently introduce small quantities of diverse candidate seed chemicals, while looking

for evidence that seeding can induce a sustained change. We predict that it will be possible to find chemicals or sets of chemicals that are more complex than any yet in the system that, when added transiently, can cause the system to transition to a new steady-state composition. Moreover, by maintaining such experiments for long periods, we believe it might be possible to detect the spontaneous emergence of evolution-like dynamics in the laboratory [51,52]. Such work would significantly advance our understanding of the origin of life by explaining how the stepwise accretion of complexity can occur even prior to the appearance of gene-based inheritance.

Our SDAS theory only requires that the system to which it is applied can be expressed in the form of a system of reaction equations with stoichiometry, with certain entities designated as the ultimate food. As a result, it may have significance beyond the origin-of-life field by providing a general framework to describe and analyze other complex systems that contain autocatalytic motifs, including biological ecosystems, social networks, and economies, which share common features such as resilience to perturbations, hierarchical organization, and scaffolding. We hope, therefore, that the algorithms and concepts developed here will be explored in other complex systems to evaluate whether seed-dependency, memory, and adaptive evolvability might also apply in these cases.

4. MATERIALS AND METHODS

4.1 Preprocessing databases of reactions

4.1.1 Abiotic reaction database

The reaction network assembled by Adam et al. (2021) [4] includes the following categories: free radical reactions, mineral geochemical reactions, amino acid production, chloride radical and polar reactions, nitrile radical and polar reactions, RNA nucleotide assembly, nuclear decay, and physicochemical reactions. We processed this database by the following steps.

First, we excluded the nuclear decay reactions because we did not plan to put radioactive atoms into the ultimate food set.

Second, with kind help from Dr. Zachary R. Adam and Dr. Albert C. Fahrenbach, we deleted duplicate reactions, added a few new reactions that were not in the original database, balanced some reaction equations, and excluded the reactions without clear stoichiometry. This is because our method requires stoichiometry of reactions.

Third, we added the formose reaction into the database. According to Breslow's mechanism [6], the formose reaction is driven by aldol and retro-aldol reactions and aldose-ketose isomerization. In combination these reactions allow low-carbon-number monosaccharides to generate high-carbon-number monosaccharides. Therefore, we added reversible aldol reactions and reversible aldose-ketose isomerization among formaldehyde, glycolaldehyde, and monosaccharides with no more than 8 carbon atoms. Optical isomers were not distinguished from each other. Formaldehyde dimerization was not added because it is very slow and its reaction mechanism is unclear but surely neither aldol/retro-aldol reaction nor aldose-ketose isomerization.

Fourth, every reaction labeled reversible was split into two unidirectional reactions.

4.1.2 Biotic reaction database

We processed the reaction database curated by Xavier et al. (2020) [5] to obtain the biotic reaction database by the following steps.

First, we removed all reactions involving chemical species that do not have specific molecular mass, such as reduced ferredoxin (KEGG: C00138), acyl-carrier protein (KEGG: C00229),

starch (KEGG: C00369), and long-chain aldehyde (KEGG: C00609), because they sometimes result in "fake" stoichiometric relationships. For example, the reaction: starch + H_2O ↔ dextrin + starch (KEGG: R02108) would make starch an infinite source of dextrin as long as H_2O is provided. Considering that glycans have both "G"-started (meaning "Glycan") and "C"-started (meaning "Compound") KEGG entries, the reactions involving "G"-started entries were also removed because these reactions are redundant.

Second, we added some obviously spontaneous reactions that were missing, such as $H_2O \leftrightarrow H^+ + OH^-$ and $H_2CO_3 \leftrightarrow H^+ + HCO_3^-$.

Third, we added the reaction KEGG R06974 into the biotic reaction database. This reaction is very similar to the reaction R06975 (Fig. S5): both reactions use HCOOH as the carbon donor to add a -CHO to -NH2 and form a -NH-CHO with ATP hydrolysis providing energy for the reaction. However, R06975 is in the network curated by Xavier et al. (2020) [5] while R06974 is not, presumably because the annotations of R06974 in the KEGG database are not as detailed as those of R06975, and thus R06974 was filtered out. We also conducted analyses without R06974. In that case, SDAS-1a, SDAS-1b, and SDAS-2a were not affected but SDAS-1c and SDAS-2b were missing. However, we opted to present results that included the plausible reaction R06974 so as to better illustrate the potential for ecosystem-level feedback without resorting to analyzing the entirety of KEGG.

Fourth, as all reactions in the KEGG biochemical reaction database are labeled reversible, every reaction was split into two unidirectional reactions. The reaction following the forward direction specified in the KEGG database has a suffix ".a" to its entry, and that of the reverse direction has a suffix ".b".

Fifth, the reactions that are labeled as multi-step were removed because each step is already a reaction in the database. Although keeping these multi-step reactions may not have big impact on the detection of SDAS existence, decreasing the number of reactions in the stoichiometric matrix should help accelerate the computation.

4.2 Network expansion

The set $\boldsymbol{R} = \{r_1, r_2, \ldots, r_i, \ldots, r_n\}$ is a set of multiple reactions r_i's that are allowed. Each r_i specifies reactants and products, and the union of all reactants and products across all r_i's is the

maximum set of chemical species $\boldsymbol{S} = \{k_1, k_2, \ldots, k_j, \ldots, k_m\}$. We define an operation called full network expansion, $\Xi(\boldsymbol{S_O}, \boldsymbol{R}) = (\boldsymbol{S_E}, \boldsymbol{R_E})$, where $\boldsymbol{S_O}$ is the subset of $\boldsymbol{S}$ where the expansion starts, $\boldsymbol{S_E}$ the set of chemical species resulting from the expansion, and $\boldsymbol{R_E}$ the expanded set of reactions resulting from the expansion. The expansion is conducted as follows:

(i) Let $\boldsymbol{R_E} = \emptyset$; define a set of reactions $\boldsymbol{R'} = \boldsymbol{R}$; let $\boldsymbol{S_E} = \boldsymbol{S_O}$.

(ii) Define a temporary set of chemical species $\boldsymbol{S'} = \emptyset$.

(iii) For a reaction r_i in $\boldsymbol{R'}$, check if the reactants required by r_i are all present in $\boldsymbol{S_E}$; if so, move r_i from $\boldsymbol{R'}$ to $\boldsymbol{R_E}$, and scan through the products of r_i to add the chemical species that are not in $\boldsymbol{S_E}$ to $\boldsymbol{S'}$. Do this for all reactions in $\boldsymbol{R'}$. Then add all chemical species in $\boldsymbol{S'}$ to $\boldsymbol{S_E}$. If during this step, no reaction in $\boldsymbol{R'}$ is moved, then the expansion is finished; otherwise, proceed to (ii).

4.3 Identifying cliques

Let us assume that $\boldsymbol{S_F}$ is a set of chemical species resulting from a full expansion within the set of allowed reactions $\boldsymbol{R}$ from a set of ultimate food $\boldsymbol{S_{UF}}$. Two non-empty seed sets of non-food chemical species $\boldsymbol{S_{P1}}$ and $\boldsymbol{S_{P2}}$ are said to be in the same clique if (a) $\Xi(\boldsymbol{S_F} \cup \boldsymbol{S_{P1}}, \boldsymbol{R}) = \Xi(\boldsymbol{S_F} \cup \boldsymbol{S_{P2}}, \boldsymbol{R})$, and (b) for any proper subset $\boldsymbol{S'_{P1}}$ of $\boldsymbol{S_{P1}}$ and any proper subset $\boldsymbol{S'_{P2}}$ of $\boldsymbol{S_{P2}}$, $\Xi(\boldsymbol{S_F} \cup \boldsymbol{S'_{P1}}, \boldsymbol{R}) \neq \Xi(\boldsymbol{S_F} \cup \boldsymbol{S_{P1}}, \boldsymbol{R})$ and $\Xi(\boldsymbol{S_F} \cup \boldsymbol{S'_{P2}}, \boldsymbol{R}) \neq \Xi(\boldsymbol{S_F} \cup \boldsymbol{S_{P2}}, \boldsymbol{R})$.

In this paper, we only investigated the cliques consisting of seeds that are individual chemical species (i.e., singleton seeds). Nonetheless, the principle could be expanded to potential seed sets comprising more than one chemical species.

4.4 Detecting Seed-Dependent Autocatalytic Systems (SDASs) by linear programming

Let us assume that a $(p\text{-}1)\times(q\text{-}1)$ stoichiometric matrix, where each row represents a chemical species and each column represents a unidirectional reaction, results from a full expansion within the set of allowed reactions $\boldsymbol{R}$. The row labels of this $(p\text{-}1)\times(q\text{-}1)$ stoichiometric matrix form a chemical species set $\boldsymbol{S_F} = \{k_1, k_2, \ldots, k_{p-1}\}$, which is defined as the external food for this detection process. Now we select a non-empty set of non-food chemical species $\boldsymbol{S_P} = \{k_p, k_{p+1}, \ldots, k_{p+h}\}$ to serve as the potential seed set. For example, we can treat all chemicals not in the external food as potential seeds and search through these one at a time. Then we conduct a

full expansion from the food set and the seed set, $\Xi(\boldsymbol{S_F} \cup \boldsymbol{S_P}, \boldsymbol{R}) = (\boldsymbol{S_{FP}}, \boldsymbol{R_{FP}})$, generating $\boldsymbol{S_{FP}} = \{k_1, k_2, \ldots, k_m\}$ and $\boldsymbol{R_{FP}} = \{r_1, r_2, \ldots, r_n\}$.

A SDAS feeding on $\boldsymbol{S_F}$ exists if there is a vector of non-negative elements $\boldsymbol{x} = (x_q, x_{q+1}, \ldots, x_n)$ such that

$$\sum_{j=q}^{n} x_j s_{ij} > 0 \ (x_j \geq 0) \quad \forall i \in [p, m], \tag{1}$$

where s_{ij} is the entry at the *i*th row and *j*th column of the stoichiometric matrix.

To determine whether such an $\boldsymbol{x}$ exists, we used the linear programming tool provided by SciPy v1.6.2 (https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.linprog.html). In addition to the constraint set by (1), this linear programming tool requires an objective function. Since we knew that the growth of an autocatalytic system feeding on the external food is unbounded when the external food is unlimited (Fig. 1), we could set an objective function to find the maximum $\sum_{j=q}^{n} x_j s_{ij}$ for an internal species k_i ($i \in [p, m]$). If this objective function was found to be unbounded, we knew that a feasible region constrained by (1) must exist, indicating that a SDAS must exist. We simply let $i = p$ and set

$$\max_{x_q, \ \ldots, \ x_n} \sum_{j=q}^{n} x_j s_{pj} \tag{2}$$

as the objective function.

We used the "highs" method [53] to confirm the existence of SDASs. Once the SDAS was confirmed to exist, we ran the integer programming process, described in Section 4.5, to find autocatalytic cores within the SDAS, subject to further specific constraints.

4.5 Detecting minimum-reaction-number autocatalytic cores by integer programming

We can use linear programming to identify a SDAS, but we also desire to find small autocatalytic cores within each SDAS, because these are easier to visualize and could potentially guide future experimental studies. Specifically, considering that there were likely few types of catalysts in the prebiotic world, we focus on finding the autocatalytic cores with as few reactions as possible. As a result, we want to minimize the number of positive components of $\boldsymbol{x}$ while the

reactions corresponding to positive x_j's still form an autocatalytic core feeding on the external food. In fact, the original SDAS may contain multiple autocatalytic cores. In this section we describe a method based on integer linear programming to enumerate small-cardinality autocatalytic cores within a given SDAS.

If there exists a vector $\boldsymbol{x}$, such that (1) is fulfilled, then by scaling $\boldsymbol{x}$, we may ensure that

$$\sum_{j=q}^{n} x_j s_{ij} \geq 1 \left(x_j \geq 0\right) \quad \forall i \in [p, m]. \tag{3}$$

Thus, without loss of generality, we may use (3) as the constraint. There may be many possible $\boldsymbol{x}$'s that fulfill (3), and we used integer programming to seek and enumerate SDASs with desirable properties.

To find smaller systems, we seek a set of columns $\boldsymbol{T} \subseteq [q, n]$ such that if $j \in \boldsymbol{T}$ and $\exists i \in [p, m]$ which makes $s_{ij} \neq 0$, then $\exists x_j \in \mathbb{R}_+^{n-q+1}$ such that $\sum_{j=q}^{n} x_j s_{ij} \geq 1$. Of course, the set $\boldsymbol{T}$ should have $|\boldsymbol{T}| \geq 1$. Finding such a set $\boldsymbol{T}$ can be accomplished in a systematic manner by seeking solutions to a linear-inequality system wherein some of the variables are required to take integer values.

In the formulation, we use binary variables $z_j \in \{0, 1\}$ that take the value 1 if and only if column $j \in [q, n]$ is in the set $\boldsymbol{T}$, and we will minimize $\sum_{j=q}^{n} z_j$ to minimize the cardinality of the selected set. Because an autocatalytic core must have at least one reaction, it is obvious that $\sum_{j=q}^{n} z_j \geq 1$.

Let β_j be the upper bound on the number of the reaction r_j that can occur in an autocatalytic core. We must enforce that if column j is not selected for the autocatalytic core (i.e., $z_j = 0$), then its level of reaction x_j must also be zero, which is done with the algebraic constraints $x_j \leq \beta_j z_j$.

For the selected set of reactions $\boldsymbol{T}$ to be an autocatalytic core, let us first define a set $\boldsymbol{\Omega_T}$ that contains and only contains the row indices of all chemical species that are involved in the reactions in $\boldsymbol{T}$ and are not external species (i.e., $\boldsymbol{\Omega_T} \subseteq [p, m]$), then we must make sure that

$$\sum_{j=q}^{n} x_j s_{ij} \geq 1 \quad \forall i \in \boldsymbol{\Omega_T}. \tag{4}$$

This is done by introducing additional binary variables $y_i \in \{0, 1\}$ ($i \in [p, m]$) indicating if species k_i is involved in the autocatalytic core represented by the positive components of the vector $\boldsymbol{z} = (z_q, z_{q+1}, \dots, z_n)$. Then, the following set of linear inequalities accomplish the conditions in (4)

$$\sum_{j=q}^{n} x_j s_{ij} \geq 1 - M_i(1 - y_i) \quad \forall i \in [p, m], \tag{5}$$

where

$$M_i = 1 - \sum_{j=q:s_{ij}<0}^{n} \beta_j s_{ij} \,. \tag{6}$$

To understand how (5) works, imagine that k_i is involved in the autocatalytic core, then y_i =1 and (5) is equivalent to (4). In contrast, if k_i is not involved in the autocatalytic core, then y_i = 0 and (5) is equivalent to

$$\sum_{j=q}^{n} x_j s_{ij} \geq \sum_{j=q:s_{ij}<0}^{n} \beta_j s_{ij} \,. \tag{7}$$

Because β_j is the upper bound on x_j, (7) should always hold and thus it is redundant in the linear system, representing the fact that (4) does not need to be considered for a k_i that is not included in the autocatalytic core.

It is necessary to link a reaction and the chemical species that are involved in the reaction. If a reaction r_j ($j \in [q, n]$) is selected for an autocatalytic core (i.e., z_j = 1), any chemical species k_i that is involved in r_j must also exist in the autocatalytic core (i.e., y_i = 1). Therefore, for any reaction r_j ($j \in [q, n]$), we define a set $\boldsymbol{\Omega_j}$ that contains and only contains the row indices of all chemical species that are involved in r_j and are not external species (i.e., $\boldsymbol{\Omega_j} \subseteq [p, m]$). Then we apply the constraint

$$y_i \geq z_j \quad \forall j \in [q, n], \forall i \in \boldsymbol{\Omega_j} \tag{8}$$

which guarantees that once r_j is included in an autocatalytic core (i.e., z_j = 1), k_i need to be sustainably synthesized (i.e., y_i = 1).

This gives a full integer programming formulation for finding a minimum-cardinality autocatalytic core among the reactions $\{r_q, r_{q+1}, \ldots, r_n\}$. We set the integer programming problem as to find

$$\min_{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}} \sum_{j=q}^{n} z_j \tag{9}$$

which is constrained by

$$\begin{cases} x_j \geq 0 \quad \forall j \in [q,n] \\ z_j \in \{0,1\} \quad \forall j \in [q,n] \\ y_i \in \{0,1\} \quad \forall i \in [p,m] \\ \sum_{j=q}^{n} z_j \geq 1 \\ x_j \leq \beta_j z_j \quad \forall j \in [q,n] \\ \sum_{j=q}^{n} x_j s_{ij} \geq 1 - M_i(1-y_i) \quad \forall i \in [p,m] \\ y_i \geq z_j \quad \forall j \in [q,n], \forall i \in \boldsymbol{\Omega_j} \end{cases} . \tag{10}$$

This formulation can be solved computationally using a state-of-the-art integer programming software, such as Gurobi. The positive elements of the binary solution vector $\boldsymbol{z}$ indicate the reaction set $\boldsymbol{T}$ in a minimum-cardinality autocatalytic core chosen from the set of reactions $\{r_q, r_{q+1}, \ldots, r_n\}$.

It should be noted that multiple autocatalytic cores with the same number of reactions may exist, and our integer programming can enumerate all minimum-cardinality autocatalytic cores by solving a sequence of integer programs. After a binary solution vector $\boldsymbol{z}$ and its associated reaction set $\boldsymbol{T}$ are identified, the constraint $\sum_{j=q:j\in\boldsymbol{T}}^{n} z_j \leq |\boldsymbol{T}| - 1$ may be added to (10), and then the process repeats. Furthermore, if we want to find minimum-cardinality autocatalytic cores with at least D reactions, we can do it by replacing $\sum_{j=q}^{n} z_j \geq 1$ with $\sum_{j=q}^{n} z_j \geq D$ in (10).

ACKNOWLEDGEMENTS

We thank Dr. Zachary R. Adam and Dr. Albert C. Fahrenbach for their kind help with curating the abiotic reaction database and the following for useful discussions: Alyssa Adams, Stephanie Colón-Santos, Emily Dolson, Praful Gagrani, Juan Perez Mercader, Alex Plum, Daniel Segrè, D. Eric Smith, Lena Vincent, and participants in the Evolving Chemical Systems workshop at the Santa Fe Institute (Nov. 2019). We thank NASA-NSF CESPOoL (Chemical Ecosystem Selection Paradigm for the Origins of Life) Ideas Lab grant (NASA-80NSSC17K0296) and University of Wisconsin Vice-Chancellor for Research and Graduate Education for supporting this project.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

REFERENCES

[1] J. Ladyman, J. Lambert, and K. Wiesner, *What Is a Complex System?*, Euro. Jnl. Phil. Sci. **3**, 33 (2013).

[2] M. Mitchell and M. Newman, *Complex Systems Theory and Evolution*, in *Encyclopedia of Evolution*, edited by M. Pagel (Oxford University Press, New York, 2002).

[3] Z. Peng, A. M. Plum, P. Gagrani, and D. A. Baum, *An Ecological Framework for the Analysis of Prebiotic Chemical Reaction Networks*, J. Theor. Biol. **507**, 110451 (2020).

[4] Z. R. Adam, A. C. Fahrenbach, S. M. Jacobson, B. Kacar, and D. Y. Zubarev, *Radiolysis Generates a Complex Organosynthetic Chemical Network*, Sci. Rep. **11**, 1 (2021).

[5] J. C. Xavier, W. Hordijk, S. Kauffman, M. Steel, and W. F. Martin, *Autocatalytic Chemical Networks at the Origin of Metabolism*, Proc. R. Soc. B. **287**, 20192377 (2020).

[6] R. Breslow, *On the Mechanism of the Formose Reaction,* Tetrahedron Letters **1**, 22 (1959).

[7] H. J. (Jim) Cleaves, *Formose Reaction*, in *Encyclopedia of Astrobiology*, edited by M. Gargaud, R. Amils, J. C. Quintanilla, H. J. (Jim) Cleaves, W. M. Irvine, D. L. Pinti, and M. Viso (Springer, Berlin, Heidelberg, 2011), pp. 600–605.

[8] M. Haas, S. Lamour, S. B. Christ, and O. Trapp, *Mineral-Mediated Carbohydrate Synthesis by Mechanical Forces in a Primordial Geochemical Setting*, Commun. Chem. **3**, 1 (2020).

[9] M. Kanehisa, *Toward Understanding the Origin and Evolution of Cellular Organisms*, Protein Sci. **28**, 1947 (2019).

[10] M. Kanehisa, M. Furumichi, Y. Sato, M. Ishiguro-Watanabe, and M. Tanabe, *KEGG: Integrating Viruses and Cellular Organisms*, Nucleic Acids Res. **49**, D545 (2021).

[11] M. Kanehisa and S. Goto, *KEGG: Kyoto Encyclopedia of Genes and Genomes*, Nucleic Acids Res. **28**, 27 (2000).

[12] B. Damer and D. Deamer, *Coupled Phases and Combinatorial Selection in Fluctuating Hydrothermal Pools: A Scenario to Guide Experimental Approaches to the Origin of Cellular Life,* Life **5**, 872 (2015).

[13] D. J. Donaldson, H. Tervahattu, A. F. Tuck, and V. Vaida, *Organic Aerosols and the Origin of Life: An Hypothesis*, Orig. Life Evol. Biosph. **34**, 57 (2004).

[14] R. Lathe, *Fast Tidal Cycling and the Origin of Life*, Icarus **168**, 18 (2004).

[15] M. R. Marín-Yaseli, E. González-Toril, C. Mompeán, and M. Ruiz-Bermejo, *The Role of Aqueous Aerosols in the "Glyoxylate Scenario": An Experimental Approach*, Chem. Eur. J. **22**, 12785 (2016).

[16] W. Martin, J. Baross, D. Kelley, and M. J. Russell, *Hydrothermal Vents and the Origin of Life*, Nat. Rev. Microbiol. **6**, 11 (2008).

[17] S. Maruyama, K. Kurokawa, T. Ebisuzaki, Y. Sawaki, K. Suda, and M. Santosh, *Nine Requirements for the Origin of Earth's Life: Not at the Hydrothermal Vent, but in a Nuclear Geyser System*, Geoscience Frontiers **10**, 1337 (2019).

[18] V. Sojo, B. Herschy, A. Whicher, E. Camprubí, and N. Lane, *The Origin of Life in Alkaline Hydrothermal Vents*, Astrobiology **16**, 181 (2016).

[19] G. Wächtershäuser, *Before Enzymes and Templates: Theory of Surface Metabolism.*, Microbiol. Rev. **52**, 452 (1988).

[20] F. Westall, K. Hickman-Lewis, N. Hinman, P. Gautret, K. a. Campbell, J. g. Bréhéret, F. Foucher, A. Hubert, S. Sorieul, A. v. Dass, T. p. Kee, T. Georgelin, and A. Brack, *A Hydrothermal-Sedimentary Context for the Origin of Life*, Astrobiology **18**, 259 (2018).

[21] A. Blokhuis, D. Lacoste, and P. Nghe, *Universal Motifs and the Diversity of Autocatalytic Systems*, Proc. Natl. Acad. Sci. U.S.A. **117**, 25230 (2020).

[22] M. Steel, J. C. Xavier, and D. H. Huson, *The Structure of Autocatalytic Networks, with Application to Early Biochemistry*, J. R. Soc. Interface **17**, 20200488 (2020).

[23] T. J. Smith, K. G. Boto, S. D. Frusher, and R. L. Giddins, *Keystone Species and Mangrove Forest Dynamics: The Influence of Burrowing by Crabs on Soil Nutrient Status and Forest Productivity*, Estuar. Coast. Shelf Sci. **33**, 419 (1991).

[24] R. MacArthur, *Species Packing and Competitive Equilibrium for Many Species*, Theor. Popul. Biol. **1**, 1 (1970).

[25] R. MacArthur, *Species Packing, and What Competition Minimizes*, Proc. Natl. Acad. Sci. U.S.A. **64**, 1369 (1969).

[26] D. A. Baum, *The Origin and Early Evolution of Life in Chemical Composition Space*, J. Theor. Biol. **456**, 295 (2018).

[27] V. Vasas, C. Fernando, M. Santos, S. Kauffman, and E. Szathmáry, *Evolution before Genes*, Biol. Direct **7**, 1 (2012).

[28] D. Segré, D. Ben-Eli, and D. Lancet, *Compositional Genomes: Prebiotic Information Transfer in Mutually Catalytic Noncovalent Assemblies*, Proc. Natl. Acad. Sci. U.S.A. **97**, 4112 (2000).

[29] D. Segré, B. Shenhav, R. Kafri, and D. Lancet, *The Molecular Roots of Compositional Inheritance*, J. Theor. Biol. **213**, 481 (2001).

[30] E. Smith and H. J. Morowitz, *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere* (Cambridge University Press, Cambridge, 2016).

[31] P. Schmitt-Kopplin, Z. Gabelica, R. D. Gougeon, A. Fekete, B. Kanawati, M. Harir, I. Gebefuegi, G. Eckel, and N. Hertkorn, *High Molecular Diversity of Extraterrestrial Organic Matter in Murchison Meteorite Revealed 40 Years after Its Fall*, Proc. Natl. Acad. Sci. U.S.A. **107**, 2763 (2010).

[32] V. Vuitton, J.-Y. Bonnet, M. Frisari, R. Thissen, E. Quirico, O. Dutuit, B. Schmitt, L. L. Roy, N. Fray, H. Cottin, E. Sciamma-O'Brien, N. Carrasco, and C. Szopa, *Very High Resolution Mass Spectrometry of HCN Polymers and Tholins*, Faraday Discuss. **147**, 495 (2010).

[33] E. Wollrab, S. Scherer, F. Aubriet, V. Carré, T. Carlomagno, L. Codutti, and A. Ott, *Chemical Analysis of a "Miller-Type" Complex Prebiotic Broth Part I: Chemical Diversity, Oxygen and Nitrogen Based Polymers*, Orig. Life Evol. Biosph. **46**, 149 (2016).

[34] S. Scherer, E. Wollrab, L. Codutti, T. Carlomagno, S. G. da Costa, A. Volkmer, A. Bronja, O. J. Schmitz, and A. Ott, *Chemical Analysis of a "Miller-Type" Complex Prebiotic Broth Part II: Gas, Oil, Water and the Oil/Water-Interface*, Orig. Life Evol. Biosph. **47**, 381 (2017).

[35] A. Wołos, R. Roszak, A. Żądło-Dobrowolska, W. Beker, B. Mikulak-Klucznik, G. Spólnik, M. Dygas, S. Szymkuć, and B. A. Grzybowski, *Synthetic Connectivity, Emergence, and Self-Regeneration in the Network of Prebiotic Chemistry*, Science **369**, (2020).

[36] H. J. Cleaves, *The Prebiotic Geochemistry of Formaldehyde*, Precambrian Res. **164**, 111 (2008).

[37] L. E. Orgel, *The Implausibility of Metabolic Cycles on the Prebiotic Earth*, PLOS Biol. **6**, e18 (2008).

[38] K. B. Muchowska, S. J. Varma, E. Chevallot-Beroux, L. Lethuillier-Karl, G. Li, and J. Moran, *Metals Promote Sequences of the Reverse Krebs Cycle*, Nat. Ecol. Evol. **1**, 1716 (2017).

[39] S. J. Varma, K. B. Muchowska, P. Chatelain, and J. Moran, *Native Iron Reduces CO2 to Intermediates and End-Products of the Acetyl-CoA Pathway*, Nat. Ecol. Evol. **2**, 1019 (2018).

[40] K. B. Muchowska, S. J. Varma, and J. Moran, *Synthesis and Breakdown of Universal Metabolic Precursors Promoted by Iron*, Nature **569**, 104 (2019).

[41] K. B. Muchowska, S. J. Varma, and J. Moran, *Nonenzymatic Metabolic Reactions and Life's Origins*, Chem. Rev. **120**, 7708 (2020).

[42] M. A. Keller, D. Kampjut, S. A. Harrison, and M. Ralser, *Sulfate Radicals Enable a Non-Enzymatic Krebs Cycle Precursor*, Nat. Ecol. Evol. **1**, 4 (2017).

[43] M. Ralser, *An Appeal to Magic? The Discovery of a Non-Enzymatic Metabolism and Its Role in the Origins of Life*, Biochem. J. **475**, 2577 (2018).

[44] M. A. Keller, A. V. Turchyn, and M. Ralser, *Non-Enzymatic Glycolysis and Pentose Phosphate Pathway-like Reactions in a Plausible Archean Ocean*, Mol. Syst. Biol. **10**, 725 (2014).

[45] M. A. Keller, A. Zylstra, C. Castro, A. V. Turchyn, J. L. Griffin, and M. Ralser, *Conditional Iron and PH-Dependent Activity of a Non-Enzymatic Glycolysis and Pentose Phosphate Pathway*, Sci. Adv. **2**, e1501235 (2016).

[46] J. Sobotta, T. Geisberger, C. Moosmann, C. M. Scheidler, W. Eisenreich, G. Wächtershäuser, and C. Huber, *A Possible Primordial Acetyleno/Carboxydotrophic Core Metabolism*, Life **10**, 35 (2020).

[47] J. L. Andersen, C. Flamm, D. Merkle, and P. F. Stadler, *Chemical Transformation Motifs—Modelling Pathways as Integer Hyperflows*, IEEE/ACM Trans. Comput. Biol. Bioinform. **16**, 510 (2019).

[48] D. T. Gillespie, *Stochastic Simulation of Chemical Kinetics*, Annu. Rev. Phys. Chem. **58**, 35 (2007).

[49] S. Colón-Santos, G. J. T. Cooper, and L. Cronin, *Taming the Combinatorial Explosion of the Formose Reaction via Recursion within Mineral Environments*, ChemSystemsChem **1**, e1900033 (2019).

[50] A. J. Surman, M. Rodriguez-Garcia, Y. M. Abul-Haija, G. J. T. Cooper, P. S. Gromski, R. Turk-MacLeod, M. Mullin, C. Mathis, S. I. Walker, and L. Cronin, *Environmental Control Programs the Emergence of Distinct Functional Ensembles from Unconstrained Chemical Reactions*, Proc. Natl. Acad. Sci. U.S.A. **116**, 5387 (2019).

[51] Vincent, Berg, Krismer, Saghafi, Cosby, Sankari, Vetsigian, Ii, and Baum, *Chemical Ecosystem Selection on Mineral Surfaces Reveals Long-Term Dynamics Consistent with the Spontaneous Emergence of Mutual Catalysis*, Life **9**, 80 (2019).

[52] D. A. Baum and K. Vetsigian, *An Experimental Framework for Generating Evolvable Chemical Systems in the Laboratory*, Orig. Life Evol. Biosph. **47**, 481 (2017).

[53] Q. Huangfu and J. A. J. Hall, *Parallelizing the Dual Revised Simplex Method*, Math. Prog. Comp. **10**, 119 (2018).

FIGURE LEGENDS

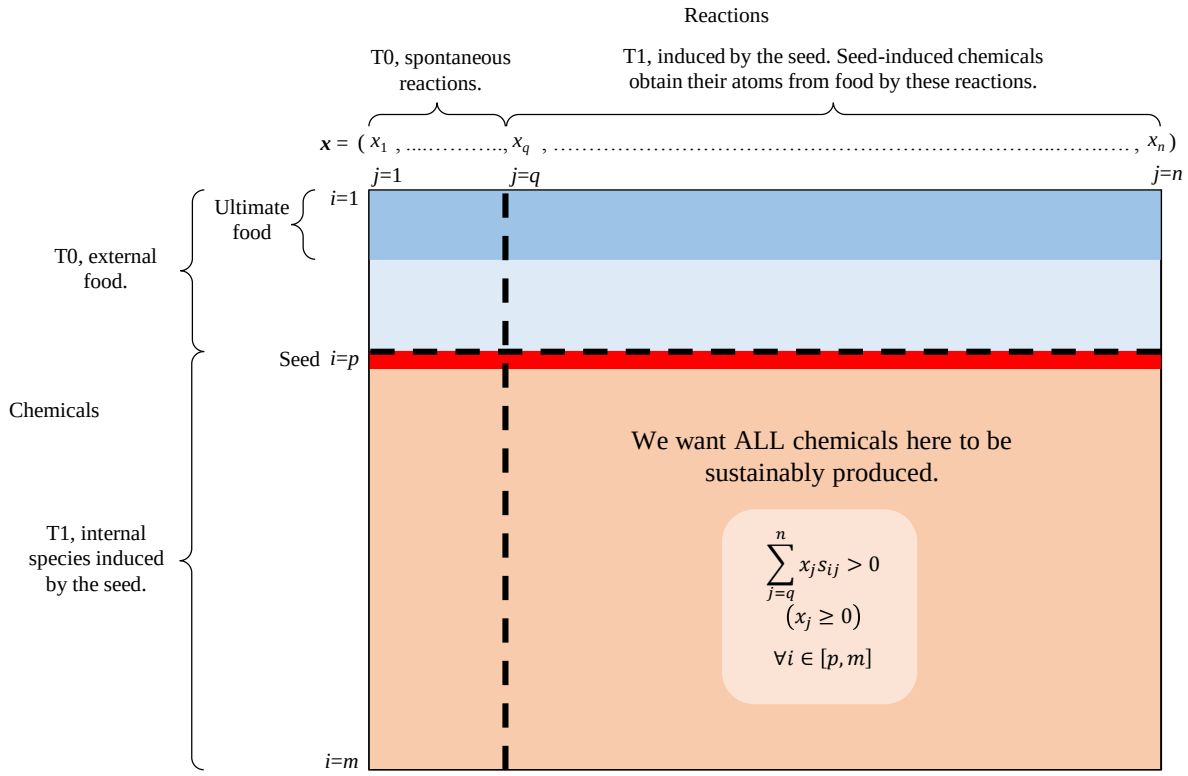

Fig. 1. SDAS detection in a stoichiometric matrix with rows representing chemical species and columns representing reactions. A stoichiometric matrix starting with an ultimate food set and a seed set can be split into four submatrices. The upper left submatrix includes reactions that would be triggered by the ultimate food. The chemicals in this submatrix represent the external food that a SDAS will need to sustain itself on. The upper right submatrix includes additional reactions induced by seeding. The lower left submatrix includes additional chemical species induced by seeding and should be empty (because these chemicals are not involved in the spontaneous reactions). The lower right submatrix includes new internal chemical species and new reactions that are induced by seeding. To determine if this case depicts a viable SDAS, we need to determine if all chemical species in the lower right submatrix can be sustainably produced by consuming the external food. Therefore, we want to find a linear combination of the reactions in this submatrix that permits every internal chemical species to have a positive net change. x_j is the number of the jth reaction in the linear combination. s_{ij} is stoichiometric coefficient of the ith chemical species in the jth reaction; s_{ij} is negative if the ith chemical species is a reactant, and positive if the ith chemical species is a product.

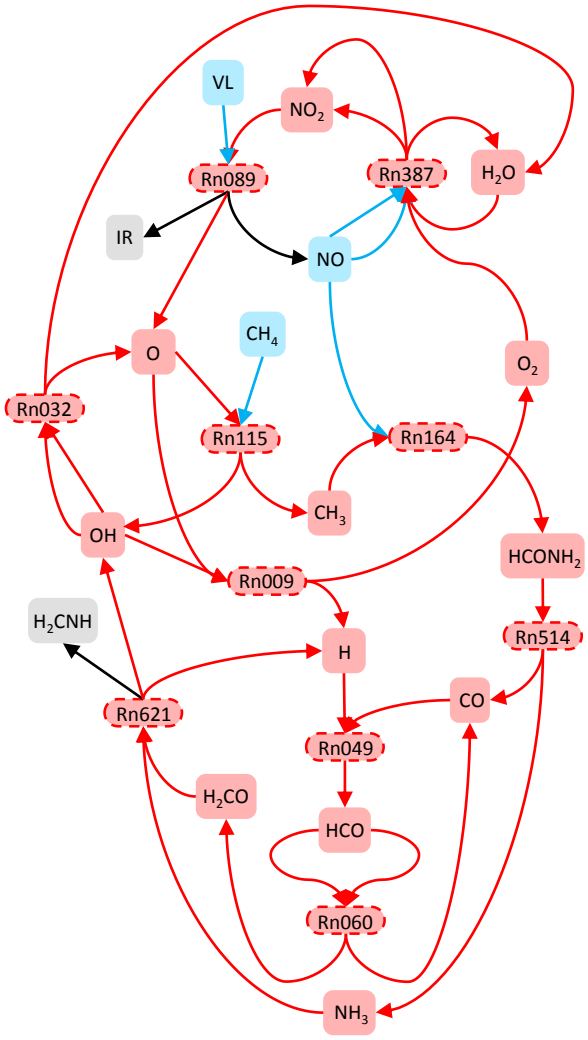

Fig. 2. A minimal autocatalytic core identified within the tier-1 SDAS in the abiotic reaction network. Rounded-corner box: chemical species. Red box: member of the autocatalytic core. Cyan box: food of the autocatalytic core. Grey box: waste of the autocatalytic core. Oval with dashed border: reaction. Cyan arrow: food consumption. Black arrow: waste production. VL: visible light. IR: infrared.

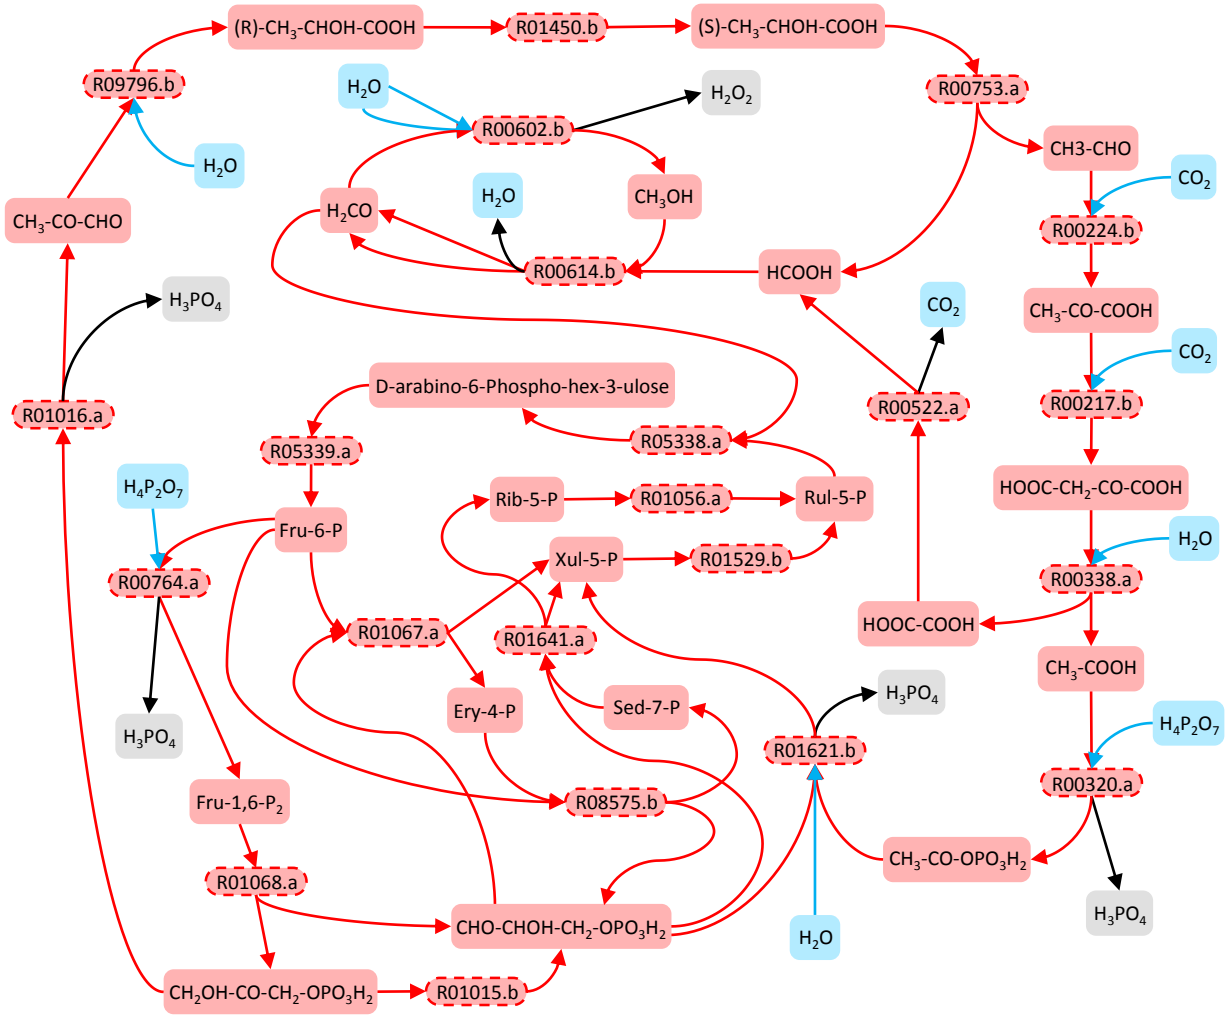

Fig. 3. A minimal autocatalytic core identified within the tier-1 SDASs in the biotic reaction network. Rounded-corner box: chemical species. Red box: member of the autocatalytic core. Cyan box: food of the autocatalytic core. Grey box: waste of the autocatalytic core. Oval with dashed border: reaction. Cyan arrow: food consumption. Black arrow: waste production. Note that some food of the entire autocatalytic core (e.g., CO_2) can be the waste of a reaction step (e.g., R00522.a).

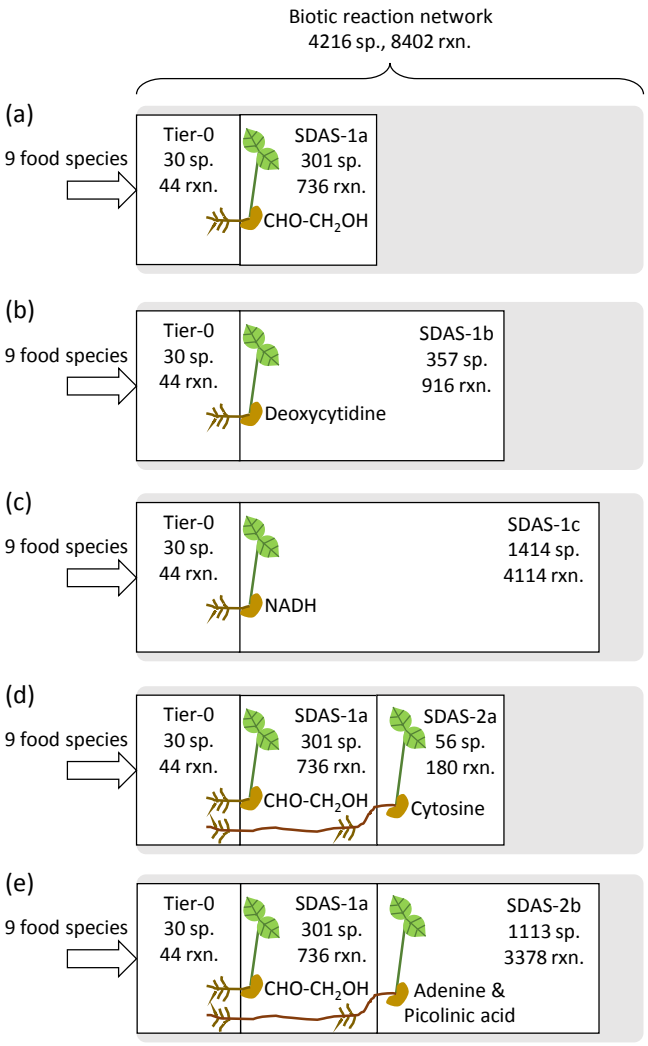

Fig. 4. Summary of sequential seeding of SDASs at different trophic levels in the biotic reaction network. (a) SDAS-1a, which can be seeded by glycolaldehyde ($CHO\text{-}CH_2OH$), feeds on the tier-0 system. **(b)** SDAS-1b, which can be seeded by deoxycytidine, feeds on the tier-0 system. **(c)** SDAS-1c, which can be seeded by NADH, feeds on the tier-0 system. **(d)** Sequential seeding by glycolaldehyde and cytosine results in SDAS-1a and SDAS-2a, a two-tier autocatalytic system that is equivalent to SDAS-1b; SDAS-2a feeds on both the tier-0 system and SDAS-1a. **(e)** Sequential seeding by glycolaldehyde, adenine and picolinic acid results in SDAS-1a and SDAS-2b, which is equivalent to SDAS-1c; SDAS-2b feeds on both the tier-0 system and SDAS-1a. sp.: abbreviation for chemical species. rxn.: abbreviation for reaction.

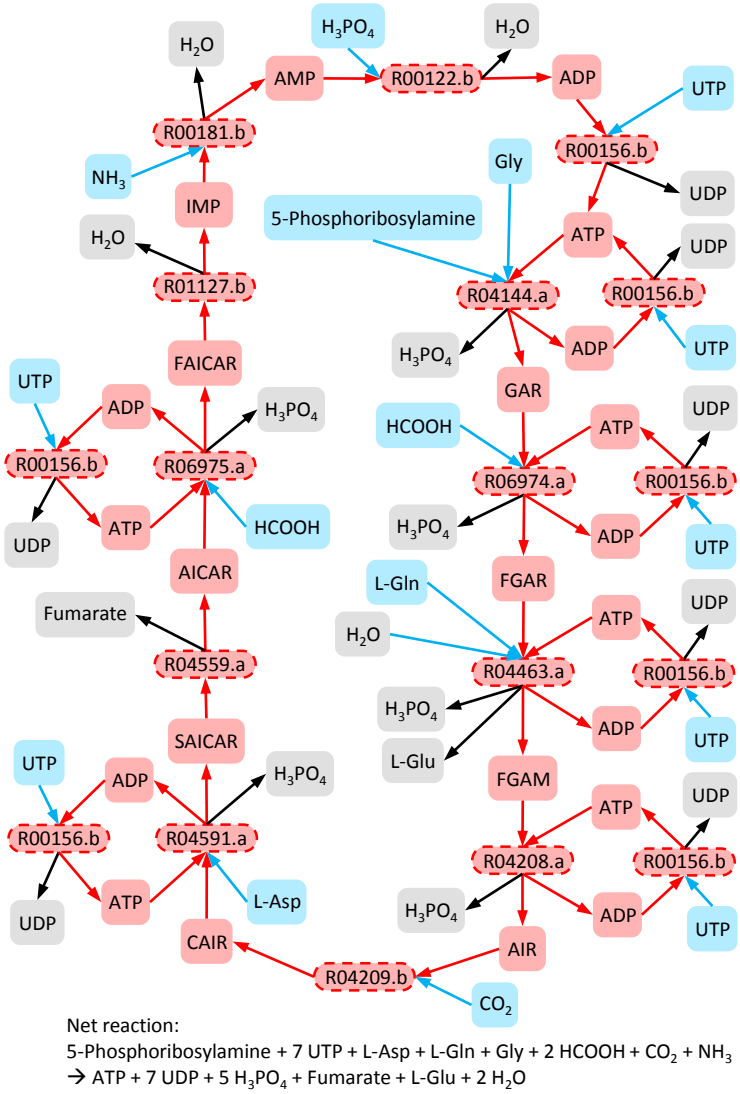

Fig. 5. An autocatalytic cycle synthesizing ATP within SDAS-2b detected in the biotic database. This autocatalytic cycle requires prior establishment of a lower-tier system able to supply chemicals such as glycine and formic acid. Note that reaction R00156.b is used repeatedly in this cycle and that chemicals, such as 5-phosphoribosylamine and UTP, are synthesized from tier-0 and SDAS-1a chemicals by SDAS-2b reactions that are not shown in this graph. Rounded-corner box: chemical species. Red box: member of the autocatalytic cycle. Cyan box: food of the autocatalytic cycle. Grey box: waste of the autocatalytic cycle. Oval with dashed border: reaction. Cyan arrow: food consumption. Black arrow: waste production. Note that some waste of the entire autocatalytic cycle (e.g., H_2O) can be the food for a reaction step (e.g., R04463.a).

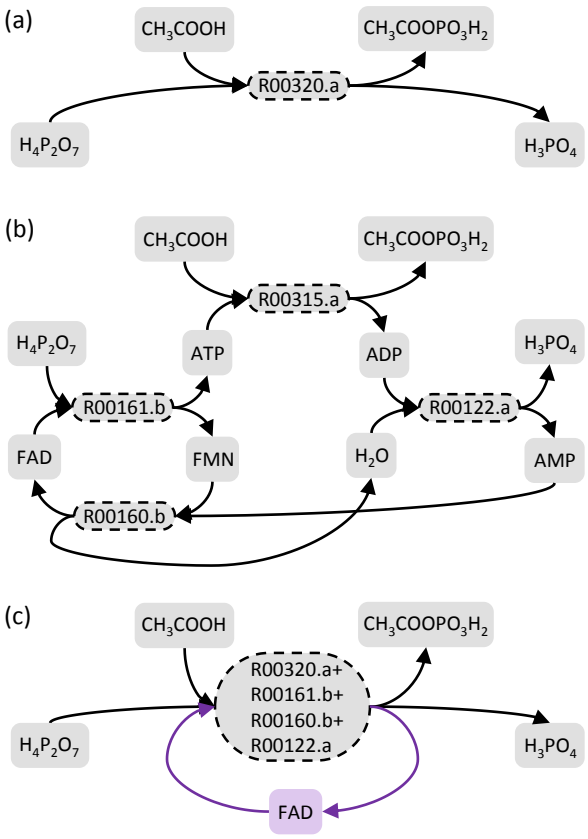

Fig. 6. An example of a higher-tier chemical being a potential catalyst of a lower-tier reaction. FAD, which is produced in biotic SDAS-2b, allows an alternative pathway for the phosphorylation of CH_3COOH, a tier-1 reaction, potentially increasing the rate at which the chemical ecosystem can assimilate phosphate. **(a)** Uncatalyzed phosphorylation of CH_3COOH in tier-1. **(b)(c)** An alternative pathway for the phosphorylation of CH_3COOH, with FAD as a potential catalyst acting via 4 intermediate reactions. Rounded-corner box: chemical species. Oval with dashed border: reaction. Purple box: catalyst. Purple arrow: catalysis.

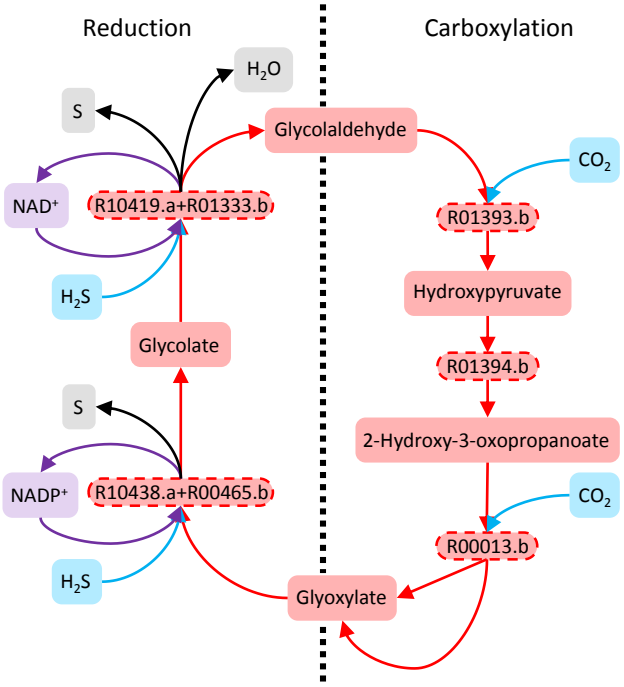

Fig. 7. A carbon-fixing autocatalytic cycle that becomes possible in the biotic SDAS-2b network. This autocatalytic cycle has two major modules: a reduction module leading from glyoxylate to glycolaldehyde, and a carboxylation module leading from glycolaldehyde to more glyoxylate. Because this cycle entails just seven reactions (note that the reduction module has two two-step reactions), it is likely to occur at a higher net rate than the pre-existing carbon fixation mechanisms in SDAS-1a, which require at least 22 reactions (Fig. 3, Table S10). Rounded-corner box: chemical species. Red box: member of the autocatalytic cycle. Cyan box: food of the autocatalytic cycle . Grey box: waste of the autocatalytic cycle. Purple box: catalyst. Oval with dashed border: reaction. Cyan arrow: food consumption. Black arrow: waste production. Purple arrow: catalysis.

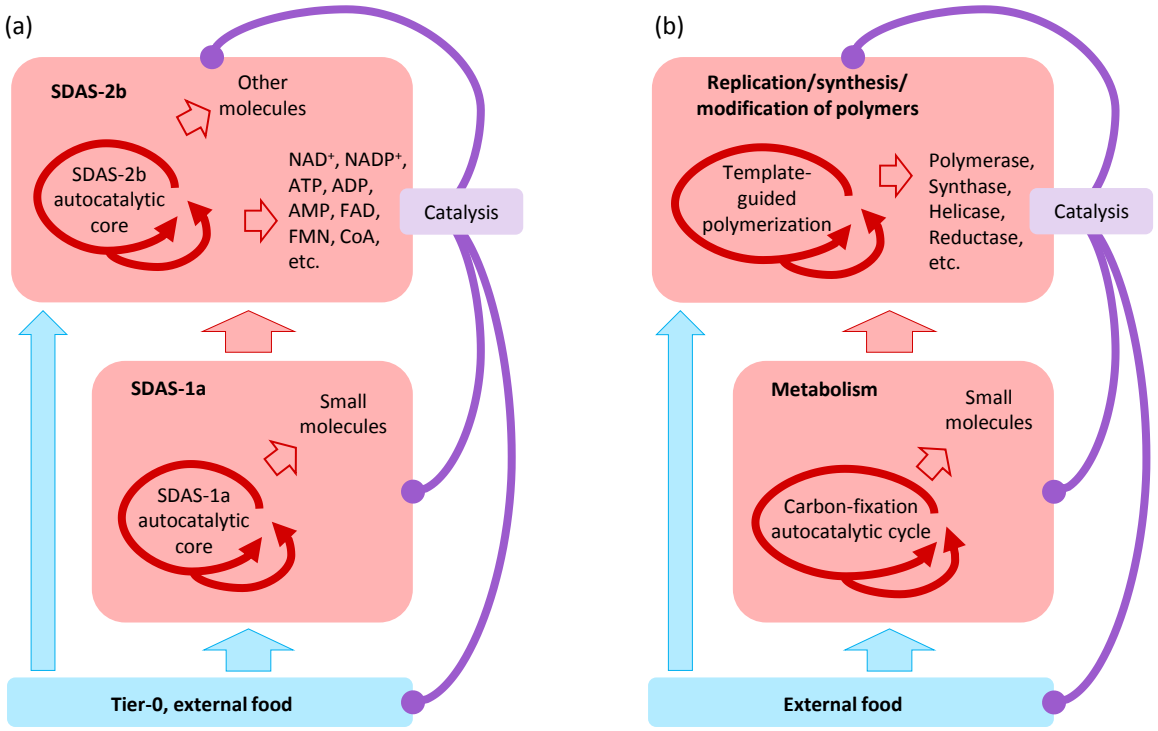

Fig. 8. The interactions between SDASs discovered in the biotic network are topologically similar to those seen in extant life. (a) Summary of the interactions between SDAS-1a and SDAS-2b. **(b)** A simplified summary of the interactions between carbon fixation metabolism, enzymes, and genetic materials in extant life.